Circuit motifs for sensory integration, learning, and the initiation of adaptive behavior in Drosophila

Inaugural-Dissertation zur Erlangung des Doktorgrades der Mathematisch-Naturwissenschaftlichen Fakultät der Universität zu Köln

Vorgelegt von: Anna-Maria Jürgensen (aus Hamburg)
Erster Gutachter: Prof. Dr. Martin Paul Nawrot
Zweiter Gutachter: Prof. Dr. Kei Ito

25. September 2023

Eidesstattliche Versicherung

Hiermit versichere ich an Eides statt, dass ich die vorliegende Dissertation selbstständig und ohne die Benutzung anderer als der angegebenen Hilfsmittel und Literatur angefertigt habe. Alle Stellen, die wörtlich oder sinngemäß aus veröffentlichten und nicht veröffentlichten Werken dem Wortlaut oder dem Sinn nach entnommen wurden, sind als solche kenntlich gemacht. Ich versichere an Eides statt, dass diese Dissertation noch keiner anderen Fakultät oder Universität zur Prüfung vorgelegen hat; dass sie, abgesehen von unten angegebenen Teilpublikationen und eingebundenen Artikeln und Manuskripten, noch nicht veröffentlicht worden ist sowie, dass ich eine Veröffentlichung der Dissertation vor Abschluss der Promotion nicht ohne Genehmigung des Promotionsausschusses vornehmen werde. Die Bestimmungen dieser Ordnung sind mir bekannt. Darüber hinaus erkläre ich hiermit, dass ich die Ordnung zur Sicherung guter wissenschaftlicher Praxis und zum Umgang mit wissenschaftlichem Fehlverhalten der Universität zu Köln gelesen und sie bei der Durchführung der Dissertation zugrundeliegenden Arbeiten und der schriftlich verfassten Dissertation beachtet habe und verpflichte mich hiermit, die dort genannten Vorgaben bei allen wissenschaftlichen Tätigkeiten zu beachten und umzusetzen. Ich versichere, dass die eingereichte elektronische Fassung der eingereichten Druckfassung vollständig entspricht.

Anna-Maria Jürgensen
Köln, den 18.09.2023

Never say never (Justin Bieber)

Abstract

Goal-directed behavior is crucial for survival in complex, dynamic environments. It requires the detection of relevant sensory stimuli and the formation of separable neuronal representations. Learning the contingencies of these sensory stimuli with innately positive or negative valent stimuli (reinforcement) forms associations, allowing the former to cue the latter. This yields cue-based predictions to upgrade the behavioral repertoire from reactive to anticipatory. In this thesis, the Trias of sensory integration, learning of contingencies, and the initiation of anticipatory behavior are studied in the framework of the fruit fly *Drosophila* olfactory pathway and mushroom body, a higher-order brain center for integrating sensory input and coincidence detection using computational network models representing the mushroom body architecture with varying degrees of abstraction. Additionally, simulations of larval locomotion were employed to investigate how the output of the mushroom body relates to behavior and to foster comparability with animal experiments. We showed that inhibitory feedback within the mushroom body produces sparse stimulus representations, increasing the separability of different sensory stimuli. This separability reduced reinforcement generalization in learning experiments through the decreased overlap of stimulus representations. Furthermore, we showed that feedback from the valence-signaling output to the reinforcement-signaling dopaminergic neurons that innervate the mushroom body could explain experimentally observed temporal dynamics of the formation of associations between sensory cues and reinforcement. This supports the hypothesis that dopaminergic neurons encode the difference between predicted and received reinforcement, which in turn drives the learning process. These dopaminergic neurons have also been argued to convey an indirect reinforcement signal in second-order learning experiments. A new sensory cue is paired with an already established one that activates dopaminergic neurons due to its association with the reinforcement. We demonstrated how different pathways for feedforward or feedback input from the mushroom body's intrinsic or output neurons can provide an indirect reinforcement signal to the dopaminergic neurons. Any direct or indirect association of sensory cues with reinforcement yielded a reinforcement expectation, biasing the fly's behavioral response towards the approach or avoidance of the respective sensory cue. We then showed that the simulated locomotory behavior of individual animals in a virtual environment depends on the biasing output of the mushroom body. In conclusion, our results contribute to understanding the implementation of mechanisms for separable stimulus representations, postulated key features of associative learning, and the link between MB output and adaptive behavior in the mushroom body and confirm their explanatory power for animal behavior.

Zusammenfassung

Zielorientiertes Verhalten ist eine überlebenswichtige Fähigkeit in komplexen und dynamischen Umwelten. Notwendig dafür sind die Erfassung und separierbare neuronale Repräsentation sensorischer Reize. Die Kontingenzen zwischen dem Auftreten sensorischer Reize mit Reizen von naturgegebener positiver oder negativer Valenz (Verstärker) wahrzunehmen erlaubt es Assoziationen zwischen ihnen herzustellen. Dadurch kann der sensorische Reiz zum Hinweis auf den zu erwartenden Verstärker werden, was antizipatorisches, statt nur reaktivem Verhalten ermöglicht. Anhand des olfaktorischen Systems und des Pilzkörpers der Fruchtfliege *Drosophila* haben wir die Trias aus sensorischenr Integration, des darauf aufbauenden Lernens von Kontingenzen und der Generierung von Verhaltensimpulsen untersucht. Der Pilzkörper ist eine zentrale Verarbeitungs- und Integrationsstelle multisensorischer Reize und ermöglicht die Erfassung von Koinzidenz die anhand von Netzwerkmodellen untersucht wurden, die die Mechanismen im Pilzkörper auf unterschiedlichen Abstraktionsebenen abbilden. Teilweise wurden sie mit Simulationen des Verhaltens von Drosophila Larven kombiniert, die vom Aktivitätsmuster der Ausgangsneurone des Pilzkörpers beeinflusst werden um Vergleich mit Verhaltensexperimenten an echten Tieren zu erleichtern. Wir haben gezeigt, dass Inhibition auf Netzwerkebene zu einer sparsen Reiz-Repräsentation beitragen, die die Separierbarkeit von Reizen erhöht und den Grad der Generalisierung in Lernexperimenten reduziert. Außerdem konnten wir zeigen, dass Feedback der Valenz kodierenden Ausgangneurone auf dopaminerge Neurone, die die Gegenwart eines Verstärkers enkodieren, die experimentell beobachteten Dynamiken des Erwerbs von Assoziationen erklären kann. Dieses Ergebnis unterstützt die Hypothese, dass die dopaminergen Neurone im Pilzkörper die Differenz zwischen erwarteter und tatsächlicher Verstärkung berechnen und diese Differenz den Lernvorgang motiviert. Diese dopaminergen Neurone sind auch Kandidaten für die Vermittlung einer indirekten Verstärkung in Lernexperimenten höherer Ordnung. Diese zeichnen sich durch die gemeinsame Präsentation eines Verstärkers mit einem sensorischen Reiz aus, der in der Folge selber als Verstärker wirken kann. Wir haben gezeigt, dass verschiedene Mechanismen, vermittelt durch die intrinsischen oder Ausgangneurone des Pilzkörpers geeignet sind diese indirekte Verstärkung an die dopaminergen Neurone zu übertragen. Allgemein erzeugte in unseren Studien jede Assoziation eines Reizes mit einem Verstärker eine Verstärkungserwartung die Verhaltensimpulse in Richtung von Annäherung oder Vermeidung beeinflusste. Zusätzlich haben wir das Verhalten simulierter Fliegenlarven in einer virtuellen Umgebung untersucht und festgestellt, dass die im Pilzkörper generierten Verhaltensimpulse sich direkt auf ihr tatsächliches Annäherungs- und Vermeidungsverhalten auswirken. Unsere Ergebnisse tragen dazu bei, die Implementierung der postulierten Mechanismen für separierbare Reiz-Repräsentation, assoziatives Lernen und den Übergang zwischen den Impulsen aus dem Pilzkörper und tatsächlichem adaptivem Verhalten zu untersuchen und zeigt deren Fähigkeit Verhalten in Tierexperimenten zu erklären.

List of Figures

Contents

1 Introduction

Animals navigate life in complex, dynamic environments that are rich in sensory input. Their survival depends on their ability to protect themselves from potential threats and seek out advantageous circumstances. To be successful in these quests, they need to process their sensory environment to identify potential sources of harm or benefit. The innate valence of a stimulus to an animal depends on its potential to provide a benefit or inflict harm. This can result in a behavioral preference to either approach or avoid these stimuli. Preference, alongside motivational internal states (like satiety or fear), the physiological capacity for behavior, and potential environmental obstacles, is one of the determining factors of animal behavior, effectively limiting the scope of feasible actions in any given situation.

Through experiences in their respective environments, animals can acquire additional information about the world, expanding their knowledge beyond the scope of innate valences. Learning processes allow them to update their representation and interpretation of the sensory environment, modifying valences or assigning them to previously neutral sensory stimuli if they experience interactions with them as helpful or harmful. In many cases, the more knowledge an animal acquires, the better its chances are of investing resources well into behavior facilitating approach or avoidance. In learning the relationships between different stimuli, they can also acquire the ability to predict future environmental changes. These relationships can be the occurrence of a stimulus usually preceded by another stimulus or a particular behavior of their own. Figure 1 provides a conceptual overview of the different elements covered in this thesis, as well as components that are not covered in any of the chapters but are essential parts of the framework of stimulus processing and learning for the initiation of goal-directed behavior as it will be presented and discussed here. The missing elements (innate valence and internal states) are covered in the discussion.
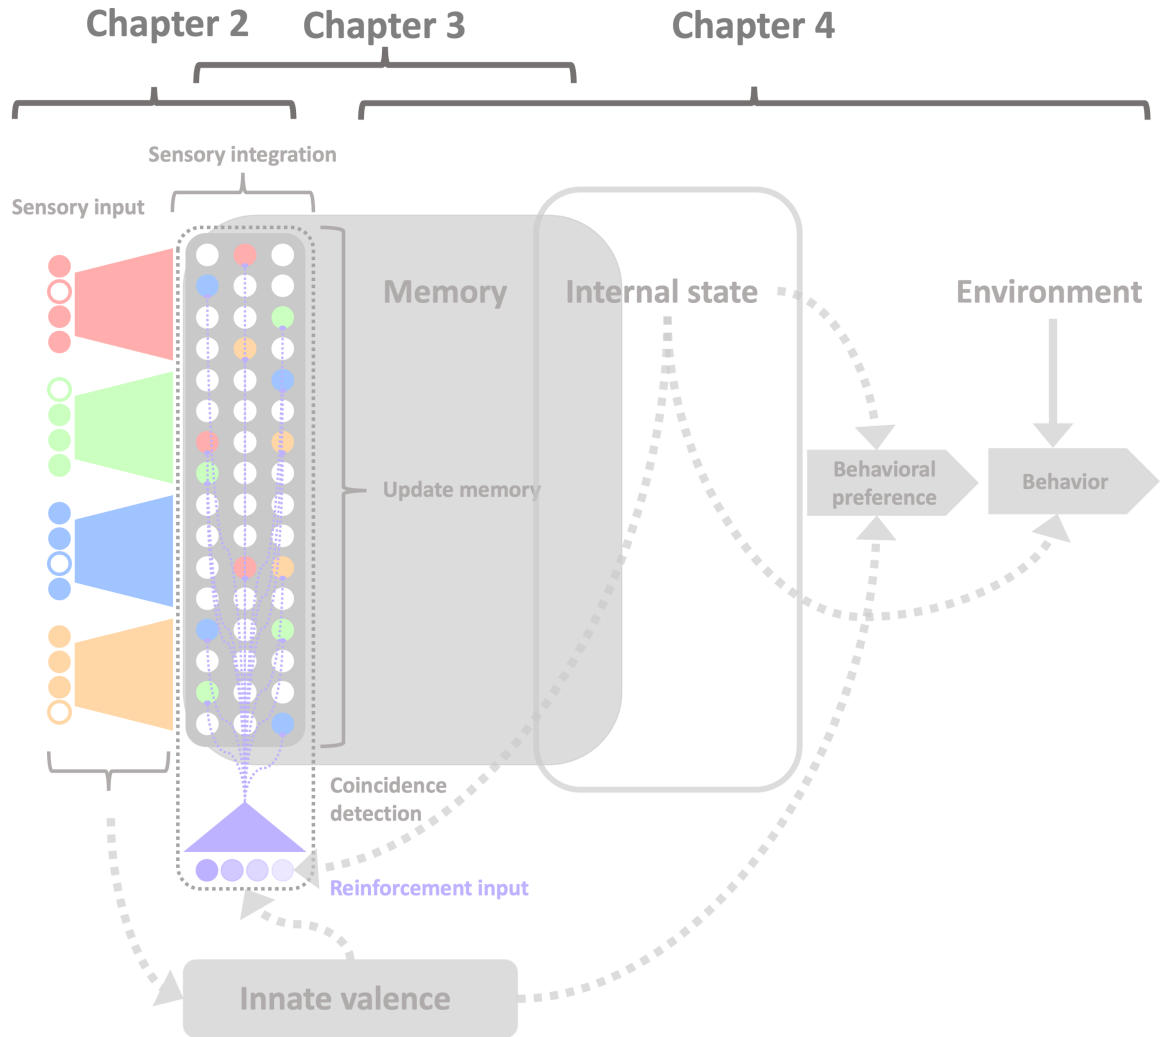
Figure 1: **The schematic relationship between sensory integration, learning, and the initiation of behavior.** Dense multisensory input is integrated into higher-order processing centers, where coincidences among them and with reinforcement are detected, and associations are formed between concurrently active sensory inputs, internal states, and reactivated prior knowledge (learned or innate). The initiation of behavior depends on the integrated output of the higher-order processing center, signaling learned valence, compared against behavioral preference based on potential innate valences. The execution of a behavioral preference can be influenced by an animal's abilities in the respective situation and by environmental factors.

1.1   Expanding knowledge through learning

Learning about relationships between stimuli is a prerequisite to effective goal-directed behavior in a dynamic world. New information needs to be acquired, stored, and retrieved at the appropriate time. The study of these mechanisms has a long scientific tradition. In 1885 Hermann Ebbinghaus was the first to publish a highly influential work on the study of memory from a non-philosophical point of view [1]. He described memory as a construct to store *…Mental states of every kind, -sensations, feelings, ideas, - which were at one time present in consciousness…*([1], page 1). He emphasized that these elements no longer persist in the conscious mind but are still stored somewhere from where they can be reproduced voluntarily or involuntarily. When integrating these concepts, he anticipated the conceptual framework of encoding, storage, and retrieval, which is still used today to refer to the major building blocks of learning and memory. These components are of interest regardless of the scientific discipline and focus on either a theoretical and conceptual [2] or mechanistic and functional perspective [3]. Across disciplines, the word memory is generally used to refer to the storage site for encoded information. From a conceptual perspective, this means retaining encoded relationships between cognitive elements. These cognitive elements can represent everything relating to a previous experience [2]. On the other hand, biology and neuroscience are concerned with the exact location of memory functions inside the brains of their varying model organisms. Memory is studied from the perspective of mechanisms in the synaptic and other interactions of neurons on a circuit, cellular, or sub-cellular level [3].

Situationally appropriate retrieval of stored information from memory allows access to knowledge about the environment that can guide decisions about behavior. Retrieval depends on the relatedness of all cues available in the retrieval situation to what has been encoded in memory (for a review, see[4]). The authors describe this process as reawakening specific memory elements through sensory cues. Memory retrieval can be based on a voluntary query or search of the stored information or happen involuntarily. Either mechanism is triggered and guided by the present, related cues [4].

In this framework, updating existing memory with acquired sensory information is referred to as learning. Two fundamental forms of learning can be differentiated: Associative and non-associative learning. In non-associative learning, the response to a single stimulus is altered after repeated or prolonged stimulation [5, 6], without any relationship being formed with other stimuli. Sensory habituation and sensitization constitute the two major forms of non-associative learning that elicit opposite responses to a stimulus [5, 6]. On the other hand, associative forms of learning rely on encoding relationships between two or more stimuli based on the contingencies of their occurrence [7, 8]. Two forms of associative learning are generally studied in different experimental paradigms: Operant and classical conditioning.

During operant conditioning, an association between a behavior (or the outcome of that behavior) and a consequence is formed. As a result, the probability of the rewarded behavior changes [9]. During classical conditioning, as originally introduced by Ivan Pavlow [10], a naturally valent stimulus is presented together with a neutral stimulus. The valent stimulus is often referred to as the unconditioned stimulus because it naturally causes a reaction (unconditioned response). Without prior learning, the neutral stimulus (conditioned stimulus) does not cause any specific reaction. In Pavlow's original work, a dog would hear a bell (conditioned stimulus) right before food is presented (unconditioned stimulus), which fuels saliva production (unconditioned re-

sponse). After learning the association between the two stimuli (bell and food), the dog would respond with a conditioned saliva response when hearing the bell. The conditioned response can be viewed as a preparation for the expected unconditioned stimulus and may, therefore, differ from the unconditioned response [10]. Classical conditioning relies on coincidence detection between two or more sensory stimuli, through which one becomes a meaningful predictor of the other.

1.2 Efficient sensory coding to support learning

Coincidence detection for the formation of associations has been argued to take place in higher-order processing centers (Figure 1) that receive sensory input of different modalities such as the mushroom body (MB) in insects [11–14] or the cerebellum in vertebrates [15–17]. Processing of sensory input is thus a prerequisite of coincidence detection and associative learning. Sensory input is relayed across several processing stages. A physical or chemical environmental stimulus interacts with a receptor at the earliest stage. A physical or chemical transformation of the receptors marks stimulus detection. It elicits a neuronal signal that is processed across various layers of relay steps towards higher-order processing centers [18, 19]. The processing of the sensory code throughout these layers can contribute to preparing the sensory representations for processing in higher-order brain areas [18, 19], for example, by integrating inputs from receptors of the same type [18, 19] and increasing efficiency and specificity of the code [18, 19]. Both the neuronal coding space and the energy available to an organism are limited resources. Through a transformation of dense code at the receptor periphery into sparse representations in higher-order processing centers, sensory systems achieve energy efficient and reliable stimulus encoding [20–22]

Any processing layer in the brain consists of a limited number of neurons, determining the upper bound for the dimensionality that can be encoded. Population sparseness refers to the representation of a stimulus across a population of neurons, such that only a few neurons are activated by any given stimulus (Figure 1), and different stimuli activate distinct sets of neurons [21]. A population-sparse stimulus representation minimizes redundancy and overlap of stimulus encodings [22–24]. Temporal sparseness is another relevant component of sparse coding. It is achieved when an individual neuron responds with only a few spikes to a specific stimulus [22, 25, 26]. When an animal interacts with changes in the sensory input in a dynamic environment, temporal sparseness allows it to encode these fluctuations efficiently [27–30].

Sparse code has been studied both in invertebrate [11, 31–34] and vertebrate species [28, 35–38]. It can increase the separability of encoded stimuli by decreasing redundancy and overlap [31, 39, 40], and supports memory formation [39–41].

1.3 The Drosophila olfactory pathway and the mushroom body as a model to study sensory coding, learning, and adaptive behavior

The *Drosophila* olfactory pathway, combined with the MB, is a popular system for studying mechanisms of sensory processing and associative learning and the initiation of behavior. The MB is a higher-order center for the integration of sensory input of different modalities in larval [42, 43], and adult [44–46] *Drosophila*. It has also been shown to be an important site for associative learning (larva: [42, 47, 48], adult: [49–

52]) and its output is involved in biasing behavior (larva:[47, 48], adult:[51, 53–56]).

For decades, the fruit fly *Drosophila* has been extensively used as a model organism in neuroscience [57]. With its short lifecycle and cheap stock management *Drosophila* is a convenient tool in labs worldwide. Over time, many genetic mutants have been developed to study specific brain circuits and functions. Additionally, the availability of tools for direct manipulation of neurons through light or temperature [58–60], together with many established experimental paradigms in both larval and adult *Drosophila* [57, 61, 62] provide researchers with unique and powerful opportunities to design learning experiments [57]. In recent years, more in-depth knowledge of the connectome of both the larval [42, 43, 47, 48, 63] and adult [44, 49, 64, 65] sensory processing and learning centers has emerged, adding to the appeal of this model organism.

1.3.1   The architecture of the Drosophila olfactory pathway and the mushroom body

The general structure and functionality of the olfactory pathway and the MB are retained between larval and adult *Drosophila* [61, 66]. Olfactory processing begins with odorants in the air surrounding the animal that are detected by receptors and encoded by the olfactory receptor neurons (larva:[61, 67, 68], adult:[69–71]). They project to the antennal lobe, which is organized in glomeruli, neuropils in which the receptor neurons synapse onto projection neurons and local interneurons. The local interneurons innervate projection neurons within the antennal lobe (larva: [61, 72], adult: [70, 73]), eliciting an inhibitory effect (larva: [61, 72], adult: [70, 74]). The projection neurons relay their input onto the dendrites of MB intrinsic neurons, the Kenyon cells (KC) (larva:[43], adult:[75, 76]).

Synapses between projection neurons and KCs are located in the calyx region of the MB (larva:[43], adult:[75, 76]). Here, a small number of glomeruli diverges onto a larger number of KCs (from ≈ 21 to ≈ 72 in the larva [43] and from ≈ 50 to ≈ 2000 in the adult [46, 77, 78]). Olfactory representations in the *Drosophila* MB are population-sparse with few KCs being activated by any given odor [78, 79]. The combined activity of a set of KCs encodes the odor identity [46, 77, 78]. Their parallel axons then form the MB lobes, where they synapse onto the MB output neurons (MBONs) (larva:[43, 80], adult:[44, 49, 75, 80]). The adult MB is divided into five lobes [44, 49], which have been credited with different roles in learning and memory retrieval [75, 81, 82]. The larval MB consists of only two lobes [83–85]. For an overview of the development of the lobes in the larva, see [84, 85]).

1.3.2   Associative learning in the Drosophila mushroom body

The MB lobes are further separated into compartments, which are defined by the MBONs that innervate them and form synapses with the KCs (larva: [43, 66], adult: [44, 49, 65, 87]). Modulatory neurons target these synapses in a compartment-specific manner (larva: [43, 66], adult: [44, 49, 65, 87, 88]). The modulatory neurons can employ dopamine (DANs) or octopamine as their neurotransmitter. Depending on the compartment, experimental protocol, and longevity of the established memory, modulatory neurons of either type modulate learning in the *Drosophila* MB. Dopamine has been shown to be involved in reward [89, 90] and punishment learning in adult[89, 91–93] and larval [42] *Drosophila*. In contrast, octopamine is involved in reward learning in both the larva [94] and the adult [90, 92]. DANs have been studied extensively in
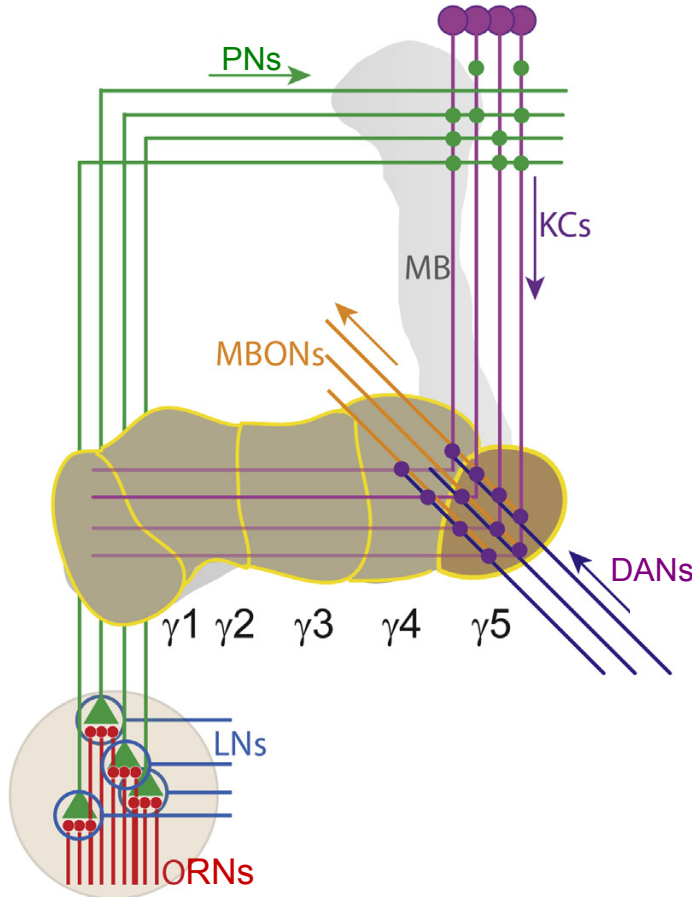
Figure 2: **Schematic of the olfactory pathway and the mushroom body in Drosophila.** Olfactory receptor neurons (ORNs) encode receptor activation that is relayed onto projection neurons (PNs) in the antennal lobe glomeruli, innervated by local interneurons (LN). The projection neurons convey this information to the mushroom body (MB) calyx, where they synapse onto Kenyon cells (KCs). Dopaminergic neurons (DANs), conveying reinforcement signals, innervate distinct subcompartments of the MB (here within the γ-lobe). The KCs project to MB output neurons (MBONs) that innervate areas downstream of the MB. Adapted from [86].

the context of learning. Their activity encodes the presence of rewards and punishments in the environment, not only in *Drosophila* (larva: [42], adult: [89, 91, 95]), but also in vertebrate [96–99] species. Within the MB compartments, modulation of the synapses between KCs and MBONs is facilitated by the activity of the respective DANs, innervating the compartment whenever a sensory stimulus and the activity of the modulatory neuron coincide (larva: [42, 47], adult: [50, 51, 93, 100]).

DANs have been suggested to target the synapses between the KCs and MBONs in two ways. In the larva, DANs synapse onto the neurons on both the pre and post-synaptic side [43]. Based on electron microscopy data obtained from the synaptic zone, it has been suggested that DANs can cause the release of neurotransmitters in both MB intrinsic and output neurons [43]. Additionally, they seem to provide reinforcement signals through neuromodulatory volume release near the pre-synaptic terminals in the KCs [43]. In the adult *Drosophila* MB, the presence of a specific type of dopamine receptor in the KCs is required to form olfactory memories [101] and DANs have been demonstrated to be pre-synaptic in connections with both KCs and MBONs in the α lobe [65]. Despite a relatively small amount of synapses between these DANs and KCs and MBONs [65]. The activity of the DANs strongly affects plasticity at the synapses between KCs and MBONs [50, 101, 102], suggesting the additional relevance of volume release [44, 65], as observed in the larva [43].

Plasticity at the synapses between the KCs and MBONs induces a shift in the activity of MB output neurons when the animal is confronted with the respective sensory cue after learning (larva: [48], adult: [50–53]). MBONs have been shown to project to pre-motor areas (larva: [42, 48], adult: [44, 49]), and learning-induced

plasticity in the MB has been suggested to modify the activity of neurons downstream of the MBONs in the adult *Drosophila* [53, 56]. When the output of an MB compartment that demonstrates modified activity due to learning is blocked, learned behavior can no longer be observed [52, 103].

1.4 Organization of this thesis

Learning and memory are a complex and multifaceted collection of several interactive processes that encompass acquisition, maintenance, updating, extinction, and sometimes even the recovery of previously extinguished information to allow for goal-directed behavior based on stored information about the environment. Several processes are required to orchestrate the transformation of input from the external world into usable memories with behavioral relevance. Sensory processing allows input to enter the brain, where it needs to be represented and coupled with simultaneously active or incoming information (Figure 1). This temporary blend of internal and external stimuli needs to be evaluated for its anchoring points within the network of the already existing memory to verify its relevance and determine the contexts within which it might later be retrieved to guide behavior (Figure 1). The ability to choose the most beneficial (anticipatory) behavior heavily depends on the situationally appropriate retrieval of information from memory, potentially compared against behavioral tendencies originating from innate valences. The execution of an emerging behavioral tendency can then be influenced by the animal's abilities in the respective situation and by environmental factors (Figure 1).

The emerging Trias of sensory representation, learning, and anticipatory, instead of purely reactive behavior, is a strong formula for survival. In this thesis, I will present four individual works that investigate the different elements of this Trias in three chapters (Figure 1):

Chapter 2: Efficient sensory coding

> 2.1 Jürgensen, A-M., Khalili, A., Chicca, E., Indiveri, G., & Nawrot, M. P. (2021). A neuromorphic model of olfactory processing and sparse coding in the Drosophila larva brain. Neuromorphic Computing and Engineering, 1(2), 024008.

Chapter 3: Circuit mechanisms of associative learning

> 3.1 Jürgensen, A-M., Sakagiannis, P., Schleyer, M., Gerber, B., & Nawrot, M. P. (2022). Prediction error drives associative olfactory learning and conditioned behavior in a spiking model of Drosophila larva. bioRxiv, 2022-12.

> 3.2 Jürgensen, A-M., Schmitt, Felix J.,& Nawrot, M. P. (2023). Minimal circuit motifs for second-order conditioning in the insect mushroom body. bioRxiv, 2023-09.

Chapter 4: Adaptive behavior

> 4.1 Sakagiannis, P., Jürgensen, A-M., & Nawrot, M. P. (2021). A realistic locomotory model of Drosophila larva for behavioral simulations. bioRxiv, 2021-07.

In the second chapter of this thesis (Chapter 2.1), we investigated three mechanisms for increasing population- and temporal sparseness in odor encoding in the *Drosophila* larva MB. By using a spiking model of the olfactory pathway and the MB and an implementation of the same network on neuromorphic hardware, we demonstrated

how sparse odor code is achieved efficiently and that it increases the separability of stimuli in this system, thereby laying the grounds for the following works on learning [104, 105].

In the third chapter, we studied mechanisms of associative learning in the larval (Chapter 3.1) and adult (Chapter 3.2) *Drosophila* MB. We investigated how animals can learn to predict upcoming reinforcement from the occurrence of sensory cues in a classical conditioning paradigm using a spiking model of the MB circuit (Chapter 3.1). In the same chapter, we next explored minimal circuits for second-order conditioning (Chapter 3.2), a paradigm in which a cue, established as a reinforcement predictor through classical conditioning, can serve as indirect reinforcement for learning associations with other sensory cues. Additionally, we introduced a tool to simulate the effects of learned preferences on larval behavior, which allows us to test our ideas about MB circuits for reinforcement prediction by comparing our simulation results to those of behavioral animal experiments (Chapter 3.1).

Finally, the fourth chapter covers our locomotory model of the larva, used in 2.1 in more depth (Chapter 4.1). It models individual larvae, controlled by architecture, that rely on intermittent forward crawling and lateral bending at its lowest level. Active sensing adds a second dimension to locomotory control by allowing navigation along odor gradients. When odor preferences acquired by associative learning were added, these preferences modified the behavior of simulated larvae.

In presenting the individual works, integrating, and discussing them, I will move around between the circuit mechanisms of learning and their theoretical and conceptual underpinnings firstly because both are equally interesting to me and have always inspired my perspective on research. Secondly, I believe that the conceptual perspective can help to remember that there is a bigger picture beyond the details and missing pieces of knowledge about a mechanism in a particular model organism.

2   Chapter: Efficient sensory coding

2.1   A neuromorphic model of olfactory processing and sparse coding in the Drosophila larva brain

NEUROMORPHIC
Computing and Engineering

A neuromorphic model of olfactory processing and sparse coding in the Drosophila larva brain

Anna-Maria Jürgensen[1,6] , Afshin Khalili[1,2,6] , Elisabetta Chicca[3,4] ,
Giacomo Indiveri[5] and Martin Paul Nawrot[1,*]

1   Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Cologne, Germany
2   Department Genetics of Learning and Memory, Leibniz Institute for Neurobiology, Magdeburg, Germany
3   Bio-Inspired Circuits and Systems Lab, Zernike Institute for Advanced Materials, University of Groningen, Groningen, The Netherlands
4   Groningen Cognitive Systems and Materials Center, University of Groningen, Groningen, The Netherlands
5   Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland
*   Authors to whom any correspondence should be addressed.
6   These authors contributed equally

E-mail: mnawrot@uni-koeln.de

Abstract

Animal nervous systems are highly efficient in processing sensory input. The neuromorphic computing paradigm aims at the hardware implementation of neural network computations to support novel solutions for building brain-inspired computing systems. Here, we take inspiration from sensory processing in the nervous system of the fruit fly larva. With its strongly limited computational resources of <200 neurons and <1.000 synapses the larval olfactory pathway employs fundamental computations to transform broadly tuned receptor input at the periphery into an energy efficient sparse code in the central brain. We show how this approach allows us to achieve sparse coding and increased separability of stimulus patterns in a spiking neural network, validated with both software simulation and hardware emulation on mixed-signal real-time neuromorphic hardware. We verify that feedback inhibition is the central motif to support sparseness in the spatial domain, across the neuron population, while the combination of spike frequency adaptation and feedback inhibition determines sparseness in the temporal domain. Our experiments demonstrate that such small, biologically realistic neural networks, efficiently implemented on neuromorphic hardware, can achieve parallel processing and efficient encoding of sensory input at full temporal resolution.

1. Introduction

Neuromorphic computing [1] is a novel paradigm that aims at emulating the naturalistic, flexible structure of animal brains on an analogous physical substrate with the potential to outperform von Neumann architectures in a range of real-world tasks [2, 3]. It can inspire novel AI solutions [4–6] and may support control of autonomous agents by spiking neural networks [7–9]. A major challenge for brain-inspired neuromorphic solutions is the identification of computational principles and circuit motifs in animal nervous systems that can be utilized on neuromorphic hardware to exploit its benefits.

Drawing inspiration from neural computation in the nervous systems of insects is particularly promising for developing neuromorphic computing paradigms. With their comparatively small brains ranging from ≈10 000 neurons in the fruit fly larva to ≈1 million neurons in the honeybee, insects are able to solve many formidable tasks such as the efficient recognition of relevant objects in a complex environment [10, 11], perceptual decision making [12–14], or the exploration of unknown terrain and navigation [15–19]. They also show simple cognitive abilities such as learning, or counting of objects [20–24]. At the same time, their compact nervous systems are optimized for energy efficient computation with limited numbers of neurons and synapses, making them ideally suited to meet current neuromorphic hardware limitations regarding network size and topology. Spiking neural networks modeled after the insect brain have been shown to support efficient

sensory processing [25], learning [7, 26], foraging and navigation [27–29], and counting [28]. Model studies also include earlier neuromorphic implementations of insect-inspired computation [4, 5, 9, 30–33].

Sparse coding [34, 35] is a fundamental principle of sensory processing, both in invertebrates [36–40] and vertebrates [41–45]. By transforming dense stimulus encoding at the receptor periphery into sparse representations in central brain areas, the sensory systems of animals achieve energy efficient and reliable stimulus encoding [35, 46], which increases separability of items [47–50]. Sparse coding in neural systems has two major components [39]. *Population sparseness* refers to the representation of a stimulus across the entire population of neurons, such that only few neurons are activated by any specific stimulus and different stimuli activate largely distinct sets of neurons. Re-coding from a dense peripheral input to a sparse code in central brain areas supports stimulus discriminability and associative memory formation by projecting stimulus features into a higher dimensional space [51–53]. *Temporal sparseness* indicates that an individual neuron responds with only a few spikes to a specific stimulus configuration [34, 54, 55] supporting the encoding of dynamic changes in the sensory environment [42, 56] and memory recall in dynamic input scenarios [28].

We are interested in the transformation of a densely coded input into a sparse representation within an olfactory pathway model of the *Drosophila* larva. As a common feature across insect species, odor information is processed across multiple network stages to generate a reliable sparse code of odor identity in the mushroom body (MB) [36, 57, 58], a central brain structure serving as a hub for multi-sensory integration, memory formation and memory recall [10, 59]. A shared characteristic of the *Drosophila* larva brain and the here-used real-time neuromorphic hardware system is their relatively small network size. With this limited capacity, computational efficiency and frugal use of the limited resources are a major constraint. Implementing evolutionary-derived mechanisms from the insect brain that allow for sparse, thus more efficient stimulus encoding on the chip could help to broaden the scope of its applications. In our network model we test the efficiency of cellular mechanisms and network motifs in producing population and temporal sparseness and test their implementation on the mixed-signal neuromorphic hardware DYNAP-SE [60] in comparison to a software simulation using the Python-based spiking neural network simulator 'Brian2' [61].
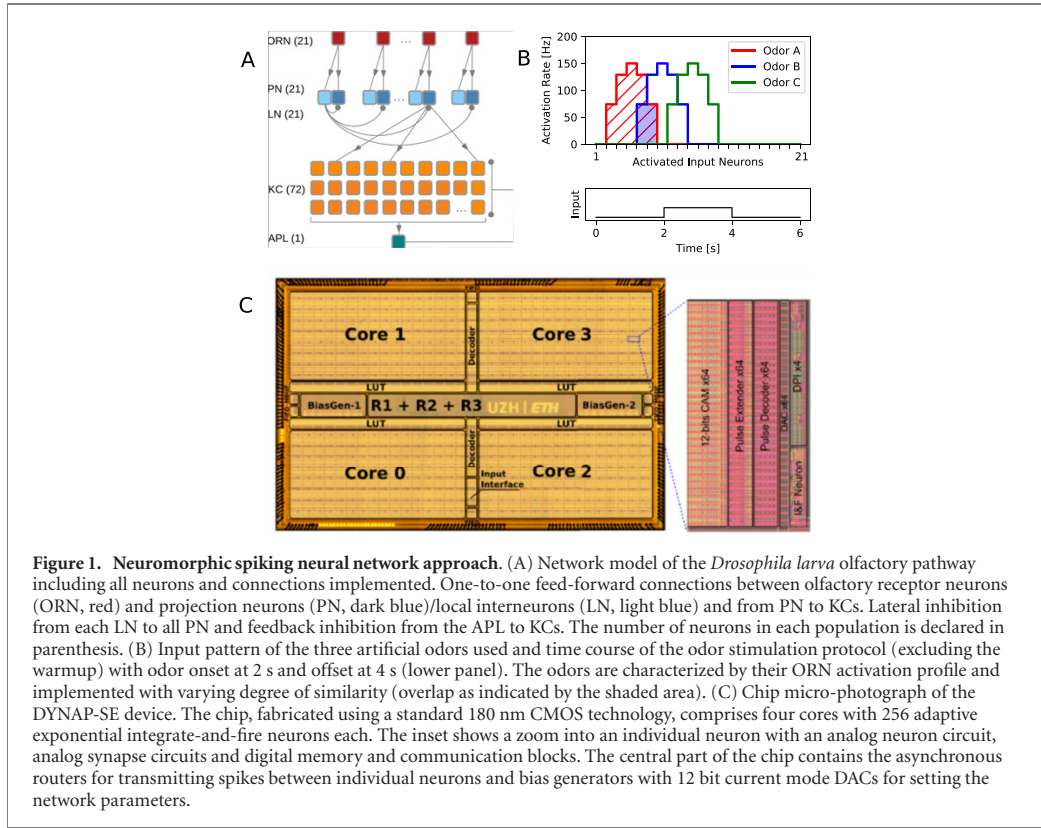
2. Methods

2.1. Spiking neural network model

The architecture of the spiking neural network model as shown in figure 1(A) uses the exact numbers of neurons in each population and the reconstructed connectivity for one hemisphere as published in the electron-microscopic study of a single animal [62, 63]. The network consists of 21 olfactory receptor neurons (ORN) at the periphery, 21 projection neurons and 21 local interneurons (LNs) in the antennal lobe and 72 Kenyon cells (KCs). In each brain hemisphere there is exactly one anterior paired lateral (APL) neuron. We hypothesize that the APL receives input from most or all mature KCs [64] included in this network model. Due to technical limitations of the DYNAP-SE chip with a maximum in-degree of 64 synapses for one neuron we randomly chose 64 KCs that provide input to the APL. This choice was fixed for the model, both on the hardware network and in the software simulation. We further hypothesize, based on evidence in the adult species, that all ORNs and all KCs have a mechanism of cellular spike frequency adaptation (SFA).

2.2. Implementation on the DYNAP-SE neuromorphic hardware

The olfactory pathway model of the *Drosophila* larva was implemented using the dynamic neuromorphic asynchronous processor (DYNAP-SE) [60] (figure 1(C)). This processor is a full-custom mixed-signal analog/digital VLSI chip, which comprises analog circuits that emulate neurons and synapses with biologically plausible neural dynamics. Given the analog nature of the circuits used, the synapses and neurons exhibit parameter variability that is characteristic also of real neurons. The analog circuits used, implement multiple aspects of neural dynamics, such as spike-frequency adaptation (implemented as a shunting inhibitory synapse), refractory periods, exponentially decaying currents, voltage-gated excitation and shunting inhibition [60, 65]. The silicon neurons circuits, similar to their biological counterparts, produce spikes. In the chip, these are stereotyped digital events which are routed to target synapses by a dedicated address event representation (AER) infrastructure [66, 67]. The conductance-based synapses are current-mode circuits [65] that produce an EPSC with biologically plausible dynamics, which are then injected into the neurons *leak* compartment. This compartment acts as a conductance block, which decreases the input current as the membrane potential increases. One of the inhibitory synapses subtracts charge directly from the membrane capacitance and provides a shunting inhibition mechanism [65]. All other synaptic currents are in turn summed together and integrated in the post-synaptic neurons leak compartment.

The model (figure 1(A)) was initially developed in software and the neural architecture was then mapped onto the mixed-signal hardware by configuring the AER routers and programming the chip digital memories to connect the silicon neurons via their corresponding synapses. The parameters of the hardware setup were

Figure 1. Neuromorphic spiking neural network approach. (A) Network model of the *Drosophila larva* olfactory pathway including all neurons and connections implemented. One-to-one feed-forward connections between olfactory receptor neurons (ORN, red) and projection neurons (PN, dark blue)/local interneurons (LN, light blue) and from PN to KCs. Lateral inhibition from each LN to all PN and feedback inhibition from the APL to KCs. The number of neurons in each population is declared in parenthesis. (B) Input pattern of the three artificial odors used and time course of the odor stimulation protocol (excluding the warmup) with odor onset at 2 s and offset at 4 s (lower panel). The odors are characterized by their ORN activation profile and implemented with varying degree of similarity (overlap as indicated by the shaded area). (C) Chip micro-photograph of the DYNAP-SE device. The chip, fabricated using a standard 180 nm CMOS technology, comprises four cores with 256 adaptive exponential integrate-and-fire neurons each. The inset shows a zoom into an individual neuron with an analog neuron circuit, analog synapse circuits and digital memory and communication blocks. The central part of the chip contains the asynchronous routers for transmitting spikes between individual neurons and bias generators with 12 bit current mode DACs for setting the network parameters.

fine-tuned using the on-chip bias generator, starting from the estimates provided by the software simulation. Computer-generated control stimuli, in the form of well defined spike trains, were provided to the chip via a custom field programmable gate array (FPGA) board. Each neuron population was implemented on a single core, using in total five cores and two chips. All the circuit biases of neurons belonging to different cores could be tuned independently. The synapses from ORN to PN, PN to KC, and KC to APL were designed as excitatory whereas the synapse from LN to PN and APL to KC were implemented as inhibitory. SFA was implemented in the ORN and KC neuron populations.

Three separate recordings were done, one for each of the three odors with 20 trial repetitions (figure 1(B)). Within each of the three experiments all conditions (different sparseness mechanisms enabled) were recorded always in the same order (LN + APL + SFA, LN + SFA, LN, APL + SFA, LN + APL, SFA).

2.3. Computer simulation of the spiking neural network

The simulations were implemented in the network simulator Brian2 [61] and run on an X86 architecture on an Ubuntu 16.04.2 Server. All neurons (figure 1(A)) were modeled as leaky integrate-and-fire neurons with conductance-based synapses. The membrane potential v_i obeys a fire-and-reset rule, being reset to the resting potential whenever the spike threshold is reached. The reset is followed by an absolute refractory period of 2 ms, during which the neuron does not integrate inputs (table 1). The membrane potential of a neuron in a particular neuron population (v^O, v^L, v^P, v^K, v^A) is governed by the respective equation. The neuron parameters can be found in table 1.

$$C_m \frac{\mathrm{d}}{\mathrm{d}t} v_i^O = g_L(E_L - v_i^O) + g_e^{\mathrm{InputO}}(E_E - v_i^O) - g_{Ia}(E_{Ia} - v_i^O) \tag{1}$$

$$C_m \frac{\mathrm{d}}{\mathrm{d}t} v_i^L = g_L(E_L - v_i^L) + g_e^{\mathrm{OL}}(E_E - v_i^L) \tag{2}$$

$$C_m \frac{\mathrm{d}}{\mathrm{d}t} v_i^P = g_L(E_L - v_i^P) + g_e^{\mathrm{OP}}(E_E - v_i^P) - g_i^{\mathrm{LP}}(E_I - v_i^P) \tag{3}$$

$$C_m \frac{\mathrm{d}}{\mathrm{d}t} v_i^K = g_L(E_L - v_i^K) - g_i^{\mathrm{APLK}}(E_I - v_i^K)$$
$$+ g_e^{\mathrm{PK}}(E_E - v_i^K) - g_{Ia}(E_{Ia} - v_i^K) \tag{4}$$

Table 1. Network simulation parameters.

Neuron parameters		
Capacitance ORN	C_m	100 pF
Capacitance PN	C_m	30 pF
Capacitance LN	C_m	50 pF
Capacitance APL	C_m	200 pF
Leak conductance ORN	g_L	5 nS
Leak conductance PN and LN	g_L	2.5 nS
Leak conductance APL	g_L	5 nS
Leak potential ORN	E_L	60 mV
Leak potential PN and LN	E_L	−60 mV
Leak potential KC and APL	E_L	−60 mV
Threshold potential ORN and KC	V_T	−35 mV
Threshold potential PN and LN	V_T	−30 mV
Threshold potential APL	V_T	−30 mV
Resting potential ORN and LN	V_r	−60 mV
Resting potential PN	V_r	−60 mV
Resting potential KC	V_r	−55 mV
Resting potential APL	V_r	−60 mV
Refractory time	τ_{ref}	2 ms
Synaptic parameters		
Excitatory potential	E_E	0 mV
Inhibitory potential	E_I	−75 mV
Excitatory time constant	τ_e	5 ms
Inhibitory time constant	τ_i	10 ms
Synaptic weights		
Weight input-ORN	wORNinputORN	3 nS
Weight ORN–PN	wORNPN	30 nS
Weight ORN–LN	wORNLN	9 nS
Weight LN–PN	wLNPN	2 nS
Weight PN–KC	wPNKC	1 nS
Weight KC–APL	wKCAPL	50 nS
Weight APL–KC	wAPLKC	100 nS
Adaptation parameters		
Adaptation time constant	τ_{Ia}	1000 ms
Adaptation reversal potential	E_{Ia}	−90 mV

$$C_m \frac{\mathrm{d}}{\mathrm{d}t} v_i^A = g_L(E_L - v_i^A) + g_e^{KAPL}(E_E - v_i^A) \tag{5}$$

where ORNs (equation (1)) and KCs (equation (4)) are equipped with an additional spike-triggered adaptation (equation (6)) where g_{Ia} is the adaptation conductance and τ_{Ia} is the decay time constant. With every spike g_{Ia} is increased in ORNs and KCs by 0.1 nS and 0.05 nS, respectively.

$$\frac{\mathrm{d}}{\mathrm{d}t} g_{Ia} = -\frac{g_{Ia}}{\tau_{Ia}}. \tag{6}$$

Note, that the neuron model used in our computer simulations is the widely used conductance based leaky integrate-and-fire neuron [68] with an additional adaptation conductance in ORNs and KCs. This model does not match perfectly well the silicon neuron physically implemented on the DYNAP-SE board, which can be modeled by a current-based adaptive exponential integrate-and-fire model [65] (see Discussion). All code for the software implementation is accessible via https://github.com/nawrotlab/DrosophilaOlfactorySparseCoding
.

2.4. Spontaneous activity

The input to the ORNs in our network model was modeled as stochastic point process realizations. It mimics the sum of spontaneous receptor activation and odor driven activation of the ORNs. On the chip, each ORN received a Poisson input to achieve a baseline firing rate ≈5 Hz. In the simulation each ORN received excitatory synaptic input modeled as a gamma process (shape parameter $k = 3$) to generate a similar baseline rate. The spontaneous firing rate of larval ORNs was previously measured in the range of 0.2−7.9 Hz, depending strongly on receptor type and odor identity [69, 70]. On the chip we measured a spontaneous ORN firing

rate of 6.2 ± 3.0 Hz. In the simulated model the average ORN baseline activity was estimated as 6.0 ± 1.4 Hz. Thus, ORNs on chip and in the simulation exhibit a similar spontaneous activity in the upper range of the empirical distribution.

2.5. Odor stimulation protocol

On the chip and in the computer simulation we included a warm up time (1.5 s and 0.3 s, respectively), which was excluded from the analyses. On the chip this restored the baseline biases following odor application. In the computer simulation this warm up period ensured that neuronal membranes and conductances were more heterogeneous at the beginning of the experiments.

We used a set of three different odors to study the effect of odor similarity. Figure 1(B) shows the activation profile (point process intensities) and overlap of all three odors across the 21 input channels. For each odor, the profile indicates the ORN-type specific activation level, mimicking the fact that each ORN expresses a genetically different receptor type. Similarity of odors is represented in the overlapping activation where odor 1 and odor 3 are distant (zero overlap), while odor 2 is constructed to have the same amount of overlap with the two other odors. The stimulation protocol assumes a 2 s odor stimulus on top of the baseline input with an activation rate according to figure 1(B).

2.6. Data analysis

2.6.1. Sparseness measure
Sparseness was quantified by the widely used modified version [71] of the Treves–Rolls measure [72].

$$S = 1 - \frac{\left(\frac{1}{N}\sum_{i=1}^{N} a_i\right)^2}{\frac{1}{N}\sum_{i=1}^{N} a_i^2} \tag{7}$$

where a_i indicates either the spike count of neuron i (population sparseness, S_{pop}), or the binned ($\Delta t = 20$ ms) population spike count (temporal sparseness, S_{tmp}) for the 2 s with odor stimulation. S assumes values between zero and one, with high values indicating sparse responses. This measure has been repeatedly used to quantify sparseness in insect olfactory processing [36, 47, 52, 54, 73–77]. We report the average and standard deviation across the three odors. We then tested the effect of excluding specific sparseness mechanisms. To test for significance of the effects of lateral inhibition and SFA, the condition with only lateral inhibition enabled was compared with the condition with only SFA present (LN vs SFA) using a t-test for related samples. To test the effect of feedback inhibition via the APL, the condition including all mechanisms (LN + APL + SFA) was compared with LN + SFA. Tests were performed independently for temporal and population sparseness.

2.6.2. Activation measure
We define the additional measure of activation as

$$A = \frac{1}{N \cdot k} \sum_{i=1}^{N} \sum_{j=1}^{k} \Theta(a_{ij}) \tag{8}$$

where $a_{i,j}$ indicates the spike count of neuron i in the time bin j and Θ is the Heaviside step function. Thus, $\Theta(a_{ij})$ indicates the binary response of neuron i in time bin j. To asses population activation A_{pop} we apply a single time bin for the complete 2 s odor stimulation time. Then A_{pop} measures the fraction across all N neurons that are odor-activated by at least one single spike. We quantify temporal activation A_{tmp} by binning the stimulus time into $k = 20$ bins of $w = 100$ ms. Thus, A_{tmp} measures the binary response probability across time bins for each neuron. Our definition of activation is related to the complementary measure of 'activity sparseness' defined in [71]. Both measures were then averaged over all 20 trials. As results we report the average and standard deviation across the three odors.

2.6.3. Distance measure
To assess the differences in odor distance between sparse and dense KC odor code we used the cosine distance (equation (9)). Vectors a and b each represent the average number of spikes evoked by all 72 KCs during the two second odor presentation across 20 independent model instances. Cosine distance between a and b was calculated as:

$$D_{\cos} = 1 - \frac{\sum_{i=1}^{n} a_i \cdot b_i}{\sqrt{\sum_{i=1}^{n} a_i^2} \cdot \sqrt{\sum_{i=1}^{n} b_i^2}}. \tag{9}$$

2.6.4. Correlation across sparseness conditions
To test for qualitatively similar effects of the different sparseness conditions on the chip and simulation we correlated the results across the six data points (sparseness conditions) between the chip and the simulation.

For significance testing we generated 100 random unique permutations of the means from the simulation and correlated these 100 data series with that of the chip (LN + APL + SFA,LN + APL, APL + SFA, LN + SFA, LN, SFA; figure 3). The average of these 100 correlations was $0.07 (sd = 0.42)$ for S_{pop} and $0.01 (sd = 0.52)$ for S_{tmp}. In both cases the distribution of correlations was normal, established using D'Agostino-Pearson test for normality. The average of these 100 correlations each was used to evaluate the similarity of the effects on the chip and in the simulation.

3. Results

The larval nervous system with its limited neural network size and low complexity lends itself to the emulation on neuromorphic hardware. We analyzed a single hemisphere olfactory network model of the first instar *Drosophila* larva with <200 neurons and <1000 synapses comparing an implementation on the neuromorphic hardware DYNAP-SE [60] with a computer simulation of the same network. We were particularly interested in the contribution of different cellular and circuit mechanisms to the transformation of a dense input pattern at the periphery into a sparse odor representation in the MB.

3.1. Olfactory pathway model

Our spiking neural network model comprises four computational layers (figure 1(A)). Its structure, the size of the neuron populations and their connectivity are based on the connectome of a single right hemisphere as reconstructed from electron-microscopic data of one individual *Drosophila larva* MB by Eichler and colleagues [62]. Peripheral processing is carried out by 21 ORNs, each expressing a different olfactory receptor type [63, 78]. ORNs make one-to-one excitatory connections with 21 PNs and with 21 LNs that together constitute the antennal lobe. Each LN forms inhibitory synapses onto all PNs, establishing lateral inhibition. The PNs make divergent random connections with a total of 72 KCs, the primary cells of the MB, where each KC receives excitatory input from 1–6 PNs. The APL receives input from all of the matured KCs [64]. All KCs with a well-developed dendrite [62] fall into this category and those are the only ones included in our circuit model. We therefore assume a dense convergent connectivity with essentially all presynaptic KCs (in our case 64 out of 72 due to technical limitations on the chip, see Methods). We further implemented inhibitory feedback from the APL onto all KCs [64]. Overall, this blueprint of the olfactory network is highly similar to that in the adult fly except for the smaller neuron numbers and reduced anatomical complexity (see Discussion). Each model instance implemented here utilizes the exact same connectivities. We thus simulate a single individual rather than an average animal.
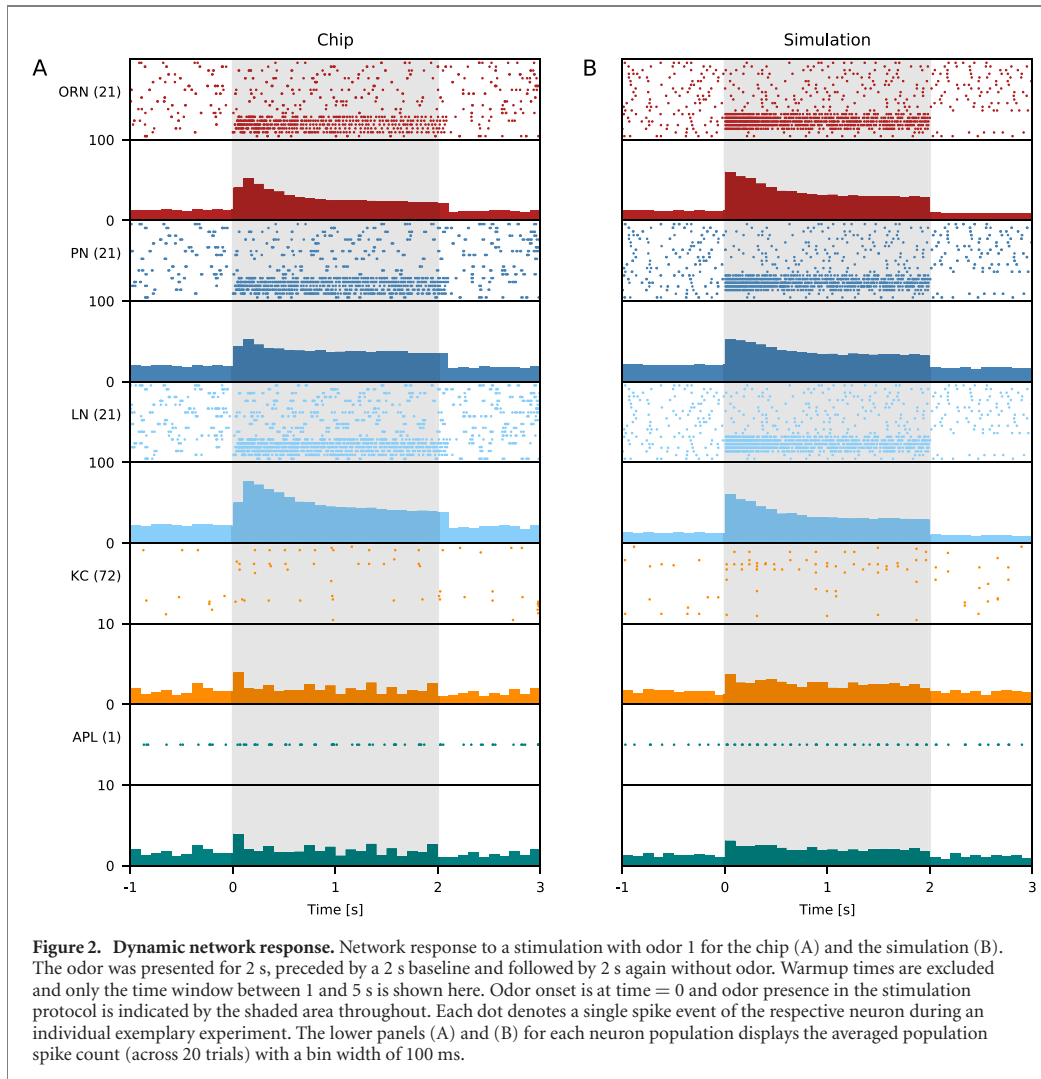
3.2. Circuit motifs and cellular adaptation

Our network model utilizes different cellular and circuit mechanisms that have been suggested to support a sparse code in the insect MB. To this end, the network topology includes three relevant motifs. First, the LN connectivity in the antennal lobe constitutes lateral inhibition as a motif that generally enhances neural contrast [34, 79] and that is implemented in the olfactory system of virtually all insects [36, 80–85], as well as in computational models thereof [25, 28, 52, 86]. Second, the random connectivity from PNs to a larger number of KCs is net divergent and sparse, expanding the dimensionality of the coding space [51, 87, 88]. Third, our model includes inhibitory feedback from the APL neuron onto all KCs. This has been shown to directly affect KC populations sparseness in the adult fly [47, 89] (see Discussion).

At the cellular level, all neurons in the network are modeled as leaky integrate-and-fire neurons. ORNs and KCs are equipped with a cellular SFA mechanism, a fundamental and ubiquitous mechanism in spiking neurons [34, 90]. ORNs have been shown to adapt during ongoing stimulation *in vivo*, both in larval [91] and adult [92, 93] *Drosophila*. The exact nature of the adaptation mechanism in the ORNs is still under investigation [92, 94, 95]. In KCs, a strong SFA conductance has conclusively been demonstrated in the cockroach [96] and the bee [97].

3.3. Dynamics of network response to odor stimulation

The response dynamics across all network stages to a single constant odor stimulation (figure 1(B)) with odor 1 is shown in figure 2(A) (chip) and figure 2(B) (simulation). At stimulus onset, a subset of all ORNs is activated according to the corresponding receptor response profile (figure 1(B), top). The ORN responses are phasic-tonic as a result of SFA with a higher firing rate at odor onset. The spike count histogram averaged across the 21 neurons of the ORN population fits the typical experimentally observed response profile observed in adult *Drosophila* [74, 92]. In the larva, little is known about stimulus adaptation in the ORNs [70]. The physical realization of SFA on chip is different from the simulation, which may partly explain the delayed response to odor onset- and offset of some neurons and the initially slower increase of the phasic response on the chip (figure 2(A), see Discussion). The off-response expressed in a prolonged silence of the odor-activated ORNs
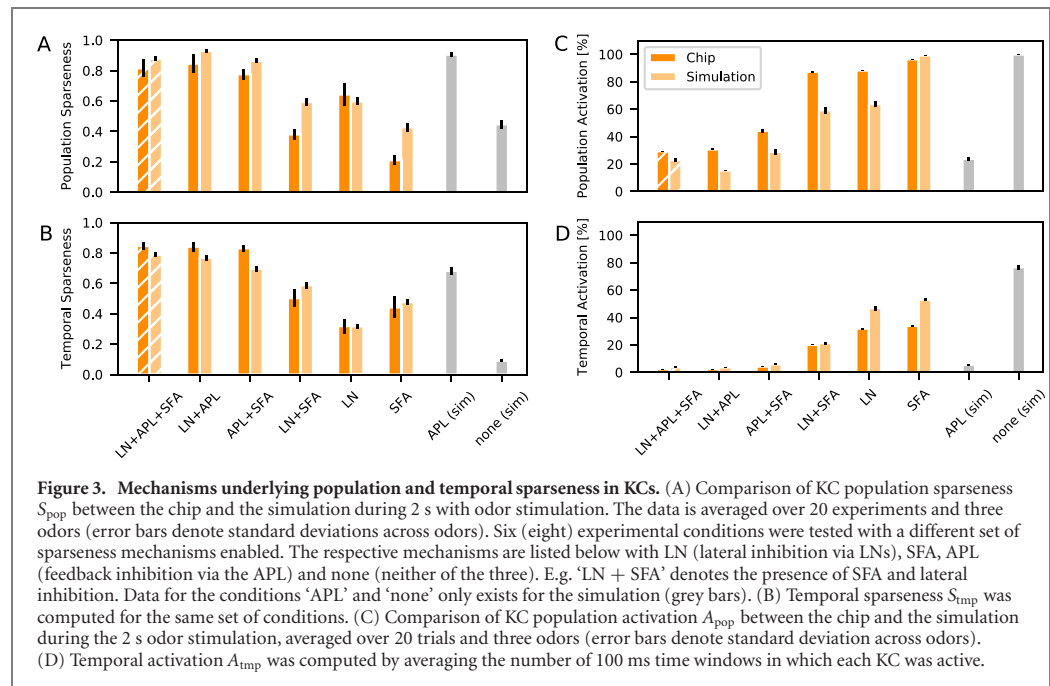
Figure 2. Dynamic network response. Network response to a stimulation with odor 1 for the chip (A) and the simulation (B). The odor was presented for 2 s, preceded by a 2 s baseline and followed by 2 s again without odor. Warmup times are excluded and only the time window between 1 and 5 s is shown here. Odor onset is at time = 0 and odor presence in the stimulation protocol is indicated by the shaded area throughout. Each dot denotes a single spike event of the respective neuron during an individual exemplary experiment. The lower panels (A) and (B) for each neuron population displays the averaged population spike count (across 20 trials) with a bin width of 100 ms.

in the simulation is an effect of SFA: the integrated adaptation current that has reached a steady state during the odor stimulation period now decays only slowly, acting in a hyperpolarizing fashion and thus reducing spiking probability [52] of the ORNs. This effect is barely visible and delayed on the chip (see Discussion). At the level of the antennal lobe both PNs (dark blue) and LNs (light blue) are excited only by the ORNs and thus follow their phasic-tonic response behavior and exhibit an inhibited off-response (figure 2), although neither neuron type is adaptive itself. The spatio-temporal response pattern of the PNs and LNs resembles the typical response pattern measured *in vivo* in adult flies and bees [81, 98, 99], including an inhibitory off-response in many neurons [92, 100, 101].

The KCs show very little spiking during spontaneous activity on the chip and in simulation. Only very few KCs do respond to odor stimulation (population sparse response) with only a single or few spikes (temporally sparse response). Spontaneous activity and response properties match well the *in vivo* situation as observed in various species [36, 54, 58]. The population spike count indicates a very brief population response within the first 100 ms, while the tonic KC response remains only slightly above the spontaneous activity level (cf [54]). Finally, the single APL driven by the excitatory KC population follows the brief phasic and weak tonic response of the KCs.

3.4. Analysis of sparsening factors in space and time

We investigate the translation from the peripheral dense code in the ORN and PN population into a central sparse code in the KC population, disentangling the contribution of the three fundamental biological mechanisms: cellular adaptation (SFA), lateral inhibition in the AL, and feedback inhibition in the MB. We

Figure 3. Mechanisms underlying population and temporal sparseness in KCs. (A) Comparison of KC population sparseness S_{pop} between the chip and the simulation during 2 s with odor stimulation. The data is averaged over 20 experiments and three odors (error bars denote standard deviations across odors). Six (eight) experimental conditions were tested with a different set of sparseness mechanisms enabled. The respective mechanisms are listed below with LN (lateral inhibition via LNs), SFA, APL (feedback inhibition via the APL) and none (neither of the three). E.g. 'LN + SFA' denotes the presence of SFA and lateral inhibition. Data for the conditions 'APL' and 'none' only exists for the simulation (grey bars). (B) Temporal sparseness S_{tmp} was computed for the same set of conditions. (C) Comparison of KC population activation A_{pop} between the chip and the simulation during the 2 s odor stimulation, averaged over 20 trials and three odors (error bars denote standard deviation across odors). (D) Temporal activation A_{tmp} was computed by averaging the number of 100 ms time windows in which each KC was active.

systematically varied the composition of the three mechanisms in our network, yielding five different conditions (figure 3) in which either one or two mechanisms were deactivated. SFA was only deactivated at the KC level and still present in ORNs. We did not vary the PN–KC connectivity pattern as this is identical to the anatomical pattern reported for the individual animal that we used as a reference.

We quantified the population activation by measuring the fraction of stimulus-activated KCs across the different conditions (see Methods) and find that it depends on the sparseness mechanisms. It is lowest in the control condition with 20.6(28.6%) responding neurons on the chip and 16.7(22.9%) in the simulation (figure 3(C)). Our results show that APL is the single crucial mechanism necessary for establishing a high population sparseness in our model. All conditions that lack feedback inhibition show strongly reduced values of S_{pop}. Lateral inhibition can to some degree recover sparseness on the chip and in the simulation.

We now consider temporal sparseness, which again reached high values in the control condition on the order of $S_{tmp} \approx 0.8$ (hatched bars in figure 3(B)). Comparing the different conditions we find that APL feedback inhibition and SFA in the KCs have a strong supporting effect for temporal sparseness. Any condition that involves the APL reached similar high values for S_{tmp}. Without the APL, SFA can partially ensure temporal sparseness on the chip and in the simulation. This is also reflected in the temporal activation measure that computes the fraction of active time bins (of 100 ms duration) for the complete 2 s stimulation time (see Methods). The results shown in figure 3(D) mirror our results in figure 3(B). In the sparse control condition KCs are active in on average only 2.3% and 3.4% of the response bins for the chip and the simulation, respectively.

Overall, we observed the same mechanistic effects on chip and in the simulation for the different combinations of activated and inactivated mechanisms (figure 3). The pattern of sparseness values across all six conditions is highly and significantly correlated between the chip and the simulation results, both for $S_{pop}(r = 0.94)$ and $S_{tmp}(r = 0.96)$ in comparison with the correlation of randomly permuted pattern of sparseness values (see Methods) with maximum correlations of 0.91 and 0.87 for S_{pop} and S_{tmp} across 100 permutations, respectively.

3.5. Sparse representation supports stimulus separation

How does the encoding of different odors at the KC level compare between the sparse control condition and a non-sparse condition? Feedback and lateral inhibition supported population sparseness in the KC population. We thus compared the control condition to the network in which both inhibitory mechanisms were disabled by quantifying the pairwise distance between KC stimulus response patterns for any two different odors. Figure 4 shows the response rates averaged over the 2 s stimulus duration for the three different stimuli for both chip (figures 4(A)–(C)) and simulation figures 4(D)–(F). Only a fraction of the KCs responded to any odor ($S_{pop} > 0.8$, figures 3(A) and (E)) in the sparse condition. However, when feedback and lateral inhibition are disabled, essentially all KCs showed an odor response to any of the three odors (figures 4(C) and (F)).

Figure 4. Response pattern overlap. Average spike frequency (over 20 trials) for Chip ((A), (B) and (C)) and simulation ((D), (E) and (F)) in response to three different odors. The odors were presented for 2 s (for information on the experimental protocol please refer to figure 1(B)). All panels display the overlap between the different odor representations either on PN level ((A) and (D)), KC ((B) and (E)) and non-population-sparse KC in a condition with only SFA enabled ((C) and (F)). Overlap indicates a low ability to differentiate between odors.

A similar result is obtained when looking at cosine distances between KC odor representations. Independently of the odor identities, average pairwise cosine distance was considerably larger in the sparse condition (chip: 0.39(0.2); simulation: 0.85(0.06)) than in the non-sparse (SFA only) condition (chip: 0.07(0.02); simulation: 0.31(0.09)), indicating a similar effect of population sparseness on odor discriminability on the chip and in the simulation.

4. Discussion

In the present manuscript we addressed two major questions. First, we asked whether the re-coding from a dense peripheral olfactory code into a sparse central brain representation of odors can be achieved in the small spiking neural network model of *Drosophila* larva. To this end we tested the relevance of three fundamental mechanisms in establishing population and temporal sparseness:

- cellular adaptation
- lateral inhibition
- feedback inhibition

Second, we explored the feasibility of applying this coding scheme on real-time analogue neuromorphic hardware by comparing hardware implementation with software simulation at the relevant levels of stimulus encoding and processing.

4.1. Neuromorphic implementation versus computer simulation

Our results show that the on-chip network implementation achieved the transformation from dense to sparse coding in space and time. We obtained the same general results on the chip and in simulation albeit small differences. What are possible factors contributing to these differences?

First, while the software simulation used identical parameters for all neurons and synapses in a given population, there is considerable heterogeneity across the physical hardware implementation due to device mismatch, which particularly affects currents and conductances [4, 102, 103]. This heterogeneity is manifest e.g. in spiking thresholds, postsynaptic current amplitudes and membrane time constants. The neuromorphic hardware heterogeneity generally matches the biological heterogeneity that is typically ignored in computer-based simulations.

Second, setting the neuron and synapse parameters is straight forward and exact in the computer simulation. On the chip, however, this requires the adjustment of various biasing currents. As a result, real parameters will differ from theoretical target parameters and across circuits, as well as after re-adjustment in the same circuit.

Third, we have used different neuron models in the hardware emulation and in the computer simulation. Thus there is no one-to-one correspondence of the biophysical neuron parameters in the software (table 1) and the set points of the electronic circuits. In an effort to validate the robustness of the architecture to the model details, we deliberately did not minimize this difference, for example by employing hardware-matching neuron models designed to mimic the electronic circuit of the DYNAP-SE (https://code.ini.uzh.ch/yigit/dynapse-simulator.git).

Our research goal in this study was to assess the robustness of function in a small neural network architecture that is supported by three specific cellular and circuit mechanisms. To this end we tested its implementation on the DYNAP-SE neuromorphic hardware in light and despite of the various differences between the exact computer simulation of homogeneous elements and the real-time processing on electronic hardware with inhomogeneous devices. Taking this perspective, the differences between hardware and software implementation strengthen the conclusion that the suggested mechanisms are robust in supporting population sparse and temporally sparse stimulus encoding.

There are a number of advantages and disadvantages in using the specific hardware solution tested here. The fact that the DYNAP-SE [60] operates in real-time makes it suitable for the spiking control of autonomous robots [104, 105] and renders computational speed independent of network size. Even for the small larval network considered here (exactly 136 neurons and 833 synapses) simulations were several times slower than real-time with 3.8 s simulation time per 1 s biological time at a resolution of 0.1 ms (single core CPU, 64 bit PC, Ubuntu 18.04.5). Simulation time can be sped up to meet real-time demand even for large network sizes on specialized systems [106, 107].

A challenge with the mixed-signal neuromorphic hardware was the sensitivity of the circuit bias currents to noise and temperature changes, and the real-time nature of the experiment emulation. As each experiment would require the real time evolution of the input patterns and of the network dynamics to produce its response, this led to complex and lengthy experiments. As there were different experimental conditions, with three different odor stimuli, each with 20 trials, a particular challenge was the time-consuming adjustment of the SFA time constant on the chip since it required post hoc estimation of the effective time constant during repeated spike recordings. We therefore made the *a priori* choice to restrict the hardware emulation experiments to only six out of eight experimental conditions (figure 3). Two more conditions in which none of the mechanisms or only the feedback inhibition via APL was active were tested in the computer simulation only (grey bars in figure 3). Still, the variability across model instances was only slightly larger on the chip than in the simulation (figure 3). In addition, new neuromorphic circuit designs will be able to compensate for these drifts by using appropriate temperature compensated bias generator circuits [108].

4.2. Mechanisms and function of population sparseness

Population sparseness at the KC level has been demonstrated for a number of species in the adult stage (see Introduction). Our model suggests that given the current knowledge of anatomical structures within the *Drosophila* larva olfactory pathway it might already be implemented at this stage with similar benefits. Different mechanisms have been suggested for the generation of population sparseness. A fundamental anatomical basis for a sparse code is the sparse and divergent connectivity between PNs and a much larger population of KCs [37, 52]. Each KC receives input from only a few PNs and thus establishes a projection from a lower into a higher dimensional space, ideally suited to generate distinct activity patterns that foster associative memory formation. Additionally, there is evidence for a low excitability of the individual KCs that require collective input from several PNs to be activated [36, 37, 96]. Connectivity in our model is based on the exact numbers from electron microscopic reconstructions of neurons and synapses in the right hemisphere of one individual brain [62]. We did not attempt to adjust excitability of KCs or PN::KC connection strength for optimal population sparseness.

Feedback inhibition has repeatedly been suggested to underlie population sparseness in several animals, including the fly larva [109]. Empirical evidence has been provided in particular in bees [100, 110] and adult flies [47, 89]. Several modeling studies have used feedback inhibition to support a sparse KC population code in larger adult KC populations [25, 27, 28, 111]. Indeed, our study shows that inhibitory feedback from the single APL neuron effectively implements a sparse code in the small population of 72 KCs (figure 3(A)). We chose to model the APL as a spiking neuron that receives input solely from KCs and inhibits KCs in a closed loop. This decision was based on experimental evidence indicating a clear polarity of the APL with input in the MB lobes and pre-synaptic densities in the calyx, presumably onto the KC dendrites [64]. Whether the APL neuron generates sodium action potentials, however, is not clear in the larva [109] and has been challenged in the adult [47]. In addition, inhibitory feedback connections within the MB have been implicated in learning through inhibitory plasticity in bees and flies, thereby modulating the sparse KC population code [47, 110, 111].

As a third factor, lateral inhibition within the *Drosophila* antennal lobe has been shown to increase population sparseness at the KC level [74, 112] and in a model thereof [52]. This model study showed a strong effect of lateral inhibition on population sparseness in a network tuned to the anatomy of the adult fly. In the present larval model we found a supportive effect. With lateral inhibition alone the model reached $S_{pop} \approx 0.6$. The interplay of feedback inhibition and lateral inhibition boosted population sparseness to $S_{pop} \approx 0.8$ (figure 3(A)). This observation is different from our previous results in a network simulation modeled after the adult fly [28] where lateral inhibition in the AL was sufficient to implement a high population sparseness and APL feedback inhibition had a mainly supporting effect. The fact that lateral inhibition is less effective in the larval than in the adult Drosophila model [52] is thus likely due to the one-to-one connectivity between the 21 ORNs and 21 PNs in the larva, which requires very strong excitatory synapses. This specific configuration establishes a dominant feed-forward component in the larval olfactory pathway (figure 1(A)).

Sparse stimulus representation across the neuronal population supports minimal overlap of and correlation across stimulus-specific spatial response patterns [34, 52, 89, 113, 114], which in turn benefits associative memory formation and increases memory capacity [47, 53, 72, 115]. We confirmed an increased inter-stimulus distance in the KC coding space on the chip and in the simulation when all sparseness mechanisms take effect.

4.3. Mechanisms and function of temporal sparseness

Temporal sparseness in the insect MB has been physiologically described in various species. It is expressed in a highly phasic stimulus response that typically consists of only a single or very few spikes and that is temporally locked to stimulus onset or to a fast transient increase in stimulus amplitude while the tonic stimulus response is almost absent [36, 54, 58]. In our model we implemented two mechanisms that can support temporal sparseness, inhibitory feedback via the spiking APL neuron and SFA. Our analysis in figure 3(B) showed that inhibitory feedback has the strongest effect, confirming experimental [77, 116] and modeling results [25, 27, 28, 117]. Cellular adaptation (SFA) showed a smaller but supporting effect in our network, which is partially in line with our previous models of the adult fly [28, 52, 77] in which we showed that SFA alone can suffice to generate high temporal sparseness.

Importantly, cellular adaptation has additional effects on stimulus coding that are not analyzed here. Being a self-inhibiting mechanism it reduces overall spiking activity, contributing to the low spontaneous and response rates in the KC population that has been repeatedly documented in various insect species [36, 54, 58]. Moreover, SFA leads to a regularization of the neuron's spike output and a reduction of the trial-to-trial variability, effectively improving response reliability [77, 118]. Finally, SFA introduces a short-term stimulus memory expressed in the conductance state of the excited neuron population, which decays with the SFA time constant [52].

Temporal sparseness was influenced strongly by SFA in the KCs and by recurrent feedback inhibition. It usually shows as longer inter-spike-intervals both in physiological data [36, 54, 100] as well as in modeling results [28, 52, 77]. Besides the prolongation of the inter-spike-intervals over the entire duration of the experiment, SFA also caused the commonly observed odor onset effect [36, 54, 58, 100, 116] in ORNs and KCs. In our data this effect was somewhat concealed in the KCs by the overall small number of spike responses. This is a tribute to the biological plausibility with respect to data collected from adult *Drosophila*, where the KC rarely show spikes at baseline [58] and a very sparse odor response pattern [58, 119]. Due to the SFA in the ORN population that was active in all experimental conditions there was a good degree of temporal sparseness in the LN-only condition as well (especially on the chip). Again we chose to accept this effect as a baseline level of sparseness to compare other conditions against. In both implementations the expected effects of SFA in the KCs could be observed.

We have previously argued that the major functional role of temporal sparseness is the rapid and reliable stimulus encoding in a temporally dynamic environment [28, 77]. Indeed, temporal dynamics is high in the natural olfactory environment and depends on air movement and on animal speed, the latter being particularly high in flying insects. As a result, adult insects during flight or locomotion may encounter a rapid on-off stimulus scenario when passing through a thin odor filament [120–124]. It remains an open question whether the SFA mechanism is at all present in the KCs during larval stages and electrophysiological approaches to neural coding in the larva is scarce. Representation of high temporal stimulus dynamics is likely of minor importance for the larva as its locomotion is slow and the natural environment suitable for larval development such as e.g. a rotten fruit likely provides little olfactory dynamics. However, larva do perform chemotaxis and thus are able to sample olfactory gradients.

4.4. Outlook

Our current research extends the present model towards a plastic spiking network model of the larva that can perform associative learning and reward prediction [19, 125] inspired by recent modeling approaches in the adult [126, 127]. Together with biologically realistic modeling of individual larva locomotion and chemotactic

behavior [16] this will allow to reproduce behavioral [128–132] and optophysiological observations [64, 133, 134] and to generate testable hypothesis at the physiological and behavioral level. In the future this may inspire modeling virtual larvae exploring and adapting to their virtual environment in a closed loop scenario and the implementation of such mini brains on compact and low-power neuromorphic hardware for the spiking control of autonomous robots [7, 28, 135, 136].

Acknowledgments

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

ORCID iDs

Anna-Maria Jürgensen ⬤ https://orcid.org/0000-0002-7871-1887
Afshin Khalili ⬤ https://orcid.org/0000-0001-5947-5551
Elisabetta Chicca ⬤ https://orcid.org/0000-0002-5518-8990
Giacomo Indiveri ⬤ https://orcid.org/0000-0002-7109-1689
Martin Paul Nawrot ⬤ https://orcid.org/0000-0003-4133-6419

References

[1] Mead C 1990 Neuromorphic electronic systems *Proc. IEEE* **78** 1629–36
[2] Indiveri G and Sandamirskaya Y 2019 The importance of space and time for signal processing in neuromorphic agents: the challenge of developing low-power, autonomous agents that interact with the environment *IEEE Signal Process. Mag.* **36** 16–28
[3] Neftci E, Binas J, Rutishauser U, Chicca E, Indiveri G and Douglas R J 2013 Synthesizing cognition in neuromorphic electronic systems *Proc. Natl Acad. Sci.* **110** 3468–76
[4] Schmuker M, Pfeil T and Nawrot M P 2014 A neuromorphic network for generic multivariate data classification *Proc. Natl Acad. Sci.* **111** 2081–6
[5] Diamond A, Nowotny T and Schmuker M 2016 Comparing neuromorphic solutions in action: implementing a bio-inspired solution to a benchmark classification task on three parallel-computing platforms *Front. Neurosci.* **9** 491
[6] Cramer B *et al* 2020 Training spiking multi-layer networks with surrogate gradients on an analog neuromorphic substrate (arXiv:2006.07239)
[7] Helgadottir L I, Haenicke J, Landgraf T, Rojas R and Nawrot M P 2013 Conditioned behavior in a robot controlled by a spiking neural network *Int. IEEE/EMBS Conf. Neural Engineering (NER)* pp 891–4
[8] Galluppi F, Denk C, Meiner M C, Stewart T C, Plana L A, Eliasmith C, Furber S and Conradt J 2014 Event-based neural computing on an autonomous mobile platform *2014 IEEE Int. Conf. Robotics and Automation (ICRA)* (IEEE) pp 2862–7
[9] Schoepe T, Janotte E, Milde M B, Bertrand O J N, Egelhaaf M and Chicca E 2021 Finding the gap: neuromorphic motion vision in cluttered environments (arXiv:2102.08417)
[10] Heisenberg M 1995 Pattern recognition in insects *Curr. Opin. Neurobiol.* **5** 475–81
[11] Carrasco D, Larsson M C and Anderson P 2015 Insect host plant selection in complex environments *Curr. Opin. Insect. Sci.* **8** 1–7
[12] Laska M, Galizia C G, Giurfa M and Menzel R 1999 Olfactory discrimination ability and odor structure activity relationships in honeybees *Chem. Senses* **24** 429–38
[13] Meckenhäuser G, Krämer S, Farkhooi F, Ronacher B and Nawrot M P 2014 Neural representation of calling songs and their behavioral relevance in the grasshopper auditory system *Front. Syst. Neurosci.* **8** 183
[14] Avarguès-Weber A, Portelli G, Benard J, Dyer A and Giurfa M 2010 Configural processing enables discrimination and categorization of face-like stimuli in honeybees *J. Exp. Biol.* **213** 593–601
[15] Collett M, Chittka L and Collett T S 2013 Spatial memory in insect navigation *Curr. Biol.* **23** 789–800
[16] Antoine W and Graham P 2012 What can we learn from studies of insect navigation? *Anim. Behav.* **84** 13–20
[17] Menzel R and Greggers U 2015 The memory structure of navigation in honeybees *J. Comp. Physiol.* A **201** 547–61
[18] Knaden M and Graham P 2016 The sensory ecology of ant navigation: from natural environments to neural mechanisms *Annu. Rev. Entomol.* **61** 63–76

[19] Sakagiannis P, Jürgensen A-M and Nawrot M P 2021 A realistic locomotory model of drosophila larva for behavioral simulations https://doi.org/10.1101/2021.07.07.451470

[20] Chittka L, Giurfa M and Riffell J A 2019 Editorial: the mechanisms of insect cognition *Front. Psychol.* **10** 2751

[21] Dacke M and Srinivasan M V 2008 Evidence for counting in insects *Anim. Cognit.* **11** 683–9

[22] Skorupski P, MaBouDi H, Samadi Galpayage Dona H and Chittka L 2018 Counting insects *Phil. Trans. R. Soc.* B **373** 20160513

[23] Howard S R, Avarguès-Weber A, Garcia J E, Greentree A D and Dyer A G 2018 Numerical ordering of zero in honey bees *Science* **360** 1124–6

[24] MaBouDi H, Samadi Galpayage Dona H, Gatto E, Loukola O J, Buckley E, Onoufriou P D, Skorupski P and Chittka L 2020 Bumblebees use sequential scanning of countable items in visual patterns to solve numerosity tasks *Integr. Comp. Biol.* **60** 929–42

[25] Assisi C, Stopfer M and Bazhenov M 2020 Optimality of sparse olfactory representations is not affected by network plasticity *PLoS Comput. Biol.* **16** e1007461

[26] Wessnitzer J, Young J M, Armstrong J D and Webb B 2012 A model of non-elemental olfactory learning in Drosophila *J. Comput. Neurosci.* **32** 197–212

[27] Ardin P, Peng F, Mangan M, Lagogiannis K and Webb B 2016 Using an insect mushroom body circuit to encode route memory in complex natural environments *PLoS Comput. Biol.* **12** e1004683

[28] Rapp H and Nawrot M P 2020 A spiking neural program for sensorimotor control during foraging in flying insects *Proc. Natl Acad. Sci.* **117** 28412–21

[29] Müller J, Nawrot M P, Menzel R and Landgraf T 2018 A neural network model for familiarity and context learning during honeybee foraging flights *Biol. Cybern.* **112** 113–26

[30] Rost T, Ramachandran H, Nawrot M P and Chicca E 2013 A neuromorphic approach to auditory pattern recognition in cricket phonotaxis *2013 European Conf. Circuit Theory and Design (ECCTD)* (IEEE) pp 1–4

[31] Dalgaty T, Vianello E, De Salvo B and Casas J 2018 Insect-inspired neuromorphic computing *Curr. Opin. Insect. Sci.* **30** 59–66

[32] Lungu I-A, Riehle A, Nawrot M P and Schmuker M 2017 Predicting voluntary movements from motor cortical activity with neuromorphic hardware *IBM J. Res. Dev.* **61** 5:1–5:12

[33] Dalgaty T, Miller J P, Vianello E and Casas J 2021 Bio-inspired architectures substantially reduce the memory requirements of neural network models *Front. Neurosci.* **15** 156

[34] Barlow H B 1959 Sensory mechanisms, the reduction of redundancy, and intelligence *Mechanisation of Thought Processes* (London:: Her Majesty's Stationery Office) pp 535–9

[35] Olshausen B and Field D 2004 Sparse coding of sensory inputs *Curr. Opin. Neurobiol.* **14** 481–7

[36] Perez-Orive J, Mazor O, Turner G C, Cassenaer S, Wilson R I and Laurent G 2002 Oscillations and sparsening of odor representations in the mushroom body *Science* **297** 359–65

[37] Jortner R A, Farivar S S and Laurent G 2007 A simple connectivity scheme for sparse coding in an olfactory system *J. Neurosci.* **27** 1659–69

[38] Finelli L A, Haney S, Bazhenov M, Stopfer M and Sejnowski T J 2008 Synaptic learning rules and sparse coding in a model sensory system *PLoS Comput. Biol.* **4** e1000062

[39] Kloppenburg P and Nawrot M P 2014 Neural coding: sparse but on time *Curr. Biol.* **24** R957–9

[40] Stopfer M 2015 Central processing in the mushroom bodies *Curr. Opin. Insect. Sci.* **6** 99–103

[41] Poo C and Isaacson J S 2009 Odor representations in olfactory cortex: 'sparse' coding, global inhibition, and oscillations *Neuron* **62** 850–61

[42] Häusler C, Susemihl A and Nawrot M P 2013 Natural image sequences constrain dynamic receptive fields and imply a sparse code *Brain Res.* **1536** 53–67

[43] Wolfe J, Houweling A R and Brecht M 2010 Sparse and powerful cortical spikes *Curr. Opin. Neurobiol.* **20** 306–12

[44] Isaacson J S 2010 Odor representations in mammalian cortical circuits *Curr. Opin. Neurobiol.* **20** 328–31

[45] Hromádka T, DeWeese M R and Zador A M 2008 Sparse representation of sounds in the unanesthetized auditory cortex *PLoS Biol.* **6** e16

[46] Laughlin S B and Sejnowski T J 2003 Communication in neuronal networks *Science* **301** 1870–4

[47] Lin A C, Bygrave A M, de Calignon A, Lee T and Miesenböck G 2014 Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination *Nat. Neurosci.* **17** 559–68

[48] Litwin-kumar A, Harris K D, Axel R, Sompolinsky H and Abbott L F 2017 Optimal degrees of synaptic connectivity *Neuron* **93** 1153–64

[49] Caron S J C, Ruta V, Abbott L F and Axel R 2013 Random convergence of olfactory inputs in the Drosophila mushroom body *Nature* **497** 113–7

[50] Huerta R, Nowotny T, García-Sanchez M, Abarbanel H D I and Rabinovich M I 2004 Learning classification in the olfactory system of insects *Neural Comput.* **16** 1601–40

[51] Huerta R and Nowotny T 2009 Fast and robust learning by reinforcement signals: explorations in the insect brain *Neural Comput.* **21** 2123–51

[52] Betkiewicz R, Lindner B and Nawrot M P 2020 Circuit and cellular mechanisms facilitate the transformation from dense to sparse coding in the insect olfactory system *eNeuro* **7** 0305–18

[53] Palm G 2013 Neural associative memories and sparse coding *Neural Netw.* **37** 165–71

[54] Ito I, Ong R C-Y, Raman B and Stopfer M 2008 Sparse odor representation and olfactory learning *Nat. Neurosci.* **11** 1177–84

[55] Herikstad R, Baker J, Lachaux J-P, Gray C M and Yen S-C 2011 Natural movies evoke spike trains with low spike time variability in cat primary visual cortex *J. Neurosci.* **31** 15844–60

[56] Haider B, Krause M R, Duque A, Yu Y, Touryan J, Mazer J A and McCormick D A 2010 Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation *Neuron* **65** 107–21

[57] Laurent G and Naraghi M 1994 Odorant-induced oscillations in the mushroom bodies of the locust *J. Neurosci.* **14** 2993–3004

[58] Turner G C, Bazhenov M and Laurent G 2008 Olfactory representations by Drosophila mushroom body neurons *J. Neurophysiol.* **99** 734–46

[59] Menzel R 2012 The honeybee as a model for understanding the basis of cognition *Nat. Rev. Neurosci.* **13** 758–68

[60] Moradi S, Qiao N, Stefanini F and Indiveri G 2018 A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs) *IEEE Trans. Biomed. Circuits Syst.* **12** 106–22

[61] Stimberg M, Brette R and Goodman D F M 2019 Brian 2, an intuitive and efficient neural simulator *eLife* **8** e47314

[62] Eichler K *et al* 2017 The complete connectome of a learning and memory centre in an insect brain *Nature* **548** 175–82

[63] Berck M E *et al* 2016 The wiring diagram of a glomerular olfactory system *eLife* **5** e14859

[64] Saumweber T *et al* 2018 Functional architecture of reward learning in mushroom body extrinsic neurons of larval Drosophila *Nat. Commun.* **9** 1104

[65] Chicca E, Stefanini F, Bartolozzi C and Indiveri G 2014 Neuromorphic electronic circuits for building autonomous cognitive systems *Proc. IEEE* **102** 1367−88

[66] Boahen K A 1998 Communicating neuronal ensembles between neuromorphic chips *Neuromorphic Systems Engineering: Neural Networks in Silicon* ed T S Lande (Norwell, MA: Kluwer Academic Publishers) pp 229−59

[67] Deiss S R, Douglas R J and Whatley A M 1998 A pulse-coded communications infrastructure for neuromorphic systems *Pulsed Neural Networks* ed W Maass and C M Bishop (Cambridge, MA: MIT Press) pp 157−78

[68] Burkitt A N 2006 A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input *Biol. Cybern.* **95** 1−19

[69] Hoare D J, Humble J, Jin D, Gilding N, Petersen R, Cobb M and McCrohan C 2011 Modeling peripheral olfactory coding in drosophila larvae *PLoS One* **6** e22996

[70] Kreher S A, Kwon J Y and Carlson J R 2005 The molecular basis of odor coding in the Drosophila larva *Neuron* **46** 445−56

[71] Willmore B and Tolhurst D J 2001 Characterizing the sparseness of neural codes *Netw. Comput. Neural Syst.* **12** 255−70

[72] Treves A and Rolls E T 1991 What determines the capacity of autoassociative memories in the brain? *Netw. Comput. Neural Syst.* **2** 371−97

[73] MaBouDi H, Shimazaki H, Giurfa M and Chittka L 2017 Olfactory learning without the mushroom bodies: spiking neural network models of the honeybee lateral antennal lobe tract reveal its capacities in odour memory tasks of varied complexities *PLoS Comput. Biol.* **13** e1005551

[74] Wilson R I, Turner G C and Laurent G 2004 Transformation of olfactory representations in the drosophila antennal lobe *Science* **303** 366−70

[75] Strube-Bloss M F, Nawrot M P and Menzel R 2011 Mushroom body output neurons encode odor-reward associations *J. Neurosci.* **31** 3129−40

[76] Brill M F, Rosenbaum T, Reus I, Kleineidam C J, Nawrot M P and Rössler W 2013 Parallel processing via a dual olfactory pathway in the honeybee *J. Neurosci.* **33** 2443−56

[77] Farkhooi F, Froese A, Muller E, Menzel R and Nawrot M P 2013 Cellular adaptation facilitates sparse and reliable coding in sensory pathways *PLoS Comput. Biol.* **9** e1003251

[78] Couto A, Alenius M and Dickson B J 2005 Molecular, anatomical, and functional organization of the Drosophila olfactory system *Curr. Biol.* **15** 1535−47

[79] Hubel D H and Wiesel T N 1959 Receptive fields of single neurones in the cat's striate cortex *J. Physiol.* **148** 574−91

[80] Martin J P, Beyerlein A, Dacks A M, Reisenman C E, Riffell J A, Lei H and Hildebrand J G 2011 The neurobiology of insect olfaction: sensory processing in a comparative context *Prog. Neurobiol.* **95** 427−47

[81] Krofczik S, Menzel R and Nawrot M P 2009 Rapid odor processing in the honeybee antennal lobe network *Front. Comput. Neurosci.* **2** 9

[82] Deisig N, Giurfa M and Sandoz J C 2010 Antennal lobe processing increases separability of odor mixture representations in the honeybee *J. Neurophysiol.* **103** 2185−94

[83] Fuscà D and Kloppenburg P 2021 Odor processing in the cockroach antennal lobe—the network components *Cell Tissue Res.* **383** 59−73

[84] Leitch B and Laurent G 1996 Gabaergic synapses in the antennal lobe and mushroom body of the locust olfactory system *J. Comp. Neurol.* **372** 487−514

[85] Olsen S R and Wilson R I 2008 Lateral presynaptic inhibition mediates gain control in an olfactory circuit *Nature* **452** 956−60

[86] Linster C and Smith B H 1997 A computational model of the response of honey bee antennal lobe circuitry to odor mixtures: overshadowing, blocking and unblocking can arise from lateral inhibition *Behav. Brain Res.* **87** 1−14

[87] Nowotny T, Huerta R, Abarbanel H D I and Rabinovich M I 2005 Self-organization in the olfactory system: one shot odor recognition in insects *Biol. Cybern.* **93** 436−46

[88] Mosqueiro T S and Huerta R 2014 Computational models to understand decision making and pattern recognition in the insect brain *Curr. Opin. Insect. Sci.* **6** 80−5

[89] Lei Z, Chen K, Li H, Liu H and Guo A 2013 The GABA system regulates the sparse coding of odors in the mushroom bodies of Drosophila *Biochem. Biophys. Res. Commun.* **436** 35−40

[90] Benda J 2021 Neural adaptation *Curr. Biol.* **31** 110−6

[91] Si G, Kanwal J K, Hu Yu, Tabone C J, Baron J, Berck M, Vignoud G and Samuel A D T 2019 Structured odorant response patterns across a complete olfactory receptor neuron population *Neuron* **101** 950.e7−62.e7

[92] Nagel K I and Wilson R I 2011 Biophysical mechanisms underlying olfactory receptor neuron dynamics *Nat. Neurosci.* **14** 208−16

[93] Gorur-Shandilya S, Demir M, Long J, Clark D A and Emonet T 2017 Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli *eLife* **6** e27670

[94] Martelli C, Carlson J R and Emonet T 2013 Intensity invariant dynamics and odor-specific latencies in olfactory receptor neuron response *J. Neurosci.* **33** 6285−97

[95] Brandão S C, Silies M and Martelli C 2021 Adaptive temporal processing of odor stimuli *Cell Tissue Res.* **383** 125−41

[96] Demmer H and Kloppenburg P 2009 Intrinsic membrane properties and inhibitory synaptic input of Kenyon cells as mechanisms for sparse coding? *J. Neurophysiol.* **102** 1538−50

[97] Kropf J and Rössler W 2018 *In situ* recording of ionic currents in projection neurons and Kenyon cells in the olfactory pathway of the honeybee *PLoS One* **13** e0191425

[98] Bhandawat V, Olsen S R, Gouwens N W, Schlief M L and Wilson R I 2007 Sensory processing in the Drosophila antennal lobe increases reliability and separability of ensemble odor representations *Nat. Neurosci.* **10** 1474−82

[99] Meyer A, Galizia C G and Nawrot M P 2013 Local interneurons and projection neurons in the antennal lobe from a spiking point of view *J. Neurophysiol.* **110** 2465−74

[100] Szyszka P, Ditzen M, Galkin A, Galizia C G and Menzel R 2005 Sparsening and temporal sharpening of olfactory representations in the honeybee mushroom bodies *J. Neurophysiol.* **94** 3303−13

[101] Nawrot M P 2012 Dynamics of sensory processing in the dual olfactory pathway of the honeybee *Apidologie* **43** 269−91

[102] Indiveri G *et al* 2011 Neuromorphic silicon neuron circuits *Front. Neurosci.* **5** 73

[103] Pfeil T, Jordan J, Tetzlaff T, Grübl A, Schemmel J, Diesmann M and Meier K 2016 Effect of heterogeneity on decorrelation mechanisms in spiking neural networks: a neuromorphic-hardware study *Phys. Rev. X* **6** 021023

[104] Liang D, Kreiser R, Nielsen C, Qiao N, Sandamirskaya Y and Indiveri G 2019 Neural state machines for robust learning and control of neuromorphic agents *IEEE J. Emerg. Sel. Topics Circuits Syst.* **9** 679−89

[105] Cao Z, Cheng L, Zhou C, Gu N, Wang X and Tan M 2015 Spiking neural network-based target tracking control for autonomous mobile robots *Neural Comput. Appl.* **26** 1839–47

[106] Knight J C and Nowotny T 2021 Larger GPU-accelerated brain simulations with procedural connectivity *Nat. Comput. Sci.* **1** 136–42

[107] Stimberg M, Goodman D F M and Nowotny T 2020 Brian2genn: accelerating spiking neural network simulations with graphics hardware *Sci. Rep.* **10** 1–12

[108] Delbruck T, Berner R, Lichtsteiner P and Dualibe C 2010 32-bit configurable bias current generator with sub-off-current capability *2010 Int. Symp. Circuits and Systems, (ISCAS)* (Paris, France) (IEEE) pp 1647–50

[109] Masuda-Nakagawa L M, Ito K, Awasaki T and O'Kane C J 2014 A single GABAergic neuron mediates feedback of odor-evoked signals in the mushroom body of larval Drosophila *Front. Neural Circ.* **8**

[110] Haenicke J, Yamagata N, Zwaka H, Nawrot M P and Menzel R 2018 Neural correlates of odor learning in the presynaptic microglomerular circuitry in the honeybee mushroom body Calyx *eNeuro* **5**

[111] Haenicke J 2015 Modeling insect inspired mechanisms of neural and behavioral plasticity *PhD Thesis* FU, Berlin https://refubium.fu-berlin.de/handle/fub188/8123 https://refubium.fu-berlin.de/handle/fub188/8123

[112] Olsen S R, Bhandawat V and Wilson R I 2010 Divisive normalization in olfactory population codes *Neuron* **66** 287–99

[113] Marr D and Thomas Thach W 1991 A theory of cerebellar cortex *From the Retina to the Neocortex* (Boston, MA: Birkhäuser) pp 11–50 https://link.springer.com/chapter/10.1007/978-1-4684-6775-8_3

[114] Albus J S 1971 A theory of cerebellar function *Math. Biosci.* **10** 25–61

[115] Nowotny T, Rabinovich M I, Huerta R and Abarbanel H D I 2003 Decoding temporal information through slow lateral excitation in the olfactory system of insects *J. Comput. Neurosci.* **15** 271–81

[116] Froese A, Szyszka P and Menzel R 2014 Effect of GABAergic inhibition on odorant concentration coding in mushroom body intrinsic neurons of the honeybee *J. Comp. Physiol.* A **200** 183–95

[117] Kee T, Sanda P, Gupta N, Stopfer M and Bazhenov M 2015 Feed-forward versus feedback inhibition in a basic olfactory circuit *PLoS Comput. Biol.* **11** e1004531

[118] Farkhooi F, Muller E and Nawrot M P 2011 Adaptation reduces variability of the neuronal population code *Phys. Rev.* E **83** 050905

[119] Honegger K S, Campbell R A A and Turner G C 2011 Cellular-resolution population imaging reveals robust sparse coding in the Drosophila mushroom body *J. Neurosci.* **31** 11772–85

[120] Pannunzi M and Nowotny T 2019 Odor stimuli: not just chemical identity *Front. Physiol.* **10** 1428

[121] Celani A, Villermaux E and Vergassola M 2014 Odor landscapes in turbulent environments *Phys. Rev.* X **4** 041015

[122] Kree M, Duplat J and Villermaux E 2013 The mixing of distant sources *Phys. Fluids* **25** 091103

[123] Demir M, Kadakia N, Anderson H D, Clark D A and Emonet T 2020 Walking drosophila navigate complex plumes using stochastic decisions biased by the timing of odor encounters *eLife* **9** e57524

[124] Vickers N J, Christensen T A, Baker T C and Hildebrand J G 2001 Odour-plume dynamics influence the brain's olfactory code *Nature* **410** 466–70

[125] Jürgensen A-M, Khalili A and Nawrot M P 2019 Reinforcement-mediated plasticity in a spiking model of the drosophila larva olfactory system *BMC Neurosci.* **20** P225

[126] Bennett J E M, Philippides A and Nowotny T 2021 Learning with reinforcement prediction errors in a model of the drosophila mushroom body *Nat. Commun.* **12** 1–14

[127] Springer M and Nawrot M P 2021 A mechanistic model for reward prediction and extinction learning in the fruit fly *eNeuro* **8** ENEURO.0549-20.2021

[128] Schleyer M, Miura D, Tanimura T and Gerber B 2015 Learning the specific quality of taste reinforcement in larval drosophila *eLife* **4** e04711

[129] Gerber B and Hendel T 2006 Outcome expectations drive learned behaviour in larval Drosophila *Proc. R. Soc.* B. **273** 2965–8

[130] Schleyer M, Reid S F, Pamir E, Saumweber T, Paisios E, Davies A, Gerber B and Louis M 2015 The impact of odor-reward memory on chemotaxis in larval Drosophila *Learn. Mem.* **22** 267–77

[131] Saumweber T, Husse J and Gerber B 2011 Innate attractiveness and associative learnability of odors can be dissociated in larval Drosophila *Chem. Senses* **36** 223–35

[132] Schleyer M, Saumweber T, Nahrendorf W, Fischer B, von Alpen D, Pauls D, Thum A and Gerber B 2011 A behavior-based circuit model of how outcome expectations organize learned behavior in larval drosophila *Learn. Mem.* **18** 639–53

[133] Schroll C *et al* 2006 Light-induced activation of distinct modulatory neurons triggers appetitive or aversive learning in Drosophila larvae *Curr. Biol.* **16** 1741–7

[134] Thum A S and Gerber B 2019 Connectomics and function of a memory network: the mushroom body of larval Drosophila *Curr. Opin. Neurobiol.* **54** 146–54

[135] Drix D and Schmuker M 2021 Resolving fast gas transients with metal oxide sensors *ACS Sens.* **6** 688–92

[136] Spaeth A, Tebyani M, Haussler D and Teodorescu M 2020 Spiking neural state machine for gait frequency entrainment in a flexible modular robot *PLoS One* **15** e0240267

3 Chapter: Circuit mechanisms of associative learning

3.1 Prediction error drives associative olfactory learning and conditioned behavior in a spiking model of Drosophila larva

Prediction error drives associative olfactory learning and conditioned behavior in a spiking model of *Drosophila larva*

Anna-Maria Jürgensen[a], Panagiotis Sakagiannis[a], Michael Schleyer[b,c], Bertram Gerber[b,d,e], and Martin Paul Nawrot[a]

[a]Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Cologne, Germany

[b]Leibniz Institute for Neurobiology (LIN), Department of Genetics, Magdeburg, Germany

[c]Institute for the Advancement of Higher Education, Faculty of Science, Hokkaido University, Japan

[d]Institute for Biology, Otto-von-Guericke University, Magdeburg, Germany

[e]Center for Brain and Behavioral Sciences (CBBS), Otto-von-Guericke University, Magdeburg, Germany

Abstract

Predicting reinforcement from the presence of environmental clues is an essential component of guiding goal-directed behavior. In insect brains, the mushroom body is central to learning the necessary associations between sensory signals and reinforcement. We propose a biologically realistic spiking network model of the Drosophila larva olfactory pathway for the association of odors and reinforcement to bias behavior towards approach or avoidance. We demonstrate that prediction error coding through the integration of currently present and expected reinforcement in dopaminergic neurons can serve as a driving force in learning that can, combined with a synaptic homeostasis mechanism, account for experimentally observed features of acquisition and loss of associations in the larva that depend on the intensity of odor and reinforcement and temporal features of their pairing. To allow direct comparisons of our simulations with behavioral data [1], we model learning-induced plasticity over the complete time course of behavioral experiments and simulate the locomotion of individual larvae towards or away from odor sources in a virtual environment.

1

Introduction

Goal-directed behavior in dynamic situations benefits from the ability to predict future conditions in the environment from the occurrence of sensory clues. In insects, the mushroom body (MB) is the central brain structure for multi-sensory integration, involved in memory formation and recall [2, 3]. It is at the core of learning and retaining valuable associations between sensory inputs and reinforcement in the synapses between the MB intrinsic and its output neurons [4–7].

One of the underlying mechanisms is associative learning, a process that gradually establishes a relationship between two previously unrelated elements. In classical conditioning, the conditioned sensory stimulus (CS) obtains behavioral relevance through its concurrence with the reinforcing unconditioned stimulus (US), an acquisition process depending dynamically on their spatiotemporal proximity. The temporal evolution of this process has been formalized in the Rescorla-Wagner (RW) model (eqn. 1) [8].

$$\Delta V = \alpha \cdot (\lambda_{\text{US}} - V(t)),$$
$$V(t + \Delta t) = V(t) + \Delta V.$$

(1)

Here, a CS obtains predictive power of concurrent or successive US [8], that depends on the strength of the already acquired association between the CS and US $V(t)$, allowing for anticipatory behavior to the expected US [9, 10]. The acquisition of this association terminates when the US is fully predicted. Until then, the change in associative strength ΔV is proportional to the difference between the maximum associative strength (or asymptote) λ_{US} and the current associative strength $V(t)$ (eqn. 1). The maximum associative strength is a property of the US, determined mainly by the intensity of the reinforcement. While the current associative strength $V(t)$ is defined by the shared learning history of CS and US [8]. The concept of prediction error (PE) [11] is a derivative of the Rescorla-Wagner model [8]. The error signal equals the difference between the current λ_{US} and the predicted value of US $V(t)$. Over the course of the memory acquisition/training phase, the pace of learning, which can be formalized as the slope of the acquisition curve, decreases as the PE is reduced, minimizing the driving force for changes of the association [8, 11]. This difference is multiplied with a learning rate parameter (α), here combined for the CS and the US (eqn. 1).

This continuous optimization of predictions, guided by the PE, could allow animals to efficiently adapt their goal-directed behavior in dynamic environments. Among the most relevant associations to be learned are those that enable the prediction of reward or punishment. Dopaminergic neurons (DANs) have long been known to encode information about reward and punishment. These types of neurons respond to the presence of rewards and punishments in the environment, both in vertebrates [12–17], as well as invertebrates [18–21]. The electrical stimulation or optogenetic activation of DANs induces approach or avoidance both in vertebrates [22–25] and invertebrates [20, 26–32]. In adult [5, 6, 33, 34] and larval [20, 35, 36] *Drosophila* this approach or avoidance learning is facilitated by

2

the modulation of MB output synapses by DAN activity. Ultimately DANs do not only signal the presence of rewards or punishments but have also been suggested to encode PE in various vertebrate species [16, 37–40] and might have a similar function in insects [19, 20, 32, 34, 41–44].

We utilize our spiking model of the *Drosophila* larva MB in one brain hemisphere that forms associations of odors with reinforcement to further test the hypothesis that PE coding within this circuit takes place in DANs that receive input from the output neurons of the MB or their downstream partners [20, 45–47], that might provide feedback to the DANs. Beyond the scope of similar models [20, 48–51](see Discussion, section: Comparison with other MB models), we demonstrate that this mechanism can reproduce the experimentally observed findings on the acquisition of associations of odors with reinforcement in a time-resolved manner [1]. To facilitate direct qualitative and quantitative comparisons with animal behavioral data, we couple this model with a realistic locomotory model of the larva [52] that captures the effects of learned associations on chemotactic behavior in individual animals.

Results

Connectome-based circuit model of the larval olfactory pathway

The network architecture of our model (fig. 1 A) is based on the anatomy of the olfactory pathway in one *Drosophila* larva brain hemisphere [20, 29, 53, 54] (for more details see Methods, section: Network model). Peripheral processing is carried out by 21 olfactory receptor neurons (ORNs), each expressing a different olfactory receptor type [53, 55, 56]. ORNs form one-to-one excitatory synaptic connections with 21 projection neurons (PNs) and 21 local interneurons (LNs) in the antennal lobe [53]. Each LN connects with all PNs via inhibitory GABAergic synapses, establishing a motif for lateral inhibition within the antennal lobe. The 72 mature larval Kenyon cells (KCs) [54] are the excitatory principal cells of the MB. Each KC receives excitatory input from 2-6 randomly selected PNs [54]. The KCs are subjected to feedback inhibition, provided via the GABAergic anterior paired lateral (APL) neuron, which receives input from all KCs [29]. Only mature KCs, characterized by a fully developed dendrite, are included in this model, yielding a complete convergent synaptic KC >APL connectivity. The output region of the MB is organized in compartments, in which the KC axons converge with the dendrites of one or few MB output neurons (MBONs) [20, 54]. Our model assumes two MBONs from two different compartments that are representative of two different categories of output neurons of the MB that mediate either approach or avoidance [4–6, 33, 35, 36, 57–59] with a single MBON each. Both MBONs receive excitatory input from all of the KCs to fully capture the information that is normally represented by the complete set of MBONs. Each compartment is also innervated by a single DAN, signaling either reward or punishment and targeting the KC>MBON synapses to facilitate learning (for a discussion of all simplifications compared to the animal brain, see Methods, section: Network model).

3

Learning through KC>MBON plasticity

We assume that the KC>MBON synapses undergo plasticity, based on strong experimental evidence in larval [20, 35, 36] and adult flies [5, 6, 34]. This plasticity requires the convergence of the sensory pathways in the form of KC activation and of the reinforcing pathway, mediated by neuromodulatory DAN signaling at the synaptic site. We employ a two-factor learning rule (eqn. 2) at each KC>MBON synapse (fig. 1 A,B). The first factor is expressed in the pre-synaptic KC activation by an odor, tagging the synapse eligible for modification. This is modeled via an exponentially decaying eligibility trace $e_i(t)$, which is set to a 1 whenever the respective KC elicits a spike (fig. 1 B). The decay time constant determines the window of opportunity for synaptic change. The presence of reinforcement (reward or punishment) constitutes the second factor and is signaled by the reward-mediating DAN$_+$ or punishment-mediating DAN$_-$, respectively. Spiking of the DAN provides a neuromodulatory reinforcement signal $R(t)$ to the synaptic site. If a DAN spike coincides with positive eligibility at the synapse, the respective synaptic weight is reduced. At each synapse i, the reduction of synaptic weight Δw_i depends on the learning rate a (table S1) and is proportional to the amplitude of the eligibility trace $e_i(t)$ (fig. 1 B):

$$\Delta w_i = -a \cdot e_i(t) \cdot R(t) \leq 0. \tag{2}$$

We introduce a synaptic homeostasis mechanism (eqn. 3) that modulates the effects of plasticity at each KC>MBON synapse to account for the experimentally observed loss of a learned association when reinforcement is omitted [41, 60, 61] and to ensure continued input to both MBONs. With each MBON spike, the current weight w_i of each respective KC>MBON synapse is increased, proportionally to the extent to which the weight differs from its original value w_{init} (table S1) and multiplied with a homeostatic regulation factor h (table S1). This mechanism serves as an implementation for the loss of the association when the reinforcement is omitted. While reinforcement is present, the learning curve will either continue to rise or remain at the asymptote if already saturated. The interaction of the two mechanisms of learning and unlearning at the level of the individual KC>MBON synapses allows to include the loss of learned associations, when continued reinforcement is omitted (see Discussion, section: A mechanistic implementation of the RW model) and also ensures continued input to the MBONs, despite the reduction of input weights over the course of the learning process (eqn. 2). The homeostatic factor h hereby serves as an implementation of a time constant of this exponential process. The interaction of synaptic plasticity and homeostatic regulation defines the magnitude of the weight at the next simulation timestep $t + \Delta t$ as

$$w_i(t + \Delta t) = w_i(t) + \Delta w_i + (w_{init} - w_i(t)) \cdot h. \tag{3}$$

It has been shown in behavioral experiments that specific MBONs encode a behavioral tendency to either approach or avoid a currently perceived stimulus, depending on the acquired stimulus valence [4–6, 36, 57, 58]. In the naive state of our model, all KC>MBON synapses have the same

129 initial weights w_{init} (table S1), and hence the spiking activity of both MBONs is highly similar.
130 Learning alters the KC>MBON synaptic weights and thus skews the previously balanced MBON
131 output. This acquired imbalance between MBON outputs biases behavior towards the approach or
132 avoidance of the conditioned odor. To quantify the effect of learning, we compute the behavioral
133 bias BB (eqn. 4) from the firing rates of both MBONs over $T = 1s$ as follows:

$$BB = \frac{MBON_+ - MBON_-}{T}. \tag{4}$$

Implementation of prediction error coding in the KC-MBON-DAN motif

135 In the larva, many DANs and other modulatory neurons receive excitatory or inhibitory input
136 from different MBONs, either in a direct manner or via one or two interneurons [20]. Based on
137 this observation, we constructed our hypothetical feedback motif (for similar models see discus-
138 sion section: Comparison with other MB models). In the model, DANs are activated by external
139 reward/punishment signals and also receive excitatory and indirect inhibitory feedback from both
140 MBONs (fig. 1 A). As the initial balance between the two MBON outputs shifts over the course of
141 the training process, the amount of excitatory and inhibitory feedback that DANs receive continues
142 to diverge, allowing the DANs to access the model's learning history. Ultimately DAN activation
143 signals the difference between the current external activation and the expected activation based
144 on prior learning, implemented as the difference between excitatory and inhibitory MBON>DAN
145 feedback. Including this feedback leads to learning curves that saturate when the reward is fully
146 predicted, and the prediction error approaches zero (fig. 2 A,D). This effect disappears, when the
147 feedback circuit is disabled (fig. 2 A). In this case the behavioral bias quickly reaches the maximum
148 value of the measure when the MBON_ elicits no more spikes and can not encode further learning.
149 Increasing reward intensity, learning rate or odor intensity (see Methods, section: Experimental
150 protocols) foster a faster acquisition of the association and increases the maximum strength of the
151 association at the same time (fig. 2 A).

152 Increasing the reward intensity after a 2.5 min (black curve), or 5 min (gray curve) of appetitive
153 training, results in a steeper slope of the learning curve and also increases the maximum during
154 training trials of 2.5min duration with increased reward intensity (fig. 2 B). Higher intensity of the
155 reward results in an average DAN spike rate of 39.14Hz($std = 1.27$(standard deviation)) compared
156 to 33.11Hz($std = 1.34$).

157 Additionally, we tested for loss of the acquired association as the reduction in behavioral memory
158 expression, over the course of prolonged exposure to the CS without the US, following initial memory
159 acquisition [8, 62]. We test this in our model experiments by presenting the odor, previously paired
160 with reward, for an extended period of time, in the absence of reinforcement. During the test
161 phase and without the presence of reward to trigger synaptic KC>MBON_ weight reduction, the
162 extinction mechanism is no longer outweighed by learning and drives each individual weight back
163 towards w_{init} (fig. 2 C, upper panel). We also demonstrate the interaction of the learning rule with

this mechanism in figure S1, where the learning rate remains constant but the magnitude of the homeostatic regulation was manipulated to show that both mechanisms need to be in balance.

Learned preference and behavior generalize to similar odors

We trained our model by pairing a reward with a single odor for 4min. After the training procedure, we tested the behavioral bias either for the same or a different odor, following the experimental approach used in the larva [63]. Mimicking the experimental data, we show that the odor preference is highest if training odor and testing odor are identical in the case of training with 3-octanol. When amylacetate is used during training, 3-octanol preference is increased (fig. 3 A). Since 3-octanol activates a subset of the ORNs activated by amylacetate (fig. 1 D), some of which with higher rates than in the case of amylacetate, we also tested for generalization using a set of ORN activation patterns with a controlled degree of overlap (see Methods, section: Sensory input, fig. 1 D) and show that with decreasing similarity, the generalization effect to a new odor is diminished (fig. 3 A). Figure 3 B shows the network response to 30sec stimulations with amylacetate and 3-octanol in a single exemplary model instance. On the level of the ORNs, 3-octanol merely activated a subset of the amylacetate-activated neurons. The uniqueness of the odor identities is enhanced in the KC population [64].

The model reproduces temporal features in trace conditioning experiments

Including an odor-evoked eligibility trace at the KC>MBON synapses allows the model to maintain the sensory odor representation for a time window, during which reinforcement will trigger synaptic change (fig. 1 B). The time window between odor and reward onset $(0, 10, 20, 30, 40, 50, 60, 120s)$ was varied for trace conditioning experiments with a 30s presentation of odor and reward that was repeated three times. A small inter-stimulus-interval (ISI) of 10 to 30s leads to an increase in behavioral bias compared to the complete overlap of odor and reinforcement (fig. 3 C), using the extended window of opportunity for synaptic change triggered by each KC spike. Long ISIs do not lead to learning as the eligibility trace declines back to zero during this time (fig. 3 C). These findings match observations from experiments in larvae [29, 65, 66] with the caveat that the trace in the real larva brain seems to extend for a slightly longer period of time, compared with our experiments.

The model reproduces paired and unpaired associative conditioning experiments

To test if learning, driven by prediction error, can account for learned larval behavior, we replicated single-trial conditioning experiments performed with larvae [1] in simulation. In these experiments, animals were trained with the odor amylacetate in a single trial of varying duration $(1 - 8 \text{ min})$. To this end, larvae were placed on a Petri dish coated with an agar-sugar substrate and the odor in

6

two small containers for diffusion in the air (paired training). Either before or following this training protocol, larvae underwent a single trial without sugar and odor. Afterward, the animals were transferred to a new dish with two odor containers placed on different sides (one of them contained amylacetate, and the other one was empty). This paired training was compared with an unpaired protocol with separate (randomized order) presentations of amylacetate and sugar. Following the paired training protocol (odor and reward are presented concurrently), the animals showed a tendency to approach the previously rewarded odor, as measured by the difference in the number of animals on each side at the end of a 3min test phase, divided by the total number of animals. Following the experimental literature, we will refer to this measure as the preference index ([1] eqn. 15). The animals' preference is relatively consistent across training trials of different duration. Prolonged paired training did not lead to an increase in preference (fig. 5 A). These experiments did not include a test for odor preference before training, but untrained larval odor preference of odors used in learning experiments has been demonstrated elsewhere [67–69]. This paired training was compared with an unpaired protocol with separate (randomized order) presentations of amylacetate and sugar. Here the extent to which animals preferred amylacetate over no odor varied with the duration of the training trial. The longer the duration of the training, the more the preference index decreased from an initially high value but saturated around 2.5min (fig. 5 A).

We aimed to replicate these behavioral experiments on two levels. Firstly, we focused on the direct model output that reflects the strength of the acquired association between amylacetate and reward (behavioral bias, eqn. 4) and later also simulated behavior based on these biases. We simulated both the paired and unpaired training protocol (fig. 4 B). While the unpaired training yielded almost no behavioral bias, the models that underwent the paired training show an increased behavioral bias, that depended on the duration of the training and saturated for longer training duration (fig. 4 B). The simulation results reported in figure 4 B were obtained using odor-naive models that exhibited no odor preference, prior to training. To account for the experimental finding that real larvae often do have an odor preference even without any training [67–69], we readjusted our experiments to include a pre-training period of 10 minutes to start the conditioning experiments with the amylacetate-reward association already established. This adaptation of the protocol leads to results (fig. 4 C) that match the results obtained in real behavioral experiments (fig. 5 A). The paired condition in figure 4 C shows that once the behavioral bias is saturated (fig. 2 A), continued pairing maintains the association, without further increasing it. Unpaired training on the other hand, causes the behavioral bias to decrease and saturate at a lower level. For a discussion of different potential causes of a reinforcement expectation prior to training, please refer to the discussion (Comparison of modeling results to experimental findings). Figure 2 A demonstrates that disabling MBON>DAN feedback leads to a learning curve that does not saturate but instead increases with a steep slope until it reaches the maximum value for the behavioral bias eqn. 4) with a MBON_ rate of 0. To verify if this PE feedback mechanism is responsible for the difference between maintenance and loss of the association in figure 4 C, we repeated the same experiment with disabled MBON>DAN feedback.

7

The behavioral bias overall is much higher, compared to the intact network (fig. 4 B). The maximum is reached before the test phase of even the shortest 1min training experiment, with no MBON spikes elicited.

Secondly, since the effect of training in lab experiments is quantified behaviorally via spatially defined, group-level metrics (preference index and performance index (eqn. 15,eqn. 16), [1]), we performed behavioral simulations of the testing phase with groups of virtual larvae for both the paired and unpaired condition [1], allowing a straightforward comparison with the animal experiments (fig. 5 A). To this end, we utilized a realistic model for the simulation of larval locomotion and chemotactic behavior [52] that uses the behavioral bias at the output of the MB model as a constant gain factor to modulate the locomotory behavior of individual larvae towards or away from a spatially placed odor source in a virtual arena (see Methods, section: Realistic modeling of larval locomotion). The resulting preference indices, acquired across groups of independently simulated larvae (fig. 5 C), can directly be compared to the experimentally obtained preference indices (fig. 5 A). We also compare performance indices from our simulated experiments (fig. 5 D) with those from the lab experiments (fig. 5 B) and find that the model can replicate these when accounting for the odor preference at the beginning of the experiment.

Discussion

Seeking rewards and avoiding punishments by predicting change in the environment is a major motivator of animal behavior. Sensory clues can acquire the necessary predictive power to guide behavior through classical conditioning, an associative learning process potentially driven by reward/punishment PE [8, 11], as observed in vertebrates [16, 37–40]. To test the biological plausibility of the proposed PE coding motif in the larval MB and test its capacity to explain behavioral data we implemented a spiking network model of the olfactory pathway, coupled with a simulation of locomotory behavior [52]. We demonstrate that our model of PE coding results in saturating group-level and individual learning curves, where the slope and maximum of the learning curve are determined by the intensity of both the reward and the odor signal. Learning is also influenced by the timing of odor and reinforcement and can be extinguished if reinforcement is omitted during the presentation of the sensory clue. After verifying that this circuit motif enables learning as predicted by the PE theory, we show that it can also explain time-resolved larval behavior in conditioning experiments.

A mechanistic implementation of the RW model

A number of predictions can be derived from the phenomenological RW model [8] and tested in our mechanistic model thereof. We found that regardless of odor/reward intensity or the model's learning rate, the strength of the odor-reward association (quantified as the behavioral bias) saturates over time (fig. 2 A), as the strength of the already acquired association $V(t)$ approaches the maxi-

mum value supported by the given reinforcement input (λ_{US}) (eqn. 1). Consequently, our model's acquisition curve saturates at a higher value when the intensity of the reinforcement is increased (fig. 2 A,B), as predicted by the RW model, in which a stronger US should result in a higher value of λ_{US} [8]. In our model, a higher reinforcement intensity relates to a higher input rate into the respective DAN (see Methods, section: Sensory input) which translates into more frequent DAN spikes within a given window of 1 second, used to compute the behavioral bias (eqn. 4). This defines the asymptote of the learning curve. According to the RW model, increasing either the intensity of the odor or the learning rate α [8] should lead to faster acquisition of the association. In our model, the learning rate directly influences the increment of each respective synaptic weight Δw^i, while an increase of the odor intensity allows for a more frequent execution of the weight update routine, by influencing the eligibility trace (eqn. 2).

The RW model predicts that the omission of reward should result in the loss of the learned association (eqn. 1, [8]). From the equation itself, we can not infer if this loss is due to extinction or forgetting. Extinction, characterized by the possibility of recovery of the association, after its temporary loss [70], has been demonstrated in adult [71, 72] but not larval *Drosophila*. To retain the association for recovery, extinction relies on the formation of parallel memory traces for the acquisition and the loss of the association [41, 60]. The mechanism implemented in our model is overwriting the association, since the homeostatic mechanism drives the synaptic weights toward their initial value, thereby deleting the learned association with no chance of recovery, but only in the presence of olfactory input, eliciting MBON spikes. The resulting behavior during the extinction phase of the experiments presents itself in a similar way, while the underlying mechanism is different.

Comparison of modeling results to experimental findings

A variety of experiments have demonstrated group-level acquisition curves that saturate over multiple training trials or with increasing duration of a single trial in olfactory conditioning [1, 51, 73–75]. To replicate larval behavior in reward learning experiments [1] with varying duration of the learning phase (fig. 5 A,B) we trained our model with an odor and reward in a paired vs. unpaired fashion (fig. 4 B). Real larvae show a strong odor preference even after a very short training and no significant increase in their preference when trained in a paired manner for longer periods of time [1, 67]. Instead, the animals trained in an unpaired protocol start out with a similarly high odor preference, which then decreases over time [1, 67]. This behavior is very counter-intuitive since the coincidence of odor and reward should yield an association of those two stimuli and thus an increased behaviorally expressed preference for the CS [8]. To resolve this contradiction, we include the observation that animals might not be naive to the training odor prior to the beginning of the experiments in the model. In that case, the animals would enter into the experiment with an already established reward prediction that would be violated during unpaired training. Three scenarios lend themselves as plausible causes of this effect: Firstly, accidental conditioning over the course of their lifespan during which they are raised on a food substrate while being exposed to air that carries

9

many different odorants. Alternatively, or in fact, additionally, the animals might exhibit an innate preference for many odors [76–78]. Finally, the presence of the reward during reward-only phases might lead to an association of the experimental context with that reward (previously discussed by Saumweber et al. [67]). The resulting reward expectation (solely based on the always present context), unmet during the odor-only phases could lead to a prediction error signal. All three candidate explanations would yield a similar projection for the unpaired experimental protocol: A reward expectation acquired prior to the actual experiment would cause a violation of that expectation during odor-only trials of the unpaired experiments. In all three cases, the animal's preferences might also generalize to a broader array of odors, leading to an overall preference for some odors, as observed experimentally. To test this hypothesis we pre-trained our model before simulating conditioning experiments (fig. 4 C) and observed that this allows us to reproduce the animal experiments (fig. 5 A,B). Including odor preference at the beginning of the experiment ensures the model not only behaves in accordance with the RW model [8], but also fits the animal experimental results [1]. A possible alternative explanation could be a sensory habituation process to the odor that might cause odor preferences to decrease over time, resulting in the observed patterns for unpaired learning. In the paired condition this effect might be abolished by the continued presentation of odor and reward together [79].

Thus far we have tested our model in experiments where the CS and US presentation were fully overlapping (paired conditions). We now consider different onset times, with the onset of the CS always preceding the onset of the US (fig. 3,C). For these experiments we used a shorter duration of $30s$ for both CS and US presentation, repeated over three acquisition trials to mimic experimental conditions in larval experiments [29, 66] that used optogenetic activation of DANs as a proxy for sugar reward. Similar to their experiments we show that the behavioral bias clearly depends on the temporal delay between CS and US (fig. 3,C). Complete temporal overlap of CS and US (ISI=0) does not seem to expend the full potential of learning the association, instead partial overlap yields stronger associations due to the extended window of opportunity for synaptic change triggered by the odor's eligibility trace. In our model, the eligibility trace e(t) represents a molecular process that maintains the odor signal locally in the KC>MBON synapses (eqn. 2). Zeng et al. [80] demonstrate that feedback from the serotonergic dorsal paired medial neuron onto the KCs directly influences the length of the KC eligibility trace, making it a candidate mechanism for associative learning with a delayed US. Appetitive and aversive trace conditioning experiments have been conducted with larvae [29, 65, 66] and adult flies and other insects [74, 81–83]. In all of these experiments where the CS is presented before the US demonstrate that longer inter-stimulus intervals abolish learning of the CS-US association when no KC odor representation persists during the reinforcement period. In the cases of shorter intervals, the experimental data is not entirely conclusive. Either the odor preference was higher for partial or no overlap, compared with complete overlap [29, 83] or highest for complete overlap [51, 66, 74].

We also looked at the extent of reinforcement generalization to novel odors. Experiments have

10

shown that associations between an odor and reinforcement generalize, to a varying extent, to other odors, as shown in experiments [63, 84]. Previous modeling experiments have also shown that reinforcement generalization depends on odor similarity in adult insects [48, 85–87]. In our larval model, we also demonstrate both generalization to other odors, as well as a loss in strength, compared to the training odor (fig. 3 A). We also show that the extent of the generalization depends on the similarity of the training and test odor, as measured by the overlap of the input patterns (fig. 1 D). The larval pathway with its relatively small coding space [53, 55] might be especially prone to such poor discriminative abilities.

0.1 Model predictions for behavioral experiments

Our approach targets two hypotheses: Firstly, symmetrical inhibitory and excitatory feedback from MBONs to DANs should yield a circuit capable of saturating learning curves as predicted by the RW model [8], due to PE [11] driving the learning process, which has also been suggested by previous models [20, 48–51]. Secondly, saturating learning curves, driven by PE should translate into (simulated) animal behavior, when comparing different training duration and intensities of reinforcement. We were able to test these hypotheses in model experiments, on the level of MB readout (behavioral bias, eqn. 4, fig. 2, 4)) and through the comparison of animal and simulated behavior of artificial larvae (fig. 5). While the simulation results fit nicely with the real larval behavior in an experiment with a varied training duration ([1], fig. 5), ultimately, the role of MBON>DAN feedback needs to be tested in behavioral animal experiments, directly manipulating this feedback. Some specific predictions that could be tested in such experiments are:

- Learning curves of individual animals should saturate over time when KC>MBON feedback is intact.

- When the MBON>DAN feedback is removed after some training, the learning curve should increase with a steeper slope and might not saturate.

- Increasing or decreasing the intensity of the odor or the reinforcement should lead to saturation on a higher or lower level, respectively.

- The removal of the KC>MBON feedback should weaken or abolish the saturation of the learning curve over time.

Based on our modeling results, we support the idea that the error computation between the prediction and reality of reinforcement is done in the DANs and relies on MBON>DAN feedback. Our hypotheses for experiments are based on this assumption. Nevertheless, some saturation, that is not based on PE, might still occur, even if MBON input to DANs is removed. The entire MB circuitry consists of many more elements than our model implementation and would presumably have additional mechanisms to ensure homeostatic balance and continued MBON input, potentially leading to some weaker form of saturation in the learning pro.

11

Comparison with other MB models

Models of the learning in the MB, based on plasticity at the MB output synapses, without PE coding, have been around for some time, both for *Drosophila* [85, 87–89] and other insects [86, 90]. In all of these models, plasticity is mediated by the activity of modulatory neurons (e.g., dopaminergic), coinciding with either KC [86, 87] or coordinated KC and MBON activity [85, 88, 89]. These models can perform associative learning of a stimulus, paired with reinforcement [85–89], as well as more complicated forms of learning such as second order conditioning [89] and matching to sample [88] or reinforcement generalization tasks, the extent of which depends on the stimulus similarity [85, 86]. Additionally, some models were successfully tested in patterning tasks [85, 86], where combinations of stimuli are reinforced, while their individual components are not or vice versa. Models in which synaptic plasticity is driven not solely by the activity of modulatory neurons, but by a prediction error signal lend themselves to studying the evolution of learning over time (either over several trials, or in a continuous manner), and its dependency on the learning history. We hypothesize that such mechanisms for PE coding in the MB involve the modulatory DANs [19, 20, 32, 34, 41–44] and are based on MBON feedback to the DANs, serving as a manifestation of previous learning. Recently a number of modeling approaches have targeted the idea of PE coding in DANs in the adult *Drosophila* [48–51] as well as in the larval MB [20]. In these models, some form of MBON>DAN feedback is implemented, allowing these models to fulfill some of the predictions of the PE theory [8, 11]. One of the most fundamental predictions is the saturation of the learning curve across time, as the prediction error decreases, demonstrated in a trial-based manner in some of those models [48–51] as well as the loss of an acquired association [20, 48–50]. Some of the previously published models include mechanisms for either permanent loss of the association in memory or extinction (parallel associations in memory). Within the MB circuitry, the formation of a parallel extinction memory involves an additional DAN of opposite valence [20, 48–50], whereas complete loss is implemented as a process of changing the KC>MBON weights in the opposite direction of the learning process [51, 89], as done in our model. Additionally, some of these models capture temporal dynamics of learning experiments to some extent by utilizing eligibility traces in the KC>MBON synapses [20, 50, 51], to our understanding, none have tested these predictions in continuous experiments with spiking dynamics. Therefore, beyond the scope of these contributions, we implemented PE coding mechanistically in a fully spiking network equipped with synaptic eligibility that we train and test in continuous experiments to allow for the assessment of dynamic change in the model's odor preference. In combination with a time-continuous behavioral simulation [52] during memory retention tests, this allowed for straightforward comparison with larval experiments.

Prediction error coding is not the only mechanism discussed in the literature to explain such phenomena in learning. Gkanias et al. tested a PE-based learning rule against a different dopamine-based learning rule that dos not require the presence of the CS as a reference point for expected reinforcement [87] in a more complex circuit model consisting of a number of interconnected micro-circuits. They show that both methods can produce a saturating learning curve across trials. Their

12

alternative learning rule, embedded in a multi-compartment structure of the MB can also explain extinction, blocking, and second order conditioning, by relying on interactions between different MBONs and DANs that encode different memory processes.

Outlook

Some experimentally observed effects in insect learning can not be captured by the RW model [8] and are thus not targeted by our model implementation. Among them are CS and US pre-exposure effects [91–94] that might be explained by changes either in attention to the CS or habituation to the CS or the US, caused by prolonged exposure prior to training, rather than changes in associative strength (for a review see [95]). Also interesting, but not directly predicted by the RW model [8] is the experimental observation of second order conditioning in adult *Drosophila* [96–99], where a second CS2 is paired with the CS, after this CS has acquired an association with the US. Through the CS2-CS pairing without the US, the CS2 acquires predictive power of the US. Different mechanisms have been proposed to be involved in causing this effect [98, 100]. Among them is an excitatory synaptic KC>DAN connection, strengthened during first order conditioning, that would allow the KC odor representation to activate the DAN as a substitute for reinforcement during the CS2-CS pairing. Exploring this phenomenon using network models could yield valuable insights into the *Drosophila* circuit, as well as aid in our general understanding of PE coding. Insect experiments have provided mixed evidence for other phenomena that can be predicted from the RW model, such as blocking [101–104] and hints at conditioned inhibition [105–107] that would be interesting to investigate. Furthermore, expanding the model to include different MB output compartments would offer a perspective to explore parallel associations regarding the same stimulus [41]. This could enable temporary loss of the learned association, while simultaneously retaining parallel memory for recovery (extinction vs. forgetting). Ultimately more possible directions arise from the major benefit of using a spiking model, which offers the potential to conduct experiments at high temporal resolution, instead of in a trial-based manner [20, 48–50]. In a future closed-loop approach that connects our continual learning MB model with the locomotory model in the full temporal resolution, we intend to simulate a behaving agent to investigate the temporal dynamics of adaptive behavior in analogy to the tracking experiments of real larva [73, 108–111].

Methods

Network model

All neurons are modeled as leaky integrate-and-fire neurons with conductance-based synapses. They elicit a spike, whenever the threshold V_T is crossed(parameters provided in table S1). Each neuronal membrane potential v_i is reset to the resting potential V_r whenever a spike occurs, followed by an absolute refractory period of $2\,\mathrm{ms}$, during which the neuron does not integrate any inputs. Any

neuron from a given population (v^O,v^P,v^L,v^K,v^A,v^M,v^D) is governed by the respective equation for ORNs, PNs, LNs, KCs, APL, MBONs and DANs (eqs. (5) to (11), fig. 1 A). Depending on the neuron type, in addition to a leak conductance g_L, the equations consist of excitatory g_e and inhibitory synaptic input g_i. In the case of the DANs, one excitatory $g_e^{M\mp D}(E_E - v_i^D)$ and inhibitory $g_i^{M\pm D}(E_I - v_i^D)$ input represent the two types of MBON feedback for the reward and punishment encoding DAN, respectively. An additional spike-triggered adaptation conductance was implemented for ORNs, KCs, MBONs, and DANs (eqn. 12, [64]), in accordance with our current knowledge of the adaptive nature of ORNs in the larva [112] and the adult fly [113, 114]. Adaptation in KCs has so far only been demonstrated in other insects [115, 116]. In the model of these neurons, the adaptation conductance g_{Ia} is increased with every spike and decays over time with τ_{Ia}. The mechanism of synaptic plasticity is described in the results section (Learning through KC>MBON plasticity).

$$C_m\frac{d}{dt}v_i^O = g_L^O(E_L^O - v_i^O) + g_e^{InputO}(E_E - v_i^O) - g_{Ia}(E_{Ia} - v_i^O) \tag{5}$$

$$C_m\frac{d}{dt}v_i^L = g_L^L(E_L^L - v_i^L) + g_e^{OL}(E_E - v_i^L) \tag{6}$$

$$C_m\frac{d}{dt}v_i^P = g_L^P(E_L^P - v_i^P) + g_e^{OP}(E_E - v_i^P) - g_i^{LP}(E_I - v_i^P) \tag{7}$$

$$C_m\frac{d}{dt}v_i^K = g_L^K(E_L^K - v_i^K) - g_i^{APLK}(E_I - v_i^K) + g_e^{PK}(E_E - v_i^K) - g_{Ia}(E_{Ia} - v_i^K) \tag{8}$$

$$C_m\frac{d}{dt}v_i^A = g_L^A(E_L^A - v_i^A) + g_e^{KAPL}(E_E - v_i^A) \tag{9}$$

$$C_m\frac{d}{dt}v_i^M = g_L^M(E_L^M - v_i^M) + g_e^{KM}(E_E - v_i^M) \tag{10}$$

$$C_m\frac{d}{dt}v_i^D = g_L^D(E_L^D - v_i^D) - g_i^{M\pm D}(E_I - v_i^D) + g_e^{M\mp D}(E_E - v_i^D) + g_e^{InputD}(E_E - v_i^D) \tag{11}$$

$$\frac{d}{dt}g_{Ia} = -\frac{g_{Ia}}{\tau_{Ia}}. \tag{12}$$

All code for the model implementation is accessible via
https://github.com/nawrotlab/PEcodingDosophilaMB

We based our circuit model on the larval connectome both in terms of connectivity as well as numbers of neurons in each population [20, 53, 54] and introduced simplifications to support the mechanic investigation of the MBON>DAN feedback circuit and its role in PE coding and excluded a number of connections that have been demonstrated in the larva. Due to limited availability of

14

469 anatomical, functional, and behavioral data most of our circuit implementation is based on the first

470 instar larva [20, 53, 54], while the information on the APL connectivity within the circuit originates

471 from studies on the third instar larva [29]. Behavioral experiments used for comparison with our

472 simulation results were also performed with third instar larvae [1, 29, 66]. We demonstrate that our

473 model based on the less developed circuit in the first instar larva is sufficient to reproduce animal

474 behavior as observed in the older animals. From the anatomy of the first instar larva we excluded

475 DAN>KC [54] and DAN>MBON synapses [54] that may play an additional role in learning-induced

476 plasticity at KC>MBON synapses [54], the details of which are not fully known. Instead, we induce

477 plasticity purely via the simulation of a neuromodulatory effect of the DANs onto the KC>MBON

478 synapses ([54]). We also neglect recurrent interactions among KC themselves [54]. Many of these

479 interactions affect KC that encode different sensory modalities, which are not included in our purely

480 olfactory model. Furthermore, we simplified the connectivity between LNs and PNs [53] and between

481 PNs and KCs to $2 - 6$ PN inputs per KC, which excludes the set of KCs in the larva that receives

482 exclusive input from only one PN [54]. This modification supported model robustness with respect

483 to odor encoding within the small set of 72 KCs. Finally, from the population of ≈ 25 larval MBONs

484 we only modeled two and correspondingly adapted KC>MBON synapses to provide both MBONs

485 with input from all KCs.

Sparse odor representation

487 We implemented four mechanisms supporting population- and temporal sparseness in the MB odor

488 representation [64]. Population sparseness is defined as the activation of only a small subset of

489 neurons by any given input [117]. In this circuit population sparseness is enhanced through lateral

490 inhibition (via LNs), inhibitory APL feedback, and the divergent connectivity from PNs to a larger

491 number of KCs [64]. Temporal sparseness indicates that an individual neuron responds with only a

492 few spikes to a specific stimulus configuration [118–120], which supports encoding dynamic changes

493 in the sensory environment [121, 122]. In our model temporal sparseness is facilitated by spike

494 frequency adaptation, an adaptive process to prolonged stimulus exposure, in ORNs and KCs and

495 by inhibitory feedback via the APL[64].

Sensory input

497 In the olfactory pathway of larval *Drosophila* any odor activates up to $\approx 1/3$ of ORNs, depending on

498 its concentration [112, 123]. We implemented receptor input with stochastic point processes to ORNs

499 via synapses to mimic the noise in a transduction process at the receptors. Each of the 21 receptor

500 inputs is modeled according to a gamma process (shape parameter k= 3). The spontaneous firing

501 rate of larval ORNs has been measured in the range of $0.2 - 7.9$ Hz, depending strongly on odor and

502 receptor type [123, 124]. ORNs in our model exhibit an average spontaneous firing rate of 8.92Hz

503 (std=0.2). We constructed realistic olfactory input across the ORN population for amylacetate and

15

3-octanol by estimating ORN spike frequency from the calcium signals measured in the receptor neurons [112] (dilution of 10^{-4} [112]), ensuring the spike rates would not exceed the rates reported by [123]. They showed that using an even stronger odor concentration (dilution 10^{-2}) ORN never exceeded a frequency of 200Hz. Due to the lower concentration used for amylacetate and 3-octanol (fig. 1 D) [112] in our experiments and because Kreher et al, 2005 measured only the first 0.5s after odor onset when the effects of spike frequency in ORNs are the weakest (leading to higher spike rates) we decided to use a maximum of 150Hz in odor activated ORNs. After generating the gamma process realizations we clipped multiple spikes occurring in each time step of the simulation discarding all but the first spike in each time step. Similar to the odor input, the presence of either reward or punishment in the experimental context was implemented as input to the DAN_+/DAN_-. Regular gamma spike trains ($k = 10$) were generated and clipped for the odor input.

To assess the effects of odor similarity on generalization we in addition created four artificial odors (A,B,C,D) (fig. 1 D) and quantified the pair-wise distances in ORN coding space using the cosine distance (eqn. 13), where vectors a and b each represent the input spike rate of two odors.

$$D_{cos} = 1 - \frac{\sum_{i=1}^{n} a_i \cdot b_i}{\sqrt{\sum_{i=1}^{n} a_i^2} \cdot \sqrt{\sum_{i=1}^{n} b_i^2}}. \tag{13}$$

The cosine distance between odors A and B equals 0.21, 0.77 between odors A and C, and 0.99 between odors A and D. The comparison of amylacetate and 3-octanol yields a distance of 0.16.

Experimental protocols

The experiments reported here belong to one of three categories. The first was performed to provide insight into the model and the effects of specific circuit functions on synaptic plasticity, and prediction error coding. To this end, we used amylacetate as the primary odor input. We varied the intensity of the reward via the frequency of gamma spike train, provided as input into the DAN_+ (either 500Hz or 550Hz, resulting in an average output spike rate of 33.11/39.14Hz), and the learning rate α(0.6nS or 0.8nS). Additionally, MBON>DAN feedback was either enabled or disabled (fig. 1 A).

Experiments belonging to the second category were designed to replicate larva lab experiments to allow for a direct comparison with our model results. With these comparisons, we aim to validate the model and show to what extent our assumptions about the circuit functions allow us to recreate experimental data (fig. 5). Replicating lab experiments also provide more insights into the circuit mechanisms and offers alternative interpretations of the phenomena observed in data from animal experiments. Our implementations of the lab experiments were set up following the general procedure described in the Maggot Learning Manual [125]. Regardless of the specific protocol used in different experiments, larvae are placed into Petri dishes in groups of 30 animals. They are allowed to move around freely on the substrate that contains reinforcing substances, such as sugar or bitter tastants. During the entire time, they are subjected to specific odorants, emitted from two small containers in the dish to create permanent and uniformly distributed odor exposure within the dish. In the

16

analogy of the experimental setting, in our simulated experiments, each model instance is trained individually through the concurrent presentation of olfactory stimulation and reward. One-minute intervals with only baseline ORN stimulation were included between training trials to simulate the time needed in the lab experiments for transferring larvae between Petri dishes. Unless otherwise specified and test phases refer to 3 min, during which only odors are presented. All simulations were implemented in the network simulator Brian2 [126].

Realistic modeling of larval locomotion

Behavior during the testing phase of the olfactory learning experiment is simulated via the freely available python-based simulation platform Larvaworld (https://github.com/nawrotlab/larvaworld, [52]). A group of 30 virtual larvae is placed with random initial orientation around the center of a 100 mm diameter Petri dish and left to freely behave for 3 minutes. The previously conditioned odor is placed at one side of the dish, 10 mm from the arena's boundary. Each larva features a bi-segmental body, supervised by a layered control architecture [52]. The basic layer of the control architecture is a locomotory model, capable of realistic autonomous exploration behavior. It consists of two coupled oscillators, one of which represents the crawling apparatus that generates forward velocity oscillations, resembling consecutive peristaltic strides [52]. The other oscillator generates alternating left and right lateral bending, manifested as oscillations of angular velocity [127]. The crawling and the bending oscillators are coupled via phase-locked suppression of lateral bending to capture the bend dependency on the stride-cycle phase during crawling (weathervaning). Finally, intermittent crawling is achieved by a superimposed intermittency module that generates alternating epochs of crawling and stationary pauses, with more headcasts for orientation during the latter [52].

Modulation of behavior due to sensory stimulation is introduced at the second, reactive layer of the control architecture. An odor signal can transiently alter both, the amplitude and frequency of the lateral bending oscillator, which biases free exploration towards approach or avoidance along an olfactory chemical gradient. This modulation of behavior is directly influenced via top-down signaling from the third, adaptive layer of the control architecture. In our approach, the spiking MB model populates the adaptive layer and its learning-dependent output, defined as the behavioral bias BB (i.e. the difference in MBON firing rates, eqn. 4), provides the top-down signal [36]. We formalize the gain of behavioral modulation as

$$G = g \cdot BB. \tag{14}$$

which is directly proportional to the behavioral bias and the additional proportionality factor $g = 0.5$.

A set of $10 * 30$ trained MB model instances is used to generate 10 groups of 30 simulated larvae. The preference index and the performance index [1] for these simulations are illustrated in figure 5.

Preference indices (Pref) are computed individually for the paired and the unpaired experiments

17

[1], based on the number of animals on each side (odor vs. empty) of the Petri dish at the end of the test phase.

$$Pref = \frac{count_{odor} - count_{no\ odor}}{count_{odor} + count_{no\ odor}}. \tag{15}$$

The Performance indices (PI) are computed from the preference indices of the paired and unpaired experiments [1].

$$PI = \frac{Pref_{paired} - Pref_{unpaired}}{2}. \tag{16}$$

18

Acknowledgements

References

[1] Aliće Weiglein, Florian Gerstner, Nino Mancini, Michael Schleyer, and Bertram Gerber. "One-trial learning in larval Drosophila". In: *Learning & Memory* 26.4 (2019), pp. 109–120. DOI: /10.1101/lm.049106.118.

[2] Martin Heisenberg. "Pattern recognition in insects". In: *Current Opinion in Neurobiology* 5.4 (1995), pp. 475–481. ISSN: 09594388. DOI: 10.1016/0959-4388(95)80008-5.

[3] Randolf Menzel. "The honeybee as a model for understanding the basis of cognition". In: *Nature Reviews Neuroscience* 13.11 (2012), pp. 758–768. DOI: 10.1038/nrn3357.

[4] Yoshinori Aso, Daisuke Hattori, Yang Yu, Rebecca M Johnston, Nirmala A Iyer, Teri-TB Ngo, Heather Dionne, LF Abbott, Richard Axel, Hiromu Tanimoto, et al. "The neuronal architecture of the mushroom body provides a logic for associative learning". In: *elife* 3 (2014), e04577. DOI: 10.7554/eLife.04577.

[5] Toshihide Hige, Yoshinori Aso, Mehrab N Modi, Gerald M Rubin, and Glenn C Turner. "Heterosynaptic plasticity underlies aversive olfactory learning in Drosophila". In: *Neuron* 88.5 (2015), pp. 985–998. DOI: /10.1016/j.neuron.2015.11.003.

[6] David Owald and Scott Waddell. "Olfactory learning skews mushroom body output pathways to steer behavioral choice in Drosophila". In: *Current opinion in neurobiology* 35 (2015), pp. 178–184. DOI: /10.1016/j.conb.2015.10.002.

[7] Martin F Strube-Bloss, Martin P Nawrot, and Randolf Menzel. "Mushroom body output neurons encode odor–reward associations". In: *Journal of Neuroscience* 31.8 (2011), pp. 3129–3140. DOI: 10.1523/JNEUROSCI.2583-10.2011.

[8] Robert A Rescorla and Allan R Wagner. "A theory of Pavlovian conditioning : Variations in the effectiveness of reinforcement and non- reinforcement". In: *Classical Conditioning 2: Current Theory and Research*. Ed. by William F. Black, Abraham H.; Prokasy. January 1972. New York, NY: Appelton-century-Crofts, 1972. Chap. 3, pp. 64–99.

[9] Peter D Balsam and Randy C Gallistel. "Temporal maps and informativeness in associative learning". In: *Trends in neurosciences* 32.2 (2009), pp. 73–78. DOI: 10.1016/j.tins.2008.10.004.

[10] Peter S Kaplan. "Importance of relative temporal parameters in trace autoshaping: From excitation to inhibition." In: *Journal of Experimental psychology: Animal behavior processes* 10.2 (1984), p. 113.

[11] Leon Kamin. "Predictability, surprise, attention and conditioning". In: *Punsihment and aversive behavior*. Ed. by Byron A. Campbell. New York: Appleton-Century-Crofts, 1969, 279–298.

[12] Wolfram Schultz. "Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey". In: *Journal of neurophysiology* 56.5 (1986), pp. 1439–1461. DOI: 10.1152/jn.1986.56.5.1439.

[13] Genela Morris, David Arkadir, Alon Nevet, Eilon Vaadia, and Hagai Bergman. "Coincident but Distinct Messages of Midbrain Dopamine and Striatal Tonically Active Neurons". In: *Neuron* 43.1 (2004), pp. 133–143. ISSN: 0896-6273. DOI: 10.1016/J.NEURON.2004.06.012.

[14] Yoriko Takikawa, Reiko Kawagoe, and Okihide Hikosaka. "A possible role of midbrain dopamine neurons in short- and long-term adaptation of saccades to position-reward mapping". In: *Journal of Neurophysiology* 92.4 (2004), pp. 2520–2529. DOI: 10.1152/jn.00238.2004.

[15] Takemasa Satoh, Sadamu Nakai, Tatsuo Sato, and Minoru Kimura. "Correlated Coding of Motivation and Outcome of Decision by Dopamine Neurons". In: *Journal of Neuroscience* 23.30 (2003), pp. 9913–9923. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.23-30-09913.2003.

[16] Jeremiah Y Cohen, Sebastian Haesler, Linh Vong, Bradford B Lowell, and Naoshige Uchida. "Neuron-type-specific signals for reward and punishment in the ventral tegmental area". In: *Nature* 482.7383 (2012), pp. 85–88. DOI: 10.1038/nature10754.

[17] Ariel Y Deutch, See-Ying Tam, and Robert H Roth. "Footshock and conditioned stress increase 3, 4-dihydroxyphenylacetic acid (DOPAC) in the ventral tegmental area but not substantia nigra". In: *Brain research* 333.1 (1985), pp. 143–146. DOI: 10.1016/0006-8993(85)90134-9.

[18] Scott Waddell. "Reinforcement signalling in Drosophila; dopamine does it all after all". In: *Current opinion in neurobiology* 23.3 (2013), pp. 324–329. DOI: 10.1016/j.conb.2013.01.005.

[19] Thomas Riemensperger, Thomas Völler, Patrick Stock, Erich Buchner, and André Fiala. "Punishment prediction by dopaminergic neurons in Drosophila". In: *Current Biology* 15.21 (2005), pp. 1953–1960. ISSN: 09609822. DOI: 10.1016/j.cub.2005.09.042.

[20] Claire Eschbach, Akira Fushiki, Michael Winding, Casey M Schneider-Mizell, Mei Shao, Rebecca Arruda, Katharina Eichler, Javier Valdes-Aleman, Tomoko Ohyama, Andreas S Thum, Bertram Gerber, Richard D Fetter, James W Truman, Ashok Litwin-Kumar, Albert Cardona, and Marta Zlatic. "Recurrent architecture for adaptive regulation of learning in the insect brain". In: *Nature Neuroscience* 23.4 (2020), pp. 544–555. ISSN: 15461726. DOI: 10.1038/s41593-020-0607-9.

[21] Martin Schwaerzel, Maria Monastirioti, Henrike Scholz, Florence Friggi-Grelin, Serge Birman, and Martin Heisenberg. "Dopamine and octopamine differentiate between aversive and appetitive olfactory memories in Drosophila". In: *Journal of Neuroscience* 23.33 (2003), pp. 10495–10502. DOI: /10.1523/JNEUROSCI.23-33-10495.2003.

[22] Dale Corbett and Roy A. Wise. "Intracranial self-stimulation in relation to the ascending dopaminergic systems of the midbrain: A moveable electrode mapping study". In: *Brain Research* 185.1 (1980), pp. 1–15. DOI: 10.1016/0006-8993(80)90666-6.

[23] William R. Stauffer, Armin Lak, Aimei Yang, Melodie Borel, Ole Paulsen, Edward S. Boyden, and Wolfram Schultz. "Dopamine Neuron-Specific Optogenetic Stimulation in Rhesus Macaques". In: *Cell* 166.6 (2016), 1564–1571.e6. DOI: 10.1016/J.CELL.2016.08.024.

[24] Roy A Wise and P-P Rompre. "Brain dopamine and reward". In: *Annual review of psychology* 40.1 (1989), pp. 191–225.

[25] Ilana B. Witten, Elizabeth E. Steinberg, Soo Yeun Lee, Thomas J. Davidson, Kelly A. Zalocusky, Matthew Brodsky, Ofer Yizhar, Saemi L. Cho, Shiaoching Gong, Charu Ramakrishnan, Garret D. Stuber, Kay M. Tye, Patricia H. Janak, and Karl Deisseroth. "Recombinase-driver rat lines: Tools, techniques, and optogenetic application to dopamine-mediated reinforcement". In: *Neuron* 72.5 (2011), pp. 721–733. DOI: 10.1016/j.neuron.2011.10.028.

[26] Chang Liu, Pierre Yves Plaaais, Nobuhiro Yamagata, Barret D. Pfeiffer, Yoshinori Aso, Anja B. Friedrich, Igor Siwanowicz, Gerald M. Rubin, Thomas Preat, and Hiromu Tanimoto. "A subset of dopamine neurons signals reward for odour memory in Drosophila". In: *Nature 2012 488:7412* 488.7412 (2012), pp. 512–516. ISSN: 1476-4687. DOI: 10.1038/nature11304.

[27] Adam Claridge-Chang, Robert D. Roorda, Eleftheria Vrontou, Lucas Sjulson, Haiyan Li, Jay Hirsh, and Gero Miesenböck. "Writing Memories with Light-Addressable Reinforcement Circuitry". In: *Cell* 139.2 (2009), pp. 405–415. ISSN: 0092-8674. DOI: 10.1016/J.CELL.2009.08.034.

[28] Christian König, Afshin Khalili, Mathangi Ganesan, Amrita P Nishu, Alejandra P Garza, Thomas Niewalda, Bertram Gerber, Yoshinori Aso, and Ayse Yarali. "Reinforcement signaling of punishment versus relief in fruit flies". In: *Learning & Memory* 25.6 (2018), pp. 247–257. DOI: 10.1101/lm.047308.118.

[29] Timo Saumweber, Astrid Rohwedder, Michael Schleyer, Katharina Eichler, Yi-Chun Chen, Yoshinori Aso, Albert Cardona, Claire Eschbach, Oliver Kobler, Anne Voigt, Archana Durairaja, Nino Mancini, Marta Zlatic, James W Truman, Andreas S Thum, and Bertram Gerber. "Functional architecture of reward learning in mushroom body extrinsic neurons of larval Drosophila". In: *Nature communications* 9.1 (2018), p. 1104. DOI: 10.1038/s41467-018-03130-1.

[30] Michael Schleyer, Aliće Weiglein, Juliane Thoener, Martin Strauch, Volker Hartenstein, Melisa Kantar Weigelt, Sarah Schuller, Timo Saumweber, Katharina Eichler, Astrid Rohwedder, et al. "Identification of dopaminergic neurons that can both establish associative memory and acutely terminate its behavioral expression". In: *Journal of Neuroscience* 40.31 (2020), pp. 5990–6006. DOI: 10.1523/JNEUROSCI.0290-20.2020.

[31] Christian Schroll, Thomas Riemensperger, Daniel Bucher, Julia Ehmer, Thomas Völler, Karen Erbguth, Bertram Gerber, Thomas Hendel, Georg Nagel, Erich Buchner, and André Fiala. "Light-induced activation of distinct modulatory neurons triggers appetitive or aversive learning in Drosophila larvae". In: *Current biology : CB* 16 (17 2006), pp. 1741–1747. DOI: 10.1016/j.cub.2006.07.023.

[32] Adithya E Rajagopalan, Ran Darshan, Karen L Hibbard, James E Fitzgerald, and Glenn C Turner. "Reward expectations direct learning and drive operant matching in Drosophila". In: *bioRxiv* (2022), pp. 2022–05. DOI: 10.1101/2022.05.24.493252.

[33] Clare E Hancock, Vahid Rostami, El Yazid Rachad, Stephan H Deimel, Martin P Nawrot, and André Fiala. "Visualization of learning-induced synaptic plasticity in output neurons of the Drosophila mushroom body γ-lobe". In: *Scientific Reports* 12.1 (2022), p. 10421. DOI: 10.1038/s41598-022-14413-5.

[34] Yoshinori Aso and Gerald M Rubin. "Dopaminergic neurons write and update memories with cell-type-specific rules". In: *Elife* 5 (2016), e16135. DOI: /10.7554/eLife.16135.

[35] Birgit Michels, Yi-chun Chen, Timo Saumweber, Dushyant Mishra, Hiromu Tanimoto, Benjamin Schmid, Olivia Engmann, and Bertram Gerber. "Cellular site and molecular mode of synapsin action in associative learning". In: *Learning & Memory* 18.5 (2011), pp. 332–344. DOI: 10.1101/lm.2101411.

[36] Claire Eschbach, Akira Fushiki, Michael Winding, Bruno Afonso, Ingrid V Andrade, Benjamin T Cocanougher, Katharina Eichler, Ruben Gepner, Guangwei Si, Javier Valdes-Aleman, et al. "Circuits for integrating learned and innate valences in the insect brain". In: *Elife* 10 (2021), e62567. DOI: /10.7554/eLife.62567.

[37] Wolfram Schultz, Paul Apicella, and Tomas Ljungberg. "Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task". In: *Journal of neuroscience* 13.3 (1993), pp. 900–913. DOI: 10.1523/JNEUROSCI.13-03-00900.1993.

22

[38] Wolfram Schultz, Peter Dayan, and P Read Montague. "Getting formal with dopamine and reward". In: *Science* 275.5306 (1997), pp. 1593–1599. DOI: 10.1126/science.275.5306.1593.

[39] Kareem A Zaghloul, Justin A Blanco, Christoph T Weidemann, Kathryn McGill, Jurg L Jaggi, Gordon H Baltuch, and Michael J Kahana. "Human substantia nigra neurons encode unexpected financial rewards". In: *Science* 323.5920 (2009), pp. 1496–1499. DOI: 10.1126/science.1167342.

[40] Wei-Xing Pan, Robert Schmidt, Jeffery R Wickens, and Brian I Hyland. "Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network". In: *Journal of Neuroscience* 25.26 (2005), pp. 6235–6242. DOI: /10.1523/JNEUROSCI.1478-05.2005.

[41] Johannes Felsenberg, Pedro F Jacob, Thomas Walker, Oliver Barnstedt, Amelia J Edmondson-Stait, Markus W Pleijzier, Nils Otto, Philipp Schlegel, Nadiya Sharifi, Emmanuel Perisse, et al. "Integration of parallel opposing memories underlies memory extinction". In: *Cell* 175.3 (2018), pp. 709–722. DOI: /10.1016/j.cell.2018.08.021.

[42] Makoto Mizunami, Kanta Terao, and Beatriz Alvarez. "Application of a prediction error theory to Pavlovian conditioning in an insect". In: *Frontiers in psychology* 9 (2018), p. 1272. DOI: /10.3389/fpsyg.2018.01272.

[43] Maria E. Villar, Miguel Pavão-Delgado, Marie Amigo, Pedro F. Jacob, Nesrine Merabet, Anthony Pinot, Sophie A. Perry, Scott Waddell, and Emmanuel Perisse. "Differential coding of absolute and relative aversive value in the Drosophila brain". In: *Current Biology* 32.21 (2022), 4576–4592.e5. ISSN: 0960-9822. DOI: 10.1016/j.cub.2022.08.058.

[44] Ilona C Grunwald Kadow and David Owald. "Decision making: Dopaminergic neurons for better or worse". In: *Current Biology* 32.21 (2022), R1237–R1240. DOI: 10.1016/j.cub.2022.09.043.

[45] Louis K Scheffer, C Shan Xu, Michal Januszewski, Zhiyuan Lu, Shin-ya Takemura, Kenneth J Hayworth, Gary B Huang, Kazunori Shinomiya, Jeremy Maitlin-Shepard, Stuart Berg, et al. "A connectome and analysis of the adult Drosophila central brain". In: *Elife* 9 (2020), e57443. DOI: /10.7554/eLife.57443.

[46] Nils Otto, Markus W Pleijzier, Isabel C Morgan, Amelia J Edmondson-Stait, Konrad J Heinz, Ildiko Stark, Georgia Dempsey, Masayoshi Ito, Ishaan Kapoor, Joseph Hsu, et al. "Input connectivity reveals additional heterogeneity of dopaminergic reinforcement in Drosophila". In: *Current Biology* 30.16 (2020), pp. 3200–3211. DOI: 10.1016/j.cub.2020.05.077.

[47] Michael Winding, Benjamin D Pedigo, Christopher L Barnes, Heather G Patsolic, Youngser Park, Tom Kazimiers, Akira Fushiki, Ingrid V Andrade, Avinash Khandelwal, Javier Valdes-Aleman, et al. "The connectome of an insect brain". In: *Science* 379.6636 (2023), eadd9330. DOI: 10.1126/science.add933.

[48] James E.M. Bennett, Andrew Philippides, and Thomas Nowotny. "Learning with reinforcement prediction errors in a model of the Drosophila mushroom body". In: *Nature Communications* 12.1 (2021). ISSN: 20411723. DOI: 10.1038/s41467-021-22592-4.

[49] Magdalena Springer and Martin P Nawrot. "A mechanistic model for reward prediction and extinction learning in the fruit fly". In: *Eneuro* 8.3 (2021). DOI: /10.1523/ENEURO.0549-20.2021.

[50] Linnie Jiang and Ashok Litwin-Kumar. "Models of heterogeneous dopamine signaling in an insect learning and memory center". In: *PLoS Computational Biology* 17.8 (2021), e1009205. DOI: /10.1371/journal.pcbi.1009205.

[51]   Chang Zhao, Yves F Widmer, Sören Diegelmann, Mihai A Petrovici, Simon G Sprecher, and Walter Senn. "Predictive olfactory learning in Drosophila". In: *Scientific reports* 11.1 (2021), pp. 1–17. DOI: 10.1038/s41598-021-85841-y.

[52]   Panagiotis Sakagiannis, Anna-Maria Jürgensen, and Martin P Nawrot. "A realistic locomotory model of Drosophila larva for behavioral simulations". In: *bioRxiv* (2021). DOI: 10.1101/2021.07.07.451470.

[53]   Matthew E Berck, Avinash Khandelwal, Lindsey Claus, Luis Hernandez-Nunez, Guangwei Si, Christopher J Tabone, Feng Li, James W Truman, Rick D Fetter, Matthieu Louis, et al. "The wiring diagram of a glomerular olfactory system". In: *Elife* 5 (2016). DOI: /10.7554/eLife.14859.

[54]   Katharina Eichler, Feng Li, Ashok Litwin-Kumar, Youngser Park, Ingrid Andrade, Casey M Schneider-Mizell, Timo Saumweber, Annina Huser, Claire Eschbach, Bertram Gerber, Richard D Fetter, James W Truman, Carey E Priebe, L F Abbott, Andreas S Thum, Marta Zlatic, and Albert Cardona. "The complete connectome of a learning and memory centre in an insect brain". In: *Nature* 548.7666 (2017), pp. 175–182. ISSN: 1476-4687. DOI: 10.1038/nature23455.

[55]   Africa Couto, Mattias Alenius, and Barry J Dickson. "Molecular, anatomical, and functional organization of the Drosophila olfactory system". In: *Current Biology* 15.17 (2005), pp. 1535–1547. DOI: /10.1016/j.cub.2005.07.034.

[56]   Leslie B Vosshall and Reinhard F Stocker. "Molecular architecture of smell and taste in Drosophila". In: *Annual review of neuroscience* 30 (2007), pp. 505–533. DOI: 10.1146/annurev.neuro.30.051606.094306.

[57]   Yoshinori Aso, Divya Sitaraman, Toshiharu Ichinose, Karla R Kaun, Katrin Vogt, Ghislain Belliart-Guérin, Pierre-Yves Plaçais, Alice A Robie, Nobuhiro Yamagata, Christopher Schnaitmann, et al. "Mushroom body output neurons encode valence and guide memory-based action selection in Drosophila". In: *Elife* 3 (2014), e04580. DOI: /10.7554/eLife.04580.

[58]   Julien Séjourné, Pierre-Yves Plaçais, Yoshinori Aso, Igor Siwanowicz, Séverine Trannoy, Vladimiros Thoma, Stevanus R Tedjakumala, Gerald M Rubin, Paul Tchénio, Kei Ito, et al. "Mushroom body efferent neurons responsible for aversive olfactory memory retrieval in Drosophila". In: *Nature neuroscience* 14.7 (2011), pp. 903–910. DOI: /10.1038/nn.2846.

[59]   Clare E Hancock, Florian Bilz, and André Fiala. "In vivo optical calcium imaging of learning-induced synaptic plasticity in Drosophila melanogaster". In: *JoVE (Journal of Visualized Experiments)* 152 (2019), e60288. DOI: 10.3791/60288.

[60]   Johannes Felsenberg, Oliver Barnstedt, Paola Cognigni, Suewei Lin, and Scott Waddell. "Re-evaluation of learned information in Drosophila". In: *Nature* 544.7649 (2017), pp. 240–244. DOI: 10.1038/nature21716.

[61]   Martin Schwaerzel, Martin Heisenberg, and Troy Zars. "Extinction antagonizes olfactory memory at the subcellular level". In: *Neuron* 35.5 (2002), pp. 951–960. DOI: 10.1016/S0896-6273(02)00832-2.

[62]   Karyn M Myers and Michael Davis. "Behavioral and neural analysis of extinction". In: *Neuron* 36.4 (2002), pp. 567–584. DOI: 10.1016/S0896-6273(02)01064-4.

[63]   Yi-chun Chen, Dushyant Mishra, Linda Schmitt, Michael Schmuker, and Bertram Gerber. "A behavioral odor similarity "space" in larval Drosophila". In: *Chemical senses* 36.3 (2011), pp. 237–249. DOI: 10.1093/chemse/bjq123.

[64] Anna-Maria Jürgensen, Afshin Khalili, Elisabetta Chicca, Giacomo Indiveri, and Martin P Nawrot. "A neuromorphic model of olfactory processing and sparse coding in the Drosophila larva brain". In: *Neuromorphic Computing and Engineering* 1.2 (2021), p. 024008. DOI: 10.1088/2634-4386/ac3ba6.

[65] Juliane Thoener, Alíce Weiglein, Bertram Gerber, and Michael Schleyer. "Optogenetically induced reward and 'frustration'memory in larval Drosophila melanogaster". In: *Journal of Experimental Biology* 225.16 (2022), jeb244565. DOI: 10.1242/jeb.244565.

[66] Alíce Weiglein, Juliane Thoener, Irina Feldbruegge, Louisa Warzog, Nino Mancini, Michael Schleyer, and Bertram Gerber. "Aversive teaching signals from individual dopamine neurons in larval Drosophila show qualitative differences in their temporal "fingerprint"". In: *Journal of Comparative Neurology* 529.7 (2021), pp. 1553–1570. DOI: /10.1002/cne.25037.

[67] Timo Saumweber, Jana Husse, and Bertram Gerber. "Innate attractiveness and associative learnability of odors can be dissociated in larval Drosophila". In: *Chemical senses* 36.3 (2011), pp. 223–235. DOI: /10.1093/chemse/bjq128.

[68] Michael Schleyer, Daisuke Miura, Teiichi Tanimura, and Bertram Gerber. "Learning the specific quality of taste reinforcement in larval Drosophila". In: *elife* 4 (2015), e04711. DOI: 10.7554/eLife.04711.

[69] Birgit Michels, Sören Diegelmann, Hiromu Tanimoto, Isabell Schwenkert, Erich Buchner, and Bertram Gerber. "A role for Synapsin in associative learning: the Drosophila larva as a study case". In: *Learning & Memory* 12.3 (2005), pp. 224–231. DOI: 10.1101/lm.92805.

[70] Mark E Bouton. "Context and behavioral processes in extinction". In: *Learning & memory* 11.5 (2004), pp. 485–494. DOI: 10.1101/lm.78804.

[71] Lingling Wang, Qi Yang, Binyan Lu, Lianzhang Wang, Yi Zhong, and Qian Li. "A behavioral paradigm to study the persistence of reward memory extinction in Drosophila". In: *Journal of genetics and genomics* 46.12 (2019), pp. 599–601. DOI: 10.1016/j.jgg.2019.11.001.

[72] Yukinori Hirano, Kunio Ihara, Tomoko Masuda, Takuya Yamamoto, Ikuko Iwata, Aya Takahashi, Hiroko Awata, Naosuke Nakamura, Mai Takakura, Yusuke Suzuki, et al. "Shifting transcriptional machinery is required for long-term memory maintenance and modification in Drosophila mushroom bodies". In: *Nature communications* 7.1 (2016), p. 13471. DOI: 10.1038/ncomms13471.

[73] Amanda Lesar, Javan Tahir, Jason Wolk, and Marc Gershow. "Switch-like and persistent memory formation in individual Drosophila larvae". In: *Elife* 10 (2021), e70317. DOI: /10.7554/eLife.70317.

[74] Tim Tully and William G Quinn. "Classical conditioning and retention in normal and mutant Drosophila melanogaster". In: *Journal of Comparative Physiology A* 157.2 (1985), pp. 263–277.

[75] Kirsa Neuser, Jana Husse, Patrick Stock, and Bertram Gerber. "Appetitive olfactory learning in Drosophila larvae: Effects of repetition, reward strength, age, gender, assay type and memory span". In: *Animal Behaviour* 69 (4 Apr. 2005), pp. 891–898. ISSN: 00033472. DOI: 10.1016/j.anbehav.2004.06.013.

[76] Dennis Mathew, Carlotta Martelli, Elizabeth Kelley-Swift, Christopher Brusalis, Marc Gershow, Aravinthan DT Samuel, Thierry Emonet, and John R Carlson. "Functional diversity among sensory receptors in a Drosophila olfactory circuit". In: *Proceedings of the National Academy of Sciences* 110.23 (2013), E2134–E2143. DOI: 10.1073/pnas.130697611.

[77] Elane Fishilevich and Leslie B Vosshall. "Genetic and functional subdivision of the Drosophila antennal lobe". In: *Current Biology* 15.17 (2005), pp. 1548–1553. DOI: 10.1016/j.cub.2005.07.066.

[78] Scott A Kreher, Dennis Mathew, Junhyong Kim, and John R Carlson. "Translation of sensory input into behavioral output via an olfactory system". In: *Neuron* 59.1 (2008), pp. 110–124. DOI: 10.1016/j.neuron.2008.06.010.

[79] Isabell Twick, John Anthony Lee, and Mani Ramaswami. "Olfactory habituation in Drosophila—odor encoding and its plasticity in the antennal lobe". In: *Progress in Brain Research* 208 (2014), pp. 3–38. DOI: 10.1016/B978-0-444-63350-7.00001-2.

[80] Jianzhi Zeng, Xuelin Li, Renzimo Zhang, Mingyue Lv, Yipan Wang, Ke Tan, Xiju Xia, Jinxia Wan, Miao Jing, Xiuning Zhang, et al. "Local 5-HT signaling bi-directionally regulates the coincidence time window for associative learning". In: *Neuron* 111.7 (2023), pp. 1118–1135. DOI: 10.1016/j.neuron.2022.12.034.

[81] Paul Szyszka, Christiane Demmler, Mariann Oemisch, Ludwig Sommer, Stephanie Biergans, Benjamin Birnbach, Ana F Silbering, and C Giovanni Galizia. "Mind the gap: olfactory trace conditioning in honeybees". In: *Journal of Neuroscience* 31.20 (2011), pp. 7229–7239. DOI: 10.1523/JNEUROSCI.6668-10.2011.

[82] Dana Shani Galili, Alja Lüdke, C Giovanni Galizia, Paul Szyszka, and Hiromu Tanimoto. "Olfactory trace conditioning in Drosophila". In: *Journal of Neuroscience* 31.20 (2011), pp. 7240–7248. DOI: 10.1523/JNEUROSCI.6667-10.2011.

[83] Hiromu Tanimoto, Martin Heisenberg, and Bertram Gerber. "Event timing turns punishment to reward". In: *Nature* 430.7003 (2004), pp. 983–983. DOI: 10.1038/430983a.

[84] Dushyant Mishra, Matthieu Louis, and Bertram Gerber. "Adaptive adjustment of the generalization-discrimination balance in larval Drosophila". In: *Journal of neurogenetics* 24.3 (2010), pp. 168–175. DOI: /10.3109/01677063.2010.498066.

[85] Jan Wessnitzer, Joanna M Young, J Douglas Armstrong, and Barbara Webb. "A model of non-elemental olfactory learning in Drosophila". In: *Journal of computational neuroscience* 32 (2012), pp. 197–212. DOI: 10.1007/s10827-011-0348-6.

[86] Fei Peng and Lars Chittka. "A simple computational model of the bee mushroom body can explain seemingly complex forms of olfactory learning and memory". In: *Current Biology* 27.2 (2017), pp. 224–230. DOI: 10.1016/j.cub.2016.10.054.

[87] Evripidis Gkanias, Li Yan McCurdy, Michael N Nitabach, and Barbara Webb. "An incentive circuit for memory dynamics in the mushroom body of Drosophila melanogaster". In: *Elife* 11 (2022), e75611. DOI: 10.7554/eLife.75611.

[88] Paolo Arena, Luca Patané, Vincenzo Stornanti, Pietro Savio Termini, Bianca Zäpf, and Roland Strauss. "Modeling the insect mushroom bodies: Application to a delayed match-to-sample task". In: *Neural Networks* 41 (2013), pp. 202–211. DOI: 10.1016/j.neunet.2012.11.013.

[89] Faramarz Faghihi, Ahmed A Moustafa, Ralf Heinrich, and Florentin Wörgötter. "A computational model of conditioning inspired by Drosophila olfactory system". In: *Neural Networks* 87 (2017), pp. 96–108. DOI: 10.1016/j.neunet.2016.11.002.

[90] Ramón Huerta and Thomas Nowotny. "Fast and robust learning by reinforcement signals: Explorations in the insect brain". In: *Neural computation* 21.8 (2009), pp. 2123–2151. ISSN: 0899-7667. DOI: 10.1162/neco.2009.03-08-733.

[91] Andrew B Barron and Sarah A Corbet. "Pre-exposure affects the olfactory response of Drosophila melanogaster to menthol". In: *Entomologia experimentalis et applicata* 90.2 (1999), pp. 175–181. DOI: 10.1046/j.1570-7458.1999.00436.x.

[92] Vanesa M Fernández, Martin Giurfa, Jean-Marc Devaud, and Walter M Farina. "Latent inhibition in an insect: the role of aminergic signaling". In: *Learning & Memory* 19.12 (2012), pp. 593–597. DOI: /10.1101/lm.028167.112.

[93] Pedro F Jacob, Paola Vargas-Gutierrez, Zeynep Okray, Stefania Vietti-Michelina, Johannes Felsenberg, and Scott Waddell. "An opposing self-reinforced odor pre-exposure memory produces latent inhibition in Drosophila". In: *bioRxiv* (2021). DOI: 10.1101/2021.02.10.430636.

[94] Sathees BC Chandra, Jay S Hosler, and Brian H Smith. "Heritable variation for latent inhibition and its correlation with reversal learning in honeybees (Apis mellifera)." In: *Journal of Comparative Psychology* 114.1 (2000), p. 86. DOI: 10.1037/0735-7036.114.1.86.

[95] RE Lubow, I Weiner, and Paul Schnur. "Conditioned attention theory". In: *Psychology of learning and motivation*. Vol. 15. Elsevier, 1981, pp. 1–49. DOI: 10.1146/annurev.neuro.30.051606.094306.

[96] Christopher J Tabone and J Steven de Belle. "Second-order conditioning in Drosophila". In: *Learning & Memory* 18.4 (2011), pp. 250–253. DOI: /10.1101/lm.2035411.

[97] Christian König, Afshin Khalili, Thomas Niewalda, Shiqiang Gao, and Bertram Gerber. "An optogenetic analogue of second-order reinforcement in Drosophila". In: *Biology Letters* 15.7 (2019), p. 20190084. DOI: 10.1101/lm.047308.118.

[98] Daichi Yamada, Daniel Bushey, Feng Li, Karen L Hibbard, Megan Sammons, Jan Funke, Ashok Litwin-Kumar, Toshihide Hige, and Yoshinori Aso. "Hierarchical architecture of dopaminergic circuits enables second-order conditioning in Drosophila". In: *Elife* 12 (2023), e79042. DOI: 10.7554/eLife.79042.

[99] Syed Abid Hussaini, Bernhard Komischke, Randolf Menzel, and Harald Lachnit. "Forward and backward second-order Pavlovian conditioning in honeybees". In: *Learning & Memory* 14.10 (2007), pp. 678–683. DOI: 10.1101/lm.471307.

[100] Isaac Cervantes-Sandoval, Anna Phan, Molee Chakraborty, and Ronald L Davis. "Reciprocal synapses between mushroom body and dopamine neurons form a positive feedback loop required for learning". In: *Elife* 6 (2017), e23789. DOI: /10.7554/eLife.23789.

[101] Kanta Terao and Makoto Mizunami. "Roles of dopamine neurons in mediating the prediction error in aversive learning in insects". In: *Scientific reports* 7.1 (2017), pp. 1–9. DOI: /10.1038/s41598-017-14473-y.

[102] Brian H Smith and Susan Cobey. "The olfactory memory of the honeybee Apis mellifera. II. Blocking between odorants in binary mixtures." In: *The Journal of experimental biology* 195.1 (1994), pp. 91–108. DOI: 10.1242/jeb.195.1.91.

[103] Robert S Thorn and Brian H Smith. "The olfactory memory of the honeybee Apis mellifera. III. Bilateral sensory input is necessary for induction and expression of olfactory blocking." In: *The Journal of experimental biology* 200.14 (1997), pp. 2045–2055. DOI: 10.1242/jeb.200.14.2045.

[104] JS Hosler and Brian H Smith. "Blocking and the detection of odor components in blends". In: *Journal of Experimental Biology* 203.18 (2000), pp. 2797–2806. DOI: 10.1242/jeb.203.18.2797.

[105] K Takeda. "Classical conditioned response in the honey bee". In: *Journal of Insect Physiology* 6.3 (1961), pp. 168–179.

[106] Kevin C Daly and Brian H Smith. "Associative olfactory learning in the moth Manduca sexta". In: *Journal of Experimental Biology* 203.13 (2000), pp. 2025–2038. DOI: 10.1242/jeb.203.13.2025.

[107] Michael Schleyer, Markus Fendt, Sarah Schuller, and Bertram Gerber. "Associative learning of stimuli paired and unpaired with reinforcement: evaluating evidence from maggots, flies, bees, and rats". In: *Frontiers in psychology* 9 (2018), p. 1494. DOI: 10.3389/fpsyg.2018.01494.

[108] David Tadres and Matthieu Louis. "PiVR: An affordable and versatile closed-loop platform to study unrestrained sensorimotor behavior". In: *PLoS biology* 18.7 (2020), e3000712. DOI: /10.1371/journal.pbio.3000712.

[109] Benjamin Risse, Dimitri Berh, Nils Otto, Christian Klämbt, and Xiaoyi Jiang. "FIMTrack: An open source tracking and locomotion analysis software for small animals". In: *PLoS computational biology* 13.5 (2017), e1005530. DOI: /10.1371/journal.pcbi.1005530.

[110] Isabell Schumann and Tilman Triphan. "The pedtracker: An automatic staging approach for drosophila melanogaster larvae". In: *Frontiers in Behavioral Neuroscience* (2020), p. 241. DOI: /10.3389/fnbeh.2020.612313.

[111] Tomoko Ohyama, Tihana Jovanic, Gennady Denisov, Tam C Dang, Dominik Hoffmann, Rex A Kerr, and Marta Zlatic. "High-throughput analysis of stimulus-evoked behaviors in Drosophila larva reveals multiple modality-specific escape strategies". In: *PloS one* 8.8 (2013), e71706. DOI: /10.1021/acscatal.9b04293.

[112] Guangwei Si, Jessleen K. Kanwal, Yu Hu, Christopher J. Tabone, Jacob Baron, Matthew Berck, Gaetan Vignoud, and Aravinthan D.T. Samuel. "Structured Odorant Response Patterns across a Complete Olfactory Receptor Neuron Population". In: *Neuron* 101.5 (2019), 950–962.e7. ISSN: 0896-6273. DOI: 10.1016/J.NEURON.2018.12.030.

[113] Katherine I Nagel and Rachel I Wilson. "Biophysical mechanisms underlying olfactory receptor neuron dynamics". In: *Nature neuroscience* 14.2 (2011), pp. 208–216. ISSN: 1546-1726. DOI: 10.1038/nn.2725.

[114] Srinivas Gorur-Shandilya, Mahmut Demir, Junjiajia Long, Damon A Clark, and Thierry Emonet. "Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli". In: *Elife* 6 (2017), e27670. DOI: 10.7554/eLife.27670.

[115] Heike Demmer and Peter Kloppenburg. "Intrinsic Membrane Properties and Inhibitory Synaptic Input of Kenyon Cells as Mechanisms for Sparse Coding?" In: *Journal of Neurophysiology* 102.3 (2009), pp. 1538–1550. ISSN: 0022-3077. DOI: 10.1152/jn.00183.2009.

[116] Jan Kropf and Wolfgang Rössler. "In-situ recording of ionic currents in projection neurons and Kenyon cells in the olfactory pathway of the honeybee". In: *PloS one* 13.1 (2018), e0191425. DOI: 10.1371/journal.pone.0191425.

[117] Bruno A. Olshausen and David J. Field. "Sparse coding of sensory inputs". In: *Current Opinion in Neurobiology* 14.4 (2004), pp. 481–487. ISSN: 09594388. DOI: 10.1016/j.conb.2004.07.007.

[118] Horace B Barlow. "Sensory mechanisms, the reduction of redundancy, and intelligence". In: *Mechanisation of thought processes*. London: Her Majesty's Stationery Office, 1959, pp. 535–539.

[119] Iori Ito, Rose Chik-Ying Ong, Baranidharan Raman, and Mark Stopfer. "Sparse odor representation and olfactory learning". In: *Nature neuroscience* 11.10 (2008), pp. 1177–1184. ISSN: 1546-1726. DOI: 10.1038/nn.2192.

[120] Roger Herikstad, Jonathan Baker, Jean Philippe Lachaux, Charles M. Gray, and Shih Cheng Yen. "Natural movies evoke spike trains with low spike time variability in cat primary visual cortex". In: *Journal of Neuroscience* 31.44 (2011), pp. 15844–15860. ISSN: 02706474. DOI: 10.1523/JNEUROSCI.5153-10.2011.

[121] Bilal Haider, Matthew R. Krause, Alvaro Duque, Yuguo Yu, Jonathan Touryan, James A. Mazer, and David A. McCormick. "Synaptic and Network Mechanisms of Sparse and Reliable Visual Cortical Activity during Nonclassical Receptive Field Stimulation". In: *Neuron* 65.1 (2010), pp. 107–121. ISSN: 08966273. DOI: 10.1016/j.neuron.2009.12.005.

[122] Chris Häusler, Alex Susemihl, and Martin Paul Nawrot. "Natural image sequences constrain dynamic receptive fields and imply a sparse code". In: *Brain research* 1536 (2013), pp. 53–67. DOI: /10.1016/j.brainres.2013.07.056.

[123] Scott A. Kreher, Jae Young Kwon, and John R. Carlson. "The molecular basis of odor coding in the Drosophila larva". In: *Neuron* 46.3 (2005), pp. 445–456. ISSN: 08966273. DOI: 10.1016/j.neuron.2005.04.007.

[124] Derek J Hoare, James Humble, Ding Jin, Niall Gilding, Rasmus Petersen, Matthew Cobb, and Catherine McCrohan. "Modeling peripheral olfactory coding in Drosophila larvae". In: *PLoS One* 6.8 (2011), e22996. DOI: /10.1371/journal.pone.0022996.

[125] Birgit Michels, Timo Saumweber, Roland Biernacki, Jeanette Thum, Rupert D.V. Glasgow, Michael Schleyer, Yi Chun Chen, Claire Eschbach, Reinhard F. Stocker, Naoko Toshima, Teiichi Tanimura, Matthieu Louis, Gonzalo Arias-Gil, Manuela Marescotti, Fabio Benfenati, and Bertram Gerber. "Pavlovian conditioning of larval Drosophila: An illustrated, multilingual, hands-on manual for odor-taste associative learning in maggots". In: *Front. Behav. Neurosci.* 11.April (2017), pp. 1–6. ISSN: 16625153. DOI: 10.3389/fnbeh.2017.00045.

[126] Marcel Stimberg, Romain Brette, and Dan FM Goodman. "Brian 2, an intuitive and efficient neural simulator". In: *Elife* 8 (2019), e47314. DOI: /10.7554/eLife.47314.

[127] Antoine Wystrach, Konstantinos Lagogiannis, and Barbara Webb. "Continuous lateral oscillations as a core mechanism for taxis in Drosophila larvae". In: *Elife* 5 (2016). DOI: 10.7554/elife.15504.
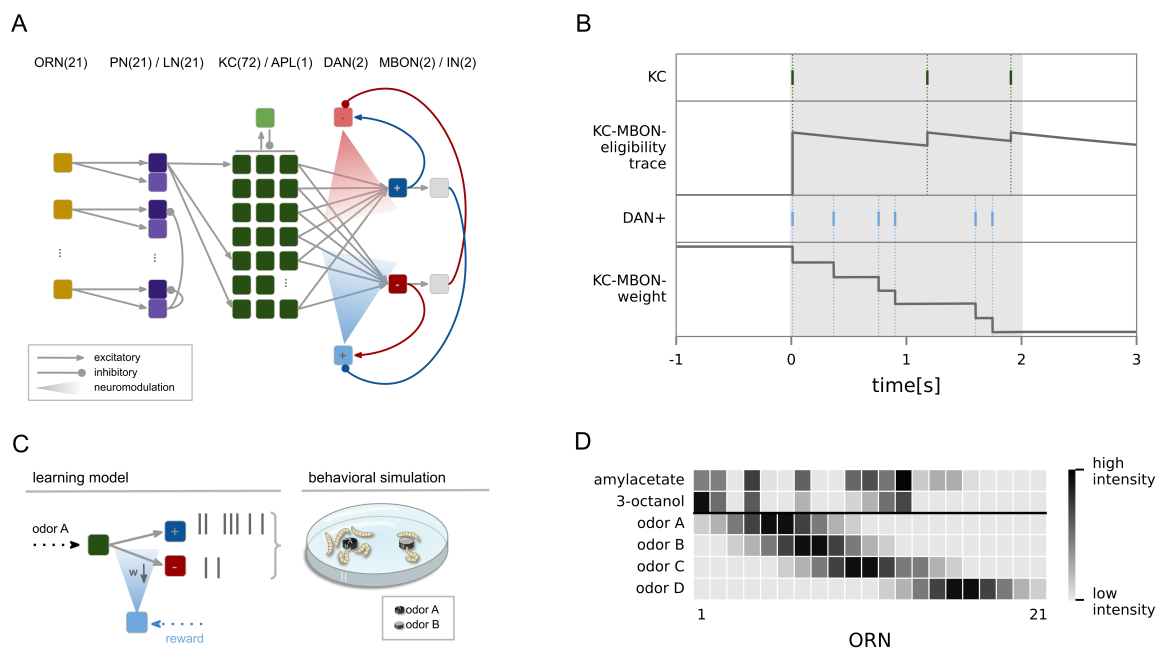
29

Figure 1: **Network mechanisms**. (A) Network model of the *Drosophila* larva olfactory pathway including all neurons and connections implemented. One-to-one feed-forward connections between 21 olfactory receptor neurons (ORN) and 21 projection neurons (PN)/local interneurons (LN) and from 2-6 PN to each of the 72 Kenyon cells (KC). Lateral inhibition from each LN innervates all PNs and recurrent feedback inhibition from the anterior paired lateral (APL) neuron is provided onto all KCs. The MB output region is organized in two distinct compartments. The upper compartment holds the approach encoding MBON$_+$ and is innervated by the punishment mediating DAN$_-$, the lower compartment holds the avoidance mediating MBON$_-$ and is innervated by the reward mediating DAN$_+$. Each DAN can exert a neuromodulatory effect on the plastic KC>MBON synapses within its compartment. MBONs provide excitatory and inhibitory (via gray interneurons) feedback to the DANs. (B) Sketch of synaptic weight change at a single KC>MBON synapse with respect to the synaptic eligibility trace elicited by KC spikes and the occurrence of reward-triggered spikes in DAN$_+$. Amylacetate is paired with a reward for 2s (gray shaded area). (C) To generate simulated larval behavior in the petri dish during the test phase of the learning experiments, we utilized our locomotory model [52], based on the behavioral bias (eqn. 4) acquired by the MB model during the training phase. The behavioral bias is used directly as input to the locomotory model. (D) All odors (see Methods, section: Sensory input) were used in the experiments. Naturalistic odor patterns for amylacetate and 3-octanol as well as four artificial patterns (odorA,odorB,odorC,odorD) with varying distances (see Methods, section: Sensory input) from odorA. Each odor activates a different set of input neurons with a different spike rate, as indicated by the color bar.
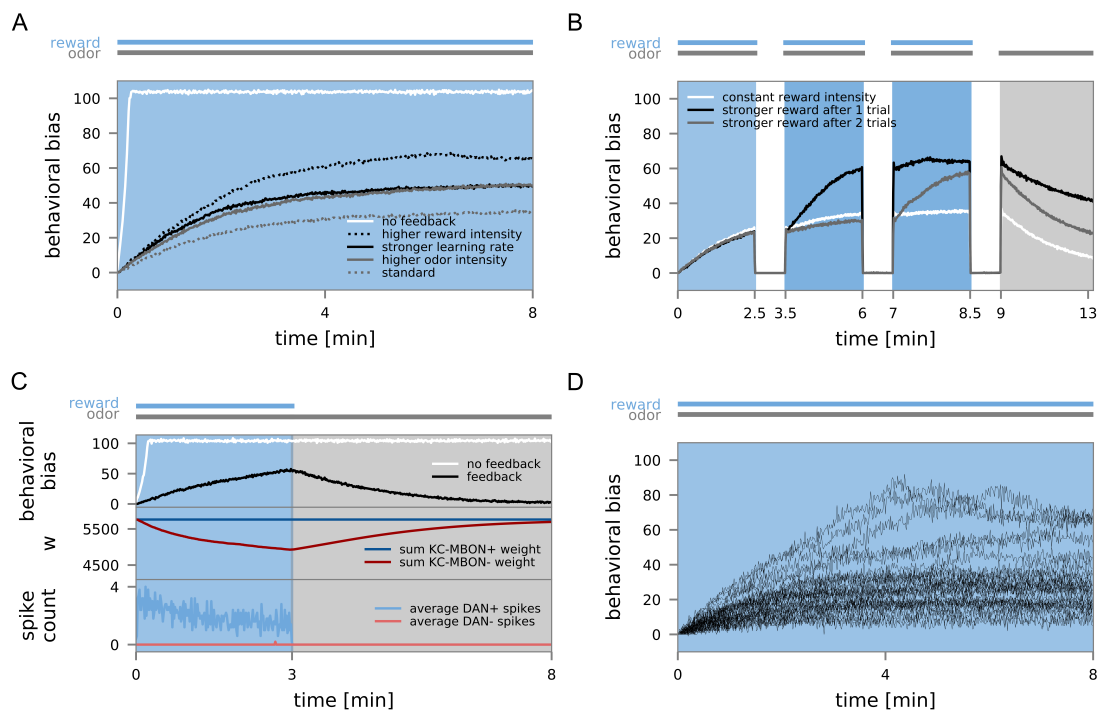
Figure 2: **Learning with prediction errors**. (A) $N = 30$ model instances were trained with the odor amylacetate (CS) and reward (US, blue background). MBON>DAN feedback, the reward/odor intensity, and the learning rate were manipulated in separate experiments. The odor preference (behavioral bias, eqn. 4) was measured continuously in windows of 1 sec and averaged over all model instances. (B) $N = 30$ model instances were trained during three trials with amylacetate and reward (blue background). Reward intensity was either constant across the three training trials (white curve), or enhanced during the third (gray) or the second and third trials (black). The training was followed by a 3 min test phase with odor only (gray background). (C) $N = 30$ model instances were trained with amylacetate and reward (blue background) and then underwent an extended test phase (gray background). (D) Individual acquisition curves for $N = 30$ model instances (standard experiment fig. 2 A).
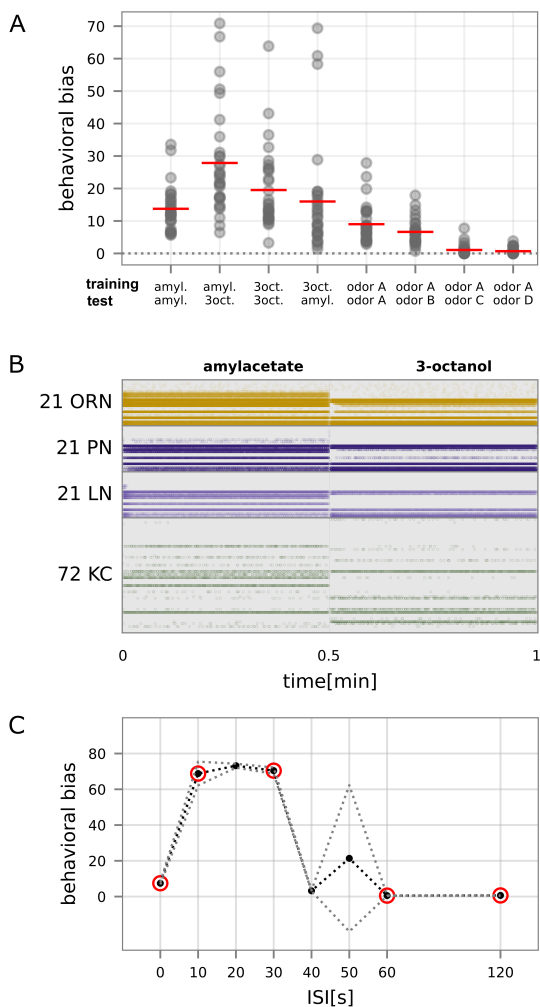
31

Figure 3: **Reward generalization and trace conditioning.** (A) The behavioral bias generalizes to odors that differ from the training odor after a 4min training (3min test phase). We conducted simulation experiments with different combinations of training and testing odor, each for 10 groups (gray circles represent the mean of a single group) of $N = 30$ larvae, and red lines indicate the mean between groups. The behavioral bias is highest when the training and the testing odor are the same. (B) Spiking activity in the network during the presentation of amylacetate (left) and 3-octanol (right) in a single naive model instance. (C) Simulated trace conditioning experiments with odor (amylacetate) and reward. Inter-stimulus interval (ISI) indicates the time between odor and reward onset. The black line displays the mean, gray lines the std over $N = 10$ groups of 30 model instances each. Conditions circled in red correspond to the conditions also used in animal experiments [29, 66]
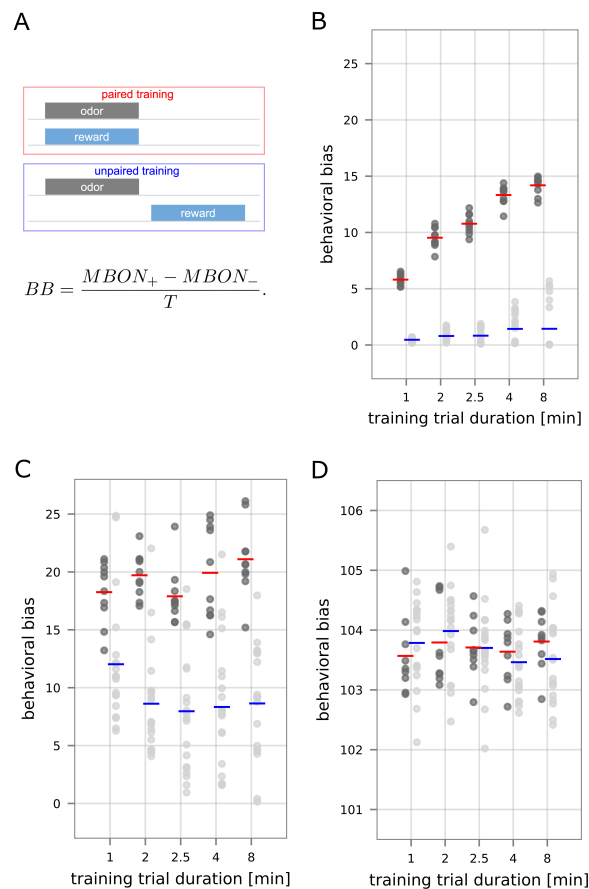
.

32

Figure 4: **Paired and unpaired learning in the MB model.** (A) Schematic overview of the paired vs. unpaired training protocol. (B) The model's behavioral bias for training with amylacetate and reward for $N = 10$ paired (dark gray, mean in red) and $N = 10$ unpaired (light gray, mean in blue) experiments with groups of 30 modeled larvae each. In the unpaired condition, half of the groups were trained with the odor preceding the reward, for the other half, the reward preceded the odor. (C) Model behavioral bias for amylacetate for $N = 30$ paired and $N = 30$ unpaired experiments with randomized order of odor and reward. Prior to the conditioning experiment the model instances underwent a 10min pre-training period, during which odor and reward were paired. (D) Model behavioral bias for amylacetate for $N = 30$ paired and $N = 30$ unpaired experiments with randomized order of odor and reward. The MBON>DAN feedback was disabled. Prior to the conditioning experiment the model instances underwent a 10min pre-training.
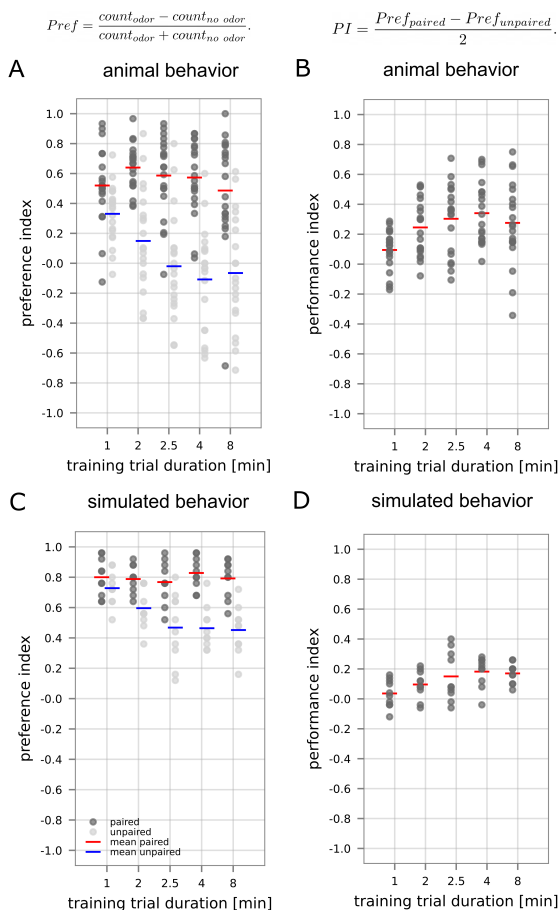
33

$$Pref = \frac{count_{odor} - count_{no\ odor}}{count_{odor} + count_{no\ odor}}.$$

$$PI = \frac{Pref_{paired} - Pref_{unpaired}}{2}.$$

Figure 5: **Replicating behavioral experiments with paired and unpaired training**. (A) Experimental preference indices for amylacetate for 20 groups of 30 real animals each for paired and unpaired experiments with randomized order of odor and reward [1]. (B) Experimental performance indices for amylacetate computed between preference in paired and unpaired real animal experiments [1]. (C) The simulated behavior is based on the protocol in A. Simulated preference indices for amylacetate for $N = 10$ paired and $N = 10$ unpaired experiments with varied order of odor and reward. (D) Simulated performance indices for amylacetate computed between preference in paired and unpaired simulation experiments.
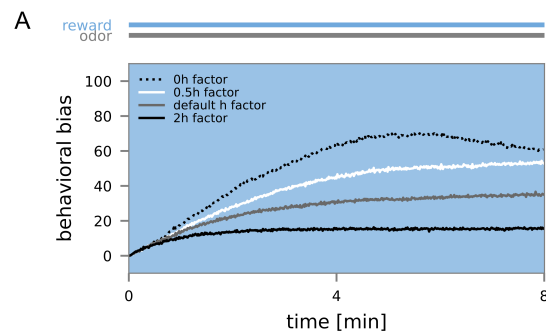
Figure S1: **The effect of the homeostatic mechanism on the learning curve**. (A) $N = 30$ model instances were trained with the odor amylacetate (CS) and reward (US, blue background). The odor preference (behavioral bias) was measured continuously in windows of 1 sec and averaged over all model instances. The learning rate was the same in all three experiments, while the magnitude of the homeostatic regulation h (eqn. 3,table S1) was either at its default value, at 0, or at half or twice the magnitude of the default value.

Neuron Parameters

Capacitance ORN	C_m^O	100pF
Capacitance PN	C_m^P	30pF
Capacitance LN	C_m^L	50pF
Capacitance KC	C_m^K	30pF
Capacitance APL	C_m^A	200pF
Capacitance MBON	C_m^M	100pF
Capacitance DAN	C_m^D	100pF
Leak Conductance ORN	g_L^O	5nS
Leak Conductance PN	g_L^P	2.5nS
Leak Conductance LN	g_L^L	2.5nS
Leak Conductance KC	g_L^K	5nS
Leak Conductance APL	g_L^A	5nS
Leak Conductance MBON	g_L^M	5nS
Leak Conductance DAN	g_L^D	5nS
Leak Potential ORN	E_L^O	-60mV
Leak Potential PN	E_L^P	-59mV
Leak Potential LN	E_L^L	-59mV
Leak Potential KC	E_L^K	-55mV
Leak Potential APL	E_L^A	-60mV
Leak Potential MBON	E_L^M	-60mV
Leak Potential DAN	E_L^D	-60mV
Threshold Potential ORN	V_T^O	-35mV
Threshold Potential PN	V_T^P	-30mV
Threshold Potential LN	V_T^L	-30mV
Threshold Potential KC	V_T^K	-35mV
Threshold Potential APL	V_T^A	-30mV
Threshold Potential MBON	V_T^M	-30mV
Threshold Potential DAN	V_T^D	-30mV
Resting Potential ORN	V_r^O	-60mV
Resting Potential PN	V_r^P	-59mV
Resting Potential LN	V_r^L	-59mV
Resting Potential KC	V_r^K	-55mV
Resting Potential APL	V_r^A	-60mV
Resting Potential MBON	V_r^M	-60mV
Resting Potential DAN	V_r^D	-60mV
Refractory Time	τ_{ref}	2ms

Synaptic Parameters

Excitatory Potential	E_E	0mV
Inhibitory Potential	E_I	-75mV
Excitatory Time Constant	τ_e	5ms
Inhibitory Time Constant	τ_i	10ms

Plasticity Parameters

Eligibility Trace Time Constant	$\tau_{eligibility}$	5s
Learning Rate	α	0.3nS

Synaptic Weights

Weight Input-ORN	wORNinputORN	3nS
Weight ORN-PN	wORNPN	10nS
Weight ORN-LN	wORNLN	4nS
Weight LN-PN	wLNPN	1nS
Weight PN-KC	wPNKC	1nS
Weight KC-APL	wKCAPL	20nS
Weight APL-KC	wAPLKC	50nS
Weight KC-MBON	wKCMBON	80nS
Weight Input-DAN	wDANinputDAN	2.5nS
Excitatory Weight MBON-DAN	wMBONDANex	4nS
Excitatory Weight MBON-local interneuron	$wMBONMBON_LN$	35nS
Inhibitory Weight local interneuron-DAN	wMBONDANin	70nS
Normalization Factor KC-MBON	$normalization_{factor}$	0.0001

Adaptation Parameters

Adaptation Time Constant	τ_{Ia}	1000ms
Adaptation Reversal Potential	E_{Ia}	-90mV
Increase of Spike Frequency Adaptation Conductance ORN	ORN_{SFA}	0.1nS
Increase of Spike Frequency Adaptation Conductance KC	KC_{SFA}	0.02nS
Increase of Spike Frequency Adaptation Conductance MBON	$MBON_{SFA}$	0.1nS
Increase of Spike Frequency Adaptation Conductance DAN	DAN_{SFA}	0.1nS

Simulation Parameters

Time Step	dt	0.1ms

Table S1: Network parameters.

3.2   Minimal circuit motifs for second-order conditioning in the insect mushroom body

Minimal circuit motifs for second-order conditioning in the insect mushroom body

Anna-Maria Jürgensen[a], Felix Johannes Schmitt[a], and Martin Paul Nawrot[a]

[a]Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Cologne, Germany

Abstract

In well-established first-order conditioning experiments, the concurrence of a sensory cue with reinforcement forms an association, allowing the cue to predict future reinforcement. Once a sensory cue is established as a predictor, it can also serve as indirect reinforcement, a phenomenon referred to as second-order conditioning. In the insect mushroom body, such associations are encoded in the plasticity of the synapses between the intrinsic and output neurons of the mushroom body, a process mediated by the activity of dopaminergic neurons that encode reinforcement signals. In second-order conditioning, a new sensory cue is paired with an already established one that presumably activates dopaminergic neurons due to its predictive power of the reinforcement. We explore minimal circuit motifs in the mushroom body for their ability to support second-order conditioning. We found that dopaminergic neurons can either be activated directly by the mushroom body's intrinsic neurons or via feedback from the output neurons via several pathways. We demonstrate that the circuit motifs differ in their computational efficiency and robustness and suggest a particular motif that relies on feedforward input of the mushroom body intrinsic neurons to dopaminergic neurons as a promising additional candidate for experimental evaluation. It differentiates well between trained and novel stimuli, demonstrating robust performance across a range of model parameters.

Introduction

By forming associations between sensory cues and reinforcement during classical conditioning (first-order conditioning, FOC), animals can learn to predict the emergence of environmental factors relevant to their survival. Once a sensory cue has been established as a predictor of such reinforcement, it can act as reinforcement in second-order conditioning (SOC). SOC has been observed across species with experiments conducted in *Drosophila* [1–4] and other invertebrate [5–8] as well as vertebrate [9–11] species. SOC experiments involve two initially neutral stimuli (stimulus 1 and stimulus 2). Stimulus 1 is first paired directly with reinforcement during FOC, whereby it acquires a valence as a cue for reinforcement. Afterward, stimulus 2 is paired with stimulus 1 (SOC), causing an expansion of the acquired valence of stimulus 1 onto stimulus 2, without stimulus 2 itself being paired with the reinforcer. Afterward, both stimuli initiate a behavioral response based on their acquired valence.

1

In *Drosophila* [12–15] and other insects [15–18], the mushroom body (MB) is a crucial brain structure for learning and encoding relationships between sensory cues and reinforcement. The Kenyon cells (KC) are the intrinsic neurons of the MB, encoding the identity of sensory inputs in *Drosophila* [19–21] as well in other insects [22–25]. In both *Drosophila* [26, 27] and other insects [24, 25, 28, 29], they relay their output onto a much smaller number of MB output neurons (MBONs). Plasticity at the KC>MBON synapses allows MBONs to encode the valence of a sensory cue, according to extensive experimental evidence from *Drosophila* [30–35]. Neuromodulators, such as dopamine mediate this plasticity [36–42]. In *Drosophila*, it has been shown that either an inherently punishing or rewarding stimulus (electric shock, sugar) [1, 3, 4] or direct optogenetic activation of dopaminergic neurons (DANs) [3] can be utilized to deliver a reinforcement signal during FOC phase of such experiments to establish second-order memory later. Experiments in *Drosophila* have suggested that stimulus 1 itself causes activation of DANs or enhances it after being paired with reinforcement [3, 4, 43, 44]. This aligns with the analogous finding that DANs in the ventral tegmental of vertebrates respond to learned cues that predict upcoming reward [45]. The mechanism inducing synaptic plasticity during both FOC and SOC likely relies on DAN activation. During FOC, DANs are activated directly by the reinforcer (Figure 1 A). Following this conditioning procedure, stimulus 1 seems to indirectly activate the DANs during SOC (Figure 1 C, [3, 4, 44]). The strength of the behaviorally expressed stimulus 1 and stimulus 2 valence after SOC can be similarly strong [1, 3] or weaker [3, 4] for stimulus 2 .

Two different circuit mechanisms lend themselves to achieving such post-FOC activation of DANs by stimulus 1: Firstly, a stimulus 1 representation among the KCs could serve as direct stimulus-specific input to the DANs [46–51]. Alternatively, the input could be supplied via MBON feedback [46, 47, 51, 52]. Their response to stimulus 1, altered by learning, could serve as a manifestation of stimulus-specific valence.

Here, we tested possible circuit motifs that could underly SOC in the insect MB using abstract and simplified network models inspired by the *Drosophila* olfactory pathway and the MB. Starting from a basic model of the MB, we explored different circuit configurations and their capacity to produce SOC in an olfactory learning protocol to identify promising candidates for experimental testing. To define our solution space, we assumed that learning in the MB depends on KC>MBON plasticity, mediated by a dopamine signal during FOC and SOC. Model circuits should be able to produce both FOC and SOC without generalizing associations with reinforcement unspecifically to novel stimuli. We tested all models in classical conditioning experiments and demonstrated their ability to support FOC and SOC. Additionally, we evaluated differences in their biological plausibility by quantifying robustness and discussing functional and anatomical evidence for the respective circuits. We found that a particular circuit that achieves DAN activation through excitatory KC input during SOC outperforms the other candidates and appears compatible with the MB anatomy. We suggest this circuit motif that differs from previously reported mechanisms [2–4] as an additional candidate for experimental tests.

Results

First-order conditioning in a basic mushroom body circuit

All models are based on rate units, each representing the activation of a single neuron. The basic network consists of 2000 KCs, each innervating two MBONs (Figure 1 A). It has been shown that MBONs receive inputs from many of the KCs in adult *Drosophila* [26, 51]. For simplicity, we

started by modeling a complete KC>MBON connectivity. Some MBONs can be categorized as approach or avoidance signaling [30, 33, 34, 53]. In the model, this corresponds to MBON$^+$ (Eqn. 2) and MBON$^-$ (Eqn. 3), respectively. Other types of MBONs were disregarded here. Initially, all synaptic KC>MBON weights were set to the same value (table 2). A single DAN is included in the model, which can be activated by the external input, representing a reinforcer in the environment (Figure 1 A). KC>MBON$^-$ synapses undergo plasticity whenever trial-based KC activation, driven by odor input, and DAN activation coincide. We employ a two-factor learning rule (Eqn. 4) at the KC>MBON$^-$ synapses, leading to a decrease in the synaptic weights with the limitation that they can not take on a negative value. DAN activation (Eqn. 1) is the sum of the model-specific external input rate Ip_{ext} (table 3) that represents reinforcement and the network internal input Ip_{int}, provided via the different circuit mechanisms. In all equations, $\overrightarrow{x}$ always denotes a vector. R represents the activity of a neuron, which can be interpreted as a spike rate of a neuron or a vector of neurons in the case of KCs, and w denotes a synaptic weight or a vector of weights. LR refers to learning rates.

$$R_{DAN}(t) = Ip_{ext}(t) + Ip_{int}(t) \tag{1}$$

$$R_{MBON^+}(t) = \overrightarrow{R}_{KC}(t) \cdot \overrightarrow{w}_{KC>MBON^+}(t) \tag{2}$$

$$R_{MBON^-}(t) = \overrightarrow{R}_{KC}(t) \cdot \overrightarrow{w}_{KC>MBON^-}(t) \tag{3}$$

$$w^i_{KC>MBON^-}(t+1) = \begin{cases} w^i_{KC>MBON^-}(t) - LR \cdot R_{DAN}(t), & \text{if } R^i_{KC}(t) > 0 \text{ and} \\ & w^i_{KC>MBON^-}(t) > (LR \cdot R_{DAN}(t)) \\ w^i_{KC>MBON^-}(t), & \text{otherwise} \end{cases} \tag{4}$$

Candidate circuits for second-order conditioning

Using the basic circuit model (Figure 1 A) as a starting point, we implemented five different extended versions of it (Figure 2). These models either rely on some form of KC>DAN input (model 1, model 2) or MBON>DAN feedback (model 3, model 4, model 5) as a means to expand the learned association of odor 1 with reinforcement to odor 2 during SOC. Unless specified otherwise, the DAN is not spontaneously active. The equations for all models can be found in table 1.

To compare the different circuit motifs in an unbiased manner, their parameters were optimized using grid search (see Methods), except the fixed parameters shared between all models (Table 2), which were kept constant to allow better comparison between the candidate mechanisms for SOC implemented in the different circuits. The goal for parameter optimization was to identify parameter combinations for each model that yield the best learning performance in an associative learning experiment that consisted of a combination of FOC and SOC learning trials (Figure 1 B). In insect learning experiments, forming a direct or indirect association with reward leads to approach behavior that can manifest in the movement toward the source of an odor or feeding-related behavior [3, 5, 6] or the avoidance of a punishment-associated odor [4]. In our model experiments, the successful acquisition of an association with reward was quantified as the normalized difference between MBON output rates (R_{MBON^+} and R_{MBON^-}), which we will refer to as the approach bias (Eqn. 5) because MBON activity has been shown to initiate approach or avoidance behavior [30, 33, 54].

$$B = \frac{R_{MBON^+} - R_{MBON^-}}{R_{MBON^+} + R_{MBON^-}} \tag{5}$$

Model 1

Model 1 includes plastic KC>KC excitatory feedback with increments of $LR_{KC>KC}$(table 3)), triggered by pre and postsynaptic KC activation (Eqn. 8) and KC>DAN synapses of a fixed strength (Eqn. 7).

Model 2

Model 2 expands the basic circuit with excitatory plastic synapses between KCs and the DAN (Eqn. 10). They are each initialized with $init_{KC>DAN}$ (table 3) and are increased by $LR_{KC>DAN}$ (table 3) when activation of the respective KC coincides with DAN activity, yielding DAN activation (Eqn. 9).

Model 3

In model 3, network input into the DAN is implemented via feedback from both MBONs (Eqn. 11). Inhibitory input with a fixed synaptic strength comes from MBON$^-$, while excitatory input is provided by the MBON$^+$ (table 3).

Model 4

Model 4 uses a spontaneously active DAN ($R_{DANbaseline}$, (table 3)) in combination with inhibitory MBON$^-$>DAN input (Eqn. 12). Both effects regulate the DAN activation in the absence of reward.

Model 5

In model 5, an excitatory plastic MBON$^+$>DAN synapse (Eqn. 14) is added to the basic circuit. When MBON$^+$ and DAN activity coincide, the synaptic strength is increased by $LR_{MBON_+>DAN}$ for each DAN spike (table 3). During FOC, this synapse is strengthened, allowing for activation of the DAN during SOC. This allows KCs to activate the DAN (Eqn. 13).

Identifying optimal parameters for each model

All models were trained in a combined FOC and SOC protocol, where odors were used as stimuli and tested for their approach bias first after completing FOC with odor 1 and reward, then after completing SOC with odors 1 and 2 (Figure 1 B, Eqn. 5). Plasticity was disabled during testing to isolate odor valence acquired during the respective training trials without the influence of the test itself. Additionally, a novel odor 3 was included in both tests to examine any generalization of the reward association that might have occurred during the FOC or SOC training processes. To assess the ability of the parameters of the different circuit motifs to support SOC, we optimized each model for the highest possible SOC performance, which translates to maximizing the approach bias for odor 2 after three SOC trials. Additionally, we introduce several criteria the models must fulfill to ensure that the learning effect for odor 2 is odor-specific and originates from the respective mechanism applied during the SOC trials. These criteria are:

- post FOC odor 1 approach bias = 1

- post FOC odor 2 approach bias = 0

- post FOC odor 3 approach bias = 0

- post SOC odor 1 approach bias = 1

- post SOC odor 3 approach bias = 0

Additionally, DAN and MBON rates should never exceed 20 Hz and 50 Hz, respectively, to stay within the biologically observed range for MBONs [34] and DANs [55]. The reported DAN spike rates are responses to an electric shock that were reported to range up to 10 Hz during a 200ms stimulation. Since the response of the DANs showed a temporal delay, we decided to use 20 Hz instead.

All parameter combinations that fulfilled the above criteria for each model were further evaluated. Among all of these parameter combinations, we selected those for each model that yielded the highest SOC performance, measured by the odor 2 approach bias after SOC. We refer to each of these as an *optimal learner*. Grid search for all models 1-5 yielded several *optimal learners*. We computed the average *optimal learner* by averaging all *optimal learners* within every parameter. We argue that this average *optimal learner* approximates the center of all equally good parameter combinations. Next, Euclidean distances were computed between all z-standardized *optimal learners* and the average *optimal learner*. The parameter combination with the smallest distance to the average parameter combination in an n-dimensional z-standardized space (n=number of optimized parameters) was selected (see Methods). We assume that parameter combinations closer to the average can be considered biologically more plausible because their central location makes them more robust to parameter deviations in all directions (see Discussion).

The basic learning model (Figure 1 A) fulfilled the criteria for FOC learning, but no parameter combination could accommodate SOC, yielding no approach bias for odor 2 after SOC. There was at least one optimal parameter combination that fulfilled the optimization criteria for each extended candidate circuit (Figure 2, Figure 3 A). All models acquired an optimal approach bias of 1 for odor 1 at both test times, indicating that the association of odor 1 and reward is learned during FOC and fully retained throughout the SOC protocol. Tests with odor 3 always yielded an approach bias of 0 for all models, indicating that the approach bias does not generalize inadmissibly to fully disjunct odors. All models, except model 1, achieved equally good SOC performance, as indicated by an approach bias of 0.33 for odor 2 after SOC. Model 1 only acquired a bias of 0.02. For each model, the maximum value of SOC performance is determined by the implementation of the compound presentation of odors 1 and 2 (see Methods). The approach bias of 0.33 for SOC (Eqn. 5) is the highest value any model can achieve in this experiment. In the experimental design, 50% of KCs were activated by odor 2 during SOC. Thus, only 50% of the KC>MBON⁻ synapses could be altered during SOC. The value of 0.33, therefore, represents the best possible SOC performance. An additional experiment was conducted as a control where the KC activation patterns for odors 1 and 2 were added during SOC (yielding 400 active KCs). The maximal performance achieved by models 2-5 was 1 in this case.

In none of the models any approach bias for the disjunct, novel odor 3 was observed. Additionally, we extended the experimental protocol (Figure 1 B) with three trials of presenting odor 3 alone and without any reward after SOC and included another test. Depending on the degree

of specificity with which the different model circuits activate the DAN, unwanted generalization of the valence to odor 3 was observed. Models 1,2,3 outperform models 4 and 5 here (Figure 2).

Learning generalizes with increasing odor overlap

The overlap of odors 1 and 2 during training and training was varied separately ($0\%, 25\%, 50\%, 75\%, 100\%$ overlap), encoded in the percentage of KCs jointly activated between the individually presented odor 1 and odor 2 (see Methods). Across all models, both FOC and SOC approach biases (Eqn. 5) increase with the overlap between odor 1 and odor 2 (Figure 3 A). Between highly overlapping odors, the reward association generalizes. This results in an approach bias for odor 2 after FOC, even though odor 2 was not presented during FOC (Figure 3 A). A joint presentation of odor 1 and odor 2 during SOC then leads to an even higher approach bias for odor 2 (Figure 3 A). All models acquire similar biases (Eqn. 5) for odor 2, depending on the degree of overlap, except model 1, where the SOC learning effect was always lower (Figure 3 A).

Keeping odor 1 and odor 2 fully disjunct, we varied the degree of overlap between odor 3 and either odor 1 (Figure 3 B, left) or odor 2 (Figure 3 B, right) during both tests. For all models, the approach bias for odor 3 after FOC scaled with the overlap and reached the same value as odor 1 if fully overlapping (f Figure 3 B, dashed purple lines). Testing again after SOC yielded the same results, thus not depicted here. When the overlap between odor 3 and odor 2 varied at both test times, no approach bias was observed after FOC since odor 2 is not presented during the FOC trials (results not shown). In a test after SOC, all models perform similarly concerning the upper bound of the approach bias at the magnitude reached by odor 2 Figure 3 B, dashed orange lines).

Robustness of second-order conditioning varies across the different model circuits

In the conditioning experiments reported thus far, all model circuits, except model 1, perform equally well (Figure 2, Figure 3). To further differentiate between them, we next examined the robustness of the model's performance to variations of their parameters using three different methods.

We first quantified the percentage of optimal parameter combinations within the searched parameter grid for each model. A real brain would likely not require a single, extremely precise combination of physiological parameters to perform any computational task, such as SOC. Since the parameters of our computational models are ultimately representations of neuronal or synaptic characteristics, the stability of SOC performance across parameter combinations could hint at the degree to which a circuit motif is biologically plausible. For each respective model, four parameters were optimized, yielding a grid with 100^4 parameter combinations. In the case of model 1, $5.9^{-5}\%$ of parameter combinations were equally optimal. Model 2 yielded 6.81%, model 3 only 0.37%, model 4 2.11% and model 5 4.29% *optimal learner* parameter combinations. From this perspective, models 2, 4, and 5 thus seem more robust compared to models 1 and 3.

As an additional measure to assessing the optimal portion of the entire solution space of possible parameter combinations, we used a method for individually sampling the four-dimensional Euclidean parameter space for each model. The four-dimensional space for each model was standardized using the range of the grid (maximum-minimum parameter value, table 3). A four-dimensional hypersphere was positioned as the point representing the average *optimal learner*, with radius=0. We then incrementally increased the radius of the hypersphere from 0 to 1 in the standardized space with 100 linearly spaced steps and, for each increase, sampled 700 data points

from its surface (see Methods, Hypersphere sampling). These sampled data points do not necessarily correspond to data points from the set of equally optimal parameter combinations found in the grid search for the respective model due to the step size used in the parameter optimization. The sampled points were transformed back into their original space and then used to simulate the respective model to examine if this parameter combination would yield FOC and SOC performance that fit the criteria for the *optimal learner*. For each radius increment, we calculated the percentage of sampled points from the hypersphere surface that exhibit the same performance as the parameter combination used as the central point. Since the parameter space was standardized for each model and the radii used were the same, the results can easily be compared between the five models. We find that the models differ in their robustness to deviations of the parameters from their optimal values. Model 2 strongly outperforms the other models. While models 4 and 5 are more robust than model 3, model 1 demonstrates no robustness. Our grid search for model 1 yielded no variability in three of the four optimized parameters ($init_{KC>KC}$, $init_{KC>DAN}$, $LR_{KC>MBON_-}$, $LR_{KC>KC}$, Table 3).

A third approach to comparing the robustness between the different model circuits is to quantify how well they retain their FOC and SOC performance when variability is introduced into the connectivity matrix or the strength of the KC>MBON synapses. We varied either the number of input KCs into each MBON (Figure 5 A,B) or the strength of the synaptic connections while retaining full connectivity (Figure 5 C,D).

We varied the number of KCs providing input to each MBON between 25% and 100% for each model instance ($N = 100$ model instances) while maintaining the magnitude of the individual weights (Table 2, $init_{KC>MBON_+}/init_{KC>MBON_-}$). For each model instance, a random number of connections was drawn from a uniform distribution and applied to each MBON (MBON$^+$, MBON$^-$) to select the same amount of random connections to be active. While FOC performance remained very robust across all models 1-5 (Figure 5 A), SOC performance was significantly impaired in model 1 and model 3 (Figure 5 B), compared to SOC performance with full connectivity (Figure 3 A).

Additionally, we evaluated the robustness of the learning performance when the strength of the KC>MBON weights was varied in a range of $\pm 5\%$ around their default weight (Table 2). Again, no significant differences were observed for FOC performance (Figure 5 C). SOC performance was retained for models 2, 4, and 5 compared to the standard model with the same strength of all synaptic weights (Figure 3 A).

While all five circuit motifs are capable of FOC (Figure 2), model 1 performs very poorly at SOC compared to the other models (Figure 2) and also shows poor robustness (Figure 4). While model circuits 3-5 fulfill the criteria for SOC, they differ in their robustness to reward generalization and variations of their parameters. Model 2, which relies on feed-forward input of KCs to the DAN, emerges as a promising candidate, in addition to the circuits that are already being explored [3, 4].

Discussion

While SOC as a phenomenon has been a target of insect learning experiments for a long time [1–3, 5, 6], the discovery of the underlying circuit mechanisms is just at its beginning [2–4]. We use mechanistic models of different variations of a basic, abstract MB circuit inspired by *Drosophila* and show that different circuit motifs, based on either KC or MBON-driven DAN activation, can support SOC. In the following, we will discuss our results in light of experimental evidence for

SOC in insects and the extent to which the MB anatomy supports the tested circuit motifs.

Second-order conditioning in insect experiments and models

Second-order conditioning has been demonstrated experimentally in honeybees [5, 6] and fruit flies [1–4, 56]. The honeybee experiments measured proboscis extension as a response to conditioning with odors and sugar. Regardless of whether the number of SOC trials was equal to [6] or 50% of the number of FOC trials [5], the conditioned response acquired during FOC was stronger than that acquired during SOC. In the fly experiments, a combination of odor and electric shock [1], odor and sugar [3], or odor and optogenetic DAN activation [3] were used. The same duration of pairing an electric shock with odor 1 and pairing odor 1 with odor 2 during SOC yielded a stronger learning effect for FOC, compared to the SOC effect [1], as observed in bees [5, 6]. Yamada et al. [3] used a protocol with longer FOC than SOC duration in an appetitive conditioning protocol. This led to similarly strong FOC and SOC effects, given a long enough FOC training duration.

Potential circuit mechanisms behind SOC were investigated in some studies, conducted in *Drosophila* [2–4]. Evidence for a mechanism based on MBON>DAN feedback comes from a study that used optogenetic silencing or activation of MBONs as an indirect punishment or reward signal for conditioning avoidance or approach of an odor, thereby circumpassing pairing of reinforcement and stimulus 1 during FOC [2]. Yamada et al. [3] also suggest an MBON>DAN pathway across two layers of interneurons as a mechanism for SOC. They show that a presentation of an odor with optogenetic DAN activation can induce suppression of the response of a particular MBON ($\alpha 1$). Decreased activity of MBON-$\alpha 1$ could cause disinhibition of multiple pathways via two interneurons, leading to a net activation of DANs that encode reward during SOC. The circuit for the disinhibition of the DAN during SOC is closely related to the motif implemented in our model 4, which performs well at SOC but appears not to be very robust to reward generalization onto novel odors. Likewise, in an aversive conditioning paradigm, it was demonstrated that the output of a particular MBON, innervating the γ-lobe (MBON-$\gamma 1$), is required during the SOC phase of the experiment to induce a learned valence of the second odor [4]. Similarly, a single MBON (MBON-α'2) innervating the α'/β'-lobes seems to play a similar role in these lobes [4]. In summary, all of these works demonstrate the importance of MBON output pathways [2–4] and seem to suggest MBON>DAN feedback as a candidate mechanism for SOC. For the sake of completeness, it has to be stated that each experiment targeted specific pathways and can not rule out the presence of different circuit motifs for SOC in other compartments.

Recent models of the adult [57, 58] or larval [47] *Drosophila* MB can accommodate SOC. They all suggest KC>MBON plasticity to learn stimulus 2 during SOC via MBON activity. Either in the form of net excitatory and inhibitory MBON>DAN feedback [47, 58] or via direct modulation of the KC>MBON synapses by altered MBON activity [57]. None of the models allow KC input to the DAN and thus do not test this alternative pathway.

Anatomical evidence for the tested circuit motifs

To evaluate the biological plausibility of our tested circuit motifs, we next assessed which DAN-activation pathways are supported by anatomical evidence. KC>DAN synapses have been found both in larval [46–48] and adult [49–51] *Drosophila*. Since KCs are known to be cholinergic [32] in the adult, it has been suggested that these connections could be excitatory [46] in the larva. This has also been confirmed in adults to affect a particular MBON ($\alpha 2\alpha'2$) [49].

Direct and indirect connections between MBONs and DANs exist in the larva [46, 47] and the adult [51, 52] within and between compartments [46, 47, 51, 52]. In the larva, excitatory and inhibitory synapses have been observed [46, 47]. In the adult *Drosophila* MB, different MBONs have been found to release excitatory or inhibitory transmitters [30], supporting the assumption that here both de- and hyperpolarizing effects of MBONs on DANs might exist.

Direct KC>KC synapses have been found in large numbers in the larva [46]. According to the transmitter released by KCs in the adult [32], they have been speculated to be excitatory [46]. In the adult, KC>KC synapses have also been demonstrated [51, 59]. KCs have been shown to express both muscarinic [59, 60] and nicotinic [61, 62] receptors, the combination of which likely enables both inhibitory [59, 60] and excitatory [63] synapses between them and rendering different interactions plausible.

Limitations

Motivated by isolating SOC as the phenomenon of interest in our experiments, we decided to reduce our circuit implementations of computational motifs to their minimum and optimize their parameters only for FOC and SOC. This allows us to determine which circuit motifs are the most efficient in computing SOC when optimized solely for this purpose. This approach neglects the surrounding network structures in the real insect MB and thus intentionally disregards other learning phenomena often addressed when studying the MB, such as prediction error, effects of stimulus exposure before learning, forgetting, or extinction. Both forgetting and extinction produce the same observable behavior in experiments, which is a decline in the response to repeated presentation of a sensory cue when reinforcement is omitted after conditioning. As opposed to forgetting, extinction is characterized by the possibility of recovery of the association after its temporary loss [64], has been demonstrated in adult *Drosophila* [65, 66]. To retain the association for recovery, extinction requires the formation of parallel memory traces for the acquisition and the decline of the association [67, 68]. The repeated presentation of stimulus 1 without reinforcement during SOC should lead to the extinction of the learned association between stimulus 1 and reinforcement. Across many trials, SOC and extinction learning should be competing phenomena, allowing SOC to occur only until the extinction process has abolished the odor approach bias. While some studies find a decline in the association between stimulus 1 and the reinforcement during SOC in honeybees [5, 6], experiments in *Drosophila* report no loss of the association between stimulus 1 and reinforcement during SOC [1, 3]. For our *Drosophila*-inspired modeling approach, we thus defined no loss of odor 1 approach bias between the tests following FOC and SOC as a criterion for our parameter optimization. By optimizing the models according to this criterion, we exclude the possibility of studying the effects of loss of an association, such as extinction or forgetting.

Additionally, the omission of mechanisms for long-term network stability in combination with the criterion of perfect odor 1 approach bias in both tests after FOC and SOC guarantee the complete down-regulation of all KC>MBON$^-$ synapses activated by odor 1 after three FOC trials of arbitrary duration would not allow for experiments with more trials. Stabilizing mechanisms for homeostasis could be introduced in a more comprehensive modeling approach, implementing the most promising circuit motif in a more naturalistic MB model.

Aside from the narrow applicability to different learning phenomena, which is the downside of our minimal circuit design, another limitation originates from the need to define the limits and the step size for the model parameter grid search. The success of it depends on selecting these limits and steps appropriately (Table 3). If ill-chosen, they could put individual models at a competitive disadvantage by not including or over-stepping their optimum.

Outlook

We demonstrated that several circuit mechanisms are potential candidates for SOC. While they vary in computational efficiency and robustness, multiple models remain good candidates, compatible with our knowledge of the MB anatomy. A pathway for KC-driven DAN activation emerges as an especially promising candidate among the circuits we tested. A valuable next step would be to integrate them into more comprehensive MB models to test how they interact with other phenomena in learning. This could also be another angle to studying their robustness. Especially interesting in this regard would be extinction, with its inherently interfering mechanism. SOC relies on maintaining the stimulus 1 valence acquired during FOC throughout SOC, which drives DAN activity. Yet, the absence of reinforcement during SOC would trigger the extinction of this very valence. It seems possible that more than one of the circuit motifs could co-exist in different MB compartments. Ultimately, not all MB compartments might be involved in SOC [3, 58], but fulfill other roles in learning.

Computational models are a highly beneficial tool for investigating the circuitry underlying SOC. Experimental validations of theoretically proposed circuit motifs would close the loop between theoretical predictions and their experimental test. However, with the available genetic tools, it is currently impossible to solely manipulate KC>DAN or MBON>DAN synapses either on the pre or post-synaptic side without affecting output onto other or input from other neurons in the network. Therefore, an experimental test of our theoretical predictions is currently difficult to achieve, underlining the importance of computational modeling.

Methods

Network Input

Odor and reward signals are provided to all models via the KCs and the DAN. Three odors are used in the experiments. In the most simple case, each exclusively activates an independent combination of (10%) of the 2000 KCs with a rate of 3 Hz to match the levels of population sparseness and low odor-response rates reported in KCs [20, 69, 70]. The first experiment combines the three odors (Figure 1 D). For each model instance, odor 1 activates a combination of 200 randomly chosen KC. KCs activated by odor 2 and odor 3 are then sequentially drawn from the combination of remaining KCs. When odors 1 and 2 are presented together during the experiments, each component of this compound odor activates a random 50% of the KCs, activated by each of the individual odors. This ensures the activated KC does not exceed 10% [19]. The second experiment aims to quantify the stability of the different circuit mechanisms against the generalization of the learned valence onto novel odors. Therefore, a varying degree of odor similarity of odor 3 with odor 1 or odor 2 is used. Odor similarity is implemented as an overlap in activated KCs for odor 1 and odor 2. If any given odor activates 200 KCs, an odor similarity of 50% would yield odors 1 and 2 to have an overlap of 100 KCs. During the joint presentation of both odors, 150 KCs would be activated.

Parameter Optimization

To increase fairness in the model comparison, all parameters that were part of the grid search were, if possible, optimized within the same boundaries and with the same step size (Table 3). Some parameters were fixed to the same value for all models to adhere to biological constraints (Table 2). KCs show very little spontaneous activity [20, 69, 70] and sparse activation [20, 69–71].

MBON rates of up to 40 Hz have been shown for one MBON [34]. We chose the initial weights for all KC>MBON synapses to achieve plausible MBON rates of no higher than 50 Hz. The upper limit of the DAN rate was 20 Hz to match the experimental literature [55].

A grid search was conducted for each model to optimize the free model parameters (Table 3). All searches contained $100^4 = 100 \cdot 10^6$ parameter combinations for testing (Table 3). The grid searches for all models were run on the same server (X86_64 architecture, Ubuntu 20.04.3). The simulation of the parameter combinations was distributed across 24 independent processes using the same random seed. The resulting data were first filtered for the fulfillment of the rate criteria for MBONs and the DAN and the performance in the FOC and SOC tests to determine all *optimal learners*, which meant identifying the parameter combinations with the same coordinates in the objective space (FOC odor 1=1, SOC odor 1=1, SOC odor 2=0.33, and SOC odor 2 0.025 for model 1).

Hypersphere sampling

Model robustness around the most central optimal parameter combination (see Results: Identifying the optimal parameters for each model) was assessed by testing random samples with increasing distance to the optimal point. We sampled the points $\mathbf{x}$ uniformly from the surface of a 4-dimensional hypersphere [72, Ch. 1.2.6] with radius r by drawing all four components independently of Gaussian distributions with a standard deviation σ and scaling with the norm $\|\mathbf{x}\|$

$$
\begin{aligned}
\sigma &\leftarrow 1 \\
\overrightarrow{x} &\leftarrow \{\mathrm{Gauss}(\sigma)\mathrm{Gauss}(\sigma), \mathrm{Gauss}(\sigma)\mathrm{Gauss}(\sigma)\} \\
\overrightarrow{x} &\leftarrow r \cdot \frac{\overrightarrow{x}}{\|\overrightarrow{x}\|}
\end{aligned}
$$

Each sample was evaluated by the indicator function

$$
\mathbf{1}_O(\overrightarrow{x}) := \begin{cases} 1 & \text{if } \overrightarrow{x} \in O \text{ optimal set, optimal performance)}, \\ 0 & \text{if } \overrightarrow{x} \notin O \ . \end{cases} \tag{6}
$$

A parameter combination $\overrightarrow{x}$ is an element of the optimal set O if it shows the same performance in the FOC and SOC tests as the central optimal parameter combination identified in the grid search for the specific model.

The code for implementing the circuit models can be obtained at https://github.com/nawrotlab/exploring_SOC_circuits.

Acknowledgements

References

[1] Christopher J Tabone and J Steven de Belle. "Second-order conditioning in Drosophila". In: *Learning & Memory* 18.4 (2011), pp. 250–253. DOI: 10.1101/lm.2035411.

[2] Christian König, Afshin Khalili, Thomas Niewalda, Shiqiang Gao, and Bertram Gerber. "An optogenetic analogue of second-order reinforcement in Drosophila". In: *Biology Letters* 15.7 (2019), p. 20190084. DOI: 10.1098/rsbl.2019.0084.

[3] Daichi Yamada, Daniel Bushey, Feng Li, Karen Hibbard, Megan Sammons, Jan Funke, Ashok Litwin-Kumar, Toshihide Hige, and Yoshinori Aso. "Hierarchical architecture of dopaminergic circuits enables second-order conditioning in Drosophila". In: *bioRxiv* (2022). DOI: 10.1101/2022.03.30.486484.

[4] El Yazid Rachad. "Neural circuit plasticity underlying learning and memory in Drosophila melanogaster: from synaptic connections to behavior". PhD thesis. Universität zu Gö"ttingen, 2923. DOI: 10.53846/goediss-9845.

[5] ME Bitterman, Randolf Menzel, Andrea Fietz, and Sabine Schäfer. "Classical conditioning of proboscis extension in honeybees (Apis mellifera)." In: *Journal of comparative psychology* 97.2 (1983), p. 107. DOI: 10.1037/0735-7036.97.2.107.

[6] Syed Abid Hussaini, Bernhard Komischke, Randolf Menzel, and Harald Lachnit. "Forward and backward second-order Pavlovian conditioning in honeybees". In: *Learning & Memory* 14.10 (2007), pp. 678–683. DOI: 10.1101/lm.471307.

[7] Makoto Mizunami, Sae Unoki, Yasuhiro Mori, Daisuke Hirashima, Ai Hatano, and Yukihisa Matsumoto. "Roles of octopaminergic and dopaminergic neurons in appetitive and aversive memory recall in an insect". In: *BMC biology* 7.1 (2009), pp. 1–16. DOI: 10.1186/1741-7007-7-46.

[8] Robert D Hawkins, Winifred Greene, and Eric R Kandel. "Classical conditioning, differential conditioning, and second-order conditioning of the Aplysia gill-withdrawal reflex in a simplified mantle organ preparation." In: *Behavioral neuroscience* 112.3 (1998), p. 636. DOI: 10.1037/0735-7044.112.3.636.

[9] Peter C Holland and Robert A Rescorla. "Second-order conditioning with food unconditioned stimulus." In: *Journal of comparative and physiological psychology* 88.1 (1975), p. 459. DOI: 10.1037/h0076219.

[10] JV Murphy and RE Miller. "Higher-order conditioning in the monkey". In: *The Journal of General Psychology* 56.1 (1957), pp. 67–72. DOI: 10.1080/00221309.1957.9918363.

[11] Michael Cook and Susan Mineka. "Second-order conditioning and overshadowing in the observational conditioning of fear in monkeys". In: *Behaviour Research and Therapy* 25.5 (1987), pp. 349–364. DOI: 10.1016/0005-7967(87)90013-1.

[12] Martin Heisenberg. "Mushroom body memoir: from maps to models". In: *Nature Reviews Neuroscience* 4.4 (2003), pp. 266–275. DOI: 10.1038/nrn1074.

[13] Martin Heisenberg. "What do the mushroom bodies do for the insect brain? An introduction". In: *Learning & memory* 5.1 (1998), pp. 1–10. DOI: 10.1101/lm.5.1.1.

[14] J Steven De Belle and Martin Heisenberg. "Associative odor learning in Drosophila abolished by chemical ablation of mushroom bodies". In: *Science* 263.5147 (1994), pp. 692–695. DOI: 10.1126/science.8303280.

[15] Troy Zars. "Behavioral functions of the insect mushroom bodies". In: *Current opinion in neurobiology* 10.6 (2000), pp. 790–795. DOI: 10.1016/S0959-4388(00)00147-1.

[16] Randolf Menzel and Martin Giurfa. "Cognitive architecture of a mini-brain: the honeybee". In: *Trends in cognitive sciences* 5.2 (2001), pp. 62–71. DOI: 10.1016/S1364-6613(00)01601-6.

[17] Makoto Mizunami, Josette M Weibrecht, and Nicholas J Strausfeld. "Mushroom bodies of the cockroach: their participation in place memory". In: *Journal of Comparative Neurology* 402.4 (1998), pp. 520–537. DOI: 10.1002/(SICI)1096-9861(19981228)402:4<520::AID-CNE6>3.0.CO;2-K.

[18] Wolfgang Rössler. "Multisensory navigation and neuronal plasticity in desert ants". In: *Trends in Neurosciences* 46.6 (2023), pp. 415–417. DOI: 10.1016/j.tins.2023.03.008.

[19] Kyle S Honegger, Robert AA Campbell, and Glenn C Turner. "Cellular-resolution population imaging reveals robust sparse coding in the Drosophila mushroom body". In: *Journal of neuroscience* 31.33 (2011), pp. 11772–11785. DOI: 10.1523/JNEUROSCI.1099-11.2011.

[20] Glenn C. Turner, Maxim Bazhenov, and Gilles Laurent. "Olfactory representations by Drosophila mushroom body neurons". In: *Journal of Neurophysiology* 99.2 (Feb. 2008), pp. 734–746. DOI: 10.1152/jn.01283.2007.

[21] Andrew C Lin, Alexei M Bygrave, Alix De Calignon, Tzumin Lee, and Gero Miesenböck. "Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination". In: *Nature neuroscience* 17.4 (2014), pp. 559–568. DOI: 10.1038/nn.3660.

[22] Makoto Mizunami, Masayuki Iwasaki, Michiko Nishikawa, and Ryuichi Okada. "Modular structures in the mushroom body of the cockroach". In: *Neuroscience letters* 229.3 (1997), pp. 153–156. DOI: 10.1016/S0304-3940(97)00438-2.

[23] Makoto Mizunami, Masayuki Iwasaki, Ryuichi Okada, and Michiko Nishikawa. "Topography of modular subunits in the mushroom bodies of the cockroach". In: *Journal of Comparative Neurology* 399.2 (1998), pp. 153–161. DOI: 10.1002/(SICI)1096-9861(19980921)399:2<153::AID-CNE1>3.0.CO;2-%23.

[24] Jürgen Rybak and Randolf Menzel. "Anatomy of the mushroom bodies in the honey bee brain: the neuronal connections of the alpha-lobe". In: *Journal of Comparative Neurology* 334.3 (1993), pp. 444–465. DOI: 10.1002/cne.903340309.

[25] Susan E Fahrbach. "Structure of the mushroom bodies of the insect brain". In: *Annu. Rev. Entomol.* 51 (2006), pp. 209–232. DOI: 10.1146/annurev.ento.51.110104.150954.

[26] Yoshinori Aso, Daisuke Hattori, Yang Yu, Rebecca M Johnston, Nirmala A Iyer, Teri-TB Ngo, Heather Dionne, LF Abbott, Richard Axel, Hiromu Tanimoto, et al. "The neuronal architecture of the mushroom body provides a logic for associative learning". In: *elife* 3 (2014), e04577. DOI: 10.7554/eLife.04577.

[27] Nobuaki K Tanaka, Hiromu Tanimoto, and Kei Ito. "Neuronal assemblies of the Drosophila mushroom body". In: *Journal of Comparative Neurology* 508.5 (2008), pp. 711–755. DOI: 10.1002/cne.21692.

[28] Martin F Strube-Bloss and Wolfgang Rössler. "Multimodal integration and stimulus categorization in putative mushroom body output neurons of the honeybee". In: *Royal Society Open Science* 5.2 (2018), p. 171785. DOI: 10.1098/rsos.171785.

[29] Yongsheng Li and Nicholas J Strausfeld. "Morphology and sensory modality of mushroom body extrinsic neurons in the brain of the cockroach, Periplaneta americana". In: *Journal of Comparative Neurology* 387.4 (1997), pp. 631–650. DOI: 10.1002/(SICI)1096-9861(19971103) 387:4<631::AID-CNE9>3.0.CO;2-3.

[30] Yoshinori Aso, Divya Sitaraman, Toshiharu Ichinose, Karla R Kaun, Katrin Vogt, Ghislain Belliart-Guérin, Pierre-Yves Plaçais, Alice A Robie, Nobuhiro Yamagata, Christopher Schnaitmann, et al. "Mushroom body output neurons encode valence and guide memory-based action selection in Drosophila". In: *Elife* 3 (2014), e04580. DOI: /10.7554/eLife.04580.

[31] David Owald and Scott Waddell. "Olfactory learning skews mushroom body output pathways to steer behavioral choice in Drosophila". In: *Current opinion in neurobiology* 35 (2015), pp. 178–184. DOI: /10.1016/j.conb.2015.10.002.

[32] Oliver Barnstedt, David Owald, Johannes Felsenberg, Ruth Brain, John-Paul Moszynski, Clifford B Talbot, Paola N Perrat, and Scott Waddell. "Memory-relevant mushroom body output synapses are cholinergic". In: *Neuron* 89.6 (2016), pp. 1237–1247. DOI: 10.1016/j. neuron.2016.02.015.

[33] David Owald, Johannes Felsenberg, Clifford B Talbot, Gaurav Das, Emmanuel Perisse, Wolf Huetteroth, and Scott Waddell. "Activity of defined mushroom body output neurons underlies learned olfactory behavior in Drosophila". In: *Neuron* 86.2 (2015), pp. 417–427. DOI: 10.1016/j.neuron.2015.03.025.

[34] Toshihide Hige, Yoshinori Aso, Mehrab N Modi, Gerald M Rubin, and Glenn C Turner. "Heterosynaptic plasticity underlies aversive olfactory learning in Drosophila". In: *Neuron* 88.5 (2015), pp. 985–998. DOI: /10.1016/j.neuron.2015.11.003.

[35] Martin F Strube-Bloss, Martin P Nawrot, and Randolf Menzel. "Mushroom body output neurons encode odor–reward associations". In: *Journal of Neuroscience* 31.8 (2011), pp. 3129–3140. DOI: 10.1523/JNEUROSCI.2583-10.2011.

[36] Yoshinori Aso and Gerald M Rubin. "Dopaminergic neurons write and update memories with cell-type-specific rules". In: *Elife* 5 (2016), e16135. DOI: /10.7554/eLife.16135.

[37] Scott Waddell. "Dopamine reveals neural circuit mechanisms of fly memory". In: *Trends in neurosciences* 33.10 (2010), pp. 457–464. DOI: /10.1016/j.tins.2010.07.001.

[38] Silke Sachse and Jennifer Beshel. "The good, the bad, and the hungry: how the central brain codes odor valence to facilitate food approach in Drosophila". In: *Current opinion in neurobiology* 40 (2016), pp. 53–58. DOI: 10.1016/j.conb.2016.06.012.

[39] Scott Waddell. "Reinforcement signalling in Drosophila; dopamine does it all after all". In: *Current opinion in neurobiology* 23.3 (2013), pp. 324–329. DOI: 10.1016/j.conb.2013.01.005.

[40] Young-Cho Kim, Hyun-Gwan Lee, Junghwa Lim, and Kyung-An Han. "Appetitive learning requires the alpha1-like octopamine receptor OAMB in the Drosophila mushroom body neurons". In: *Journal of Neuroscience* 33.4 (2013), pp. 1672–1677. DOI: 10.1523/JNEUROSCI. 3042-12.2013.

[41] Martin Schwaerzel, Maria Monastirioti, Henrike Scholz, Florence Friggi-Grelin, Serge Birman, and Martin Heisenberg. "Dopamine and octopamine differentiate between aversive and appetitive olfactory memories in Drosophila". In: *Journal of Neuroscience* 23.33 (2003), pp. 10495–10502. DOI: 10.1523/JNEUROSCI.23-33-10495.2003.

[42] Makoto Mizunami and Yukihisa Matsumoto. "Roles of octopamine and dopamine neurons for mediating appetitive and aversive signals in Pavlovian conditioning in crickets". In: *Frontiers in physiology* 8 (2017), p. 1027. DOI: 10.3389/fphys.2017.01027.

[43] Kristina V Dylla, Georg Raiser, C Giovanni Galizia, and Paul Szyszka. "Trace conditioning in Drosophila induces associative plasticity in mushroom body Kenyon cells and dopaminergic neurons". In: *Frontiers in neural circuits* 11 (2017), p. 42. DOI: 10.3389/fncir.2017.00042.

[44] Thomas Riemensperger, Thomas Völler, Patrick Stock, Erich Buchner, and André Fiala. "Punishment prediction by dopaminergic neurons in Drosophila". In: *Current Biology* 15.21 (2005), pp. 1953–1960. DOI: 10.1016/j.cub.2005.09.042.

[45] Wolfram Schultz. "Predictive reward signal of dopamine neurons". In: *Journal of neurophysiology* (1998). DOI: 10.1152/jn.1998.80.1.1.

[46] Katharina Eichler, Feng Li, Ashok Litwin-Kumar, Youngser Park, Ingrid Andrade, Casey M Schneider-Mizell, Timo Saumweber, Annina Huser, Claire Eschbach, Bertram Gerber, Richard D Fetter, James W Truman, Carey E Priebe, L F Abbott, Andreas S Thum, Marta Zlatic, and Albert Cardona. "The complete connectome of a learning and memory centre in an insect brain". In: *Nature* 548.7666 (2017), pp. 175–182. DOI: 10.1038/nature23455.

[47] Claire Eschbach, Akira Fushiki, Michael Winding, Casey M Schneider-Mizell, Mei Shao, Rebecca Arruda, Katharina Eichler, Javier Valdes-Aleman, Tomoko Ohyama, Andreas S Thum, Bertram Gerber, Richard D Fetter, James W Truman, Ashok Litwin-Kumar, Albert Cardona, and Marta Zlatic. "Recurrent architecture for adaptive regulation of learning in the insect brain". In: *Nature Neuroscience* 23.4 (2020), pp. 544–555. DOI: 10.1038/s41593-020-0607-9.

[48] Timo Saumweber, Astrid Rohwedder, Michael Schleyer, Katharina Eichler, Yi-Chun Chen, Yoshinori Aso, Albert Cardona, Claire Eschbach, Oliver Kobler, Anne Voigt, Archana Durairaja, Nino Mancini, Marta Zlatic, James W Truman, Andreas S Thum, and Bertram Gerber. "Functional architecture of reward learning in mushroom body extrinsic neurons of larval Drosophila". In: *Nature communications* 9.1 (2018), p. 1104. DOI: 10.1038/s41467-018-03130-1.

[49] Isaac Cervantes-Sandoval, Anna Phan, Molee Chakraborty, and Ronald L Davis. "Reciprocal synapses between mushroom body and dopamine neurons form a positive feedback loop required for learning". In: *Elife* 6 (2017), e23789. DOI: /10.7554/eLife.23789.

[50] Louis K Scheffer, C Shan Xu, Michal Januszewski, Zhiyuan Lu, Shin-ya Takemura, Kenneth J Hayworth, Gary B Huang, Kazunori Shinomiya, Jeremy Maitlin-Shepard, Stuart Berg, et al. "A connectome and analysis of the adult Drosophila central brain". In: *Elife* 9 (2020), e57443. DOI: /10.7554/eLife.57443.

[51] Shin-ya Takemura, Yoshinori Aso, Toshihide Hige, Allan Wong, Zhiyuan Lu, C Shan Xu, Patricia K Rivlin, Harald Hess, Ting Zhao, Toufiq Parag, et al. "A connectome of a learning and memory center in the adult Drosophila brain". In: *Elife* 6 (2017), e26975. DOI: 10.7554/eLife.26975.

[52] Toshiharu Ichinose, Yoshinori Aso, Nobuhiro Yamagata, Ayako Abe, Gerald M Rubin, and Hiromu Tanimoto. "Reward signal in a recurrent circuit drives appetitive long-term memory formation". In: *Elife* 4 (2015), e10719.

[53] Julien Séjourné, Pierre-Yves Plaçais, Yoshinori Aso, Igor Siwanowicz, Séverine Trannoy, Vladimiros Thoma, Stevanus R Tedjakumala, Gerald M Rubin, Paul Tchénio, Kei Ito, et al. "Mushroom body efferent neurons responsible for aversive olfactory memory retrieval in Drosophila". In: *Nature neuroscience* 14.7 (2011), pp. 903–910. DOI: /10.1038/nn.2846.

[54] Clare E Hancock, Vahid Rostami, El Yazid Rachad, Stephan H Deimel, Martin P Nawrot, and André Fiala. "Visualization of learning-induced synaptic plasticity in output neurons of the Drosophila mushroom body γ-lobe". In: *Scientific Reports* 12.1 (2022), p. 10421. DOI: 10.1038/s41598-022-14413-5.

[55] Cheng Huang, Junjie Luo, Seung Je Woo, Lucas Roitman, Jizhou Li, Vincent Pieribone, Madhuvanthi Kannan, Ganesh Vasan, and Mark Schnitzer. "Dopamine signals integrate innate and learnt valences to regulate memory dynamics". In: *Research Square* (2022). DOI: 10.21203/rs.3.rs-1915648/v1.

[56] Björn Brembs and Martin Heisenberg. "Conditioning with compound stimuli in Drosophila melanogaster in the flight simulator". In: *Journal of experimental biology* 204.16 (2001), pp. 2849–2859. DOI: 10.1242/jeb.204.16.2849.

[57] Faramarz Faghihi, Ahmed A Moustafa, Ralf Heinrich, and Florentin Wörgötter. "A computational model of conditioning inspired by Drosophila olfactory system". In: *Neural Networks* 87 (2017), pp. 96–108. DOI: 10.1016/j.neunet.2016.11.002.

[58] Linnie Jiang and Ashok Litwin-Kumar. "Models of heterogeneous dopamine signaling in an insect learning and memory center". In: *PLoS Computational Biology* 17.8 (2021), e1009205. DOI: /10.1371/journal.pcbi.1009205.

[59] Julia E Manoim, Andrew M Davidson, Shirley Weiss, Toshihide Hige, and Moshe Parnas. "Lateral axonal modulation is required for stimulus-specific olfactory conditioning in Drosophila". In: *Current Biology* 32.20 (2022), pp. 4438–4450. DOI: 10.1016/j.cub.2022.09.007.

[60] Noa Bielopolski, Hoger Amin, Anthi A Apostolopoulou, Eyal Rozenfeld, Hadas Lerner, Wolf Huetteroth, Andrew C Lin, and Moshe Parnas. "Inhibitory muscarinic acetylcholine receptors enhance aversive olfactory learning in adult Drosophila". In: *eLife* 8 (2019), e48264. DOI: 10.7554/eLife.48264.

[61] Amanda Crocker, Xiao-Juan Guan, Coleen T. Murphy, and Mala Murthy. "Cell-Type-Specific Transcriptome Analysis in the Drosophila Mushroom Body Reveals Memory-Related Changes in Gene Expression". In: *Cell Reports* 15.7 (2016), pp. 1580–1596. DOI: 10.1016/j.celrep.2016.04.046.

[62] Vincent Croset, Christoph D Treiber, and Scott Waddell. "Cellular diversity in the *Drosophila* midbrain revealed by single-cell transcriptomics". In: *eLife* 7 (2018), e34550. DOI: 10.7554/eLife.34550.

[63] Hailing Su and Diane K O'Dowd. "Fast synaptic currents in Drosophila mushroom body Kenyon cells are mediated by α-bungarotoxin-sensitive nicotinic acetylcholine receptors and picrotoxin-sensitive GABA receptors". In: (). DOI: 10.1523/JNEUROSCI.23-27-09246.2003.

[64] Mark E Bouton. "Context and behavioral processes in extinction". In: *Learning & memory* 11.5 (2004), pp. 485–494. DOI: 10.1101/lm.78804.

[65] Lingling Wang, Qi Yang, Binyan Lu, Lianzhang Wang, Yi Zhong, and Qian Li. "A behavioral paradigm to study the persistence of reward memory extinction in Drosophila". In: *Journal of genetics and genomics* 46.12 (2019), pp. 599–601. DOI: 10.1016/j.jgg.2019.11.001.

[66] Yukinori Hirano, Kunio Ihara, Tomoko Masuda, Takuya Yamamoto, Ikuko Iwata, Aya Takahashi, Hiroko Awata, Naosuke Nakamura, Mai Takakura, Yusuke Suzuki, Junjiro Horiuchi, Hiroyuki Okuno, and Minoru Saitoe. "Shifting transcriptional machinery is required for long-term memory maintenance and modification in Drosophila mushroom bodies". In: *Nature communications* 7.1 (2016), p. 13471. DOI: 10.1038/ncomms13471.

[67] Johannes Felsenberg, Oliver Barnstedt, Paola Cognigni, Suewei Lin, and Scott Waddell. "Re-evaluation of learned information in Drosophila". In: *Nature* 544.7649 (2017), pp. 240–244. DOI: 10.1038/nature21716.

[68] Johannes Felsenberg, Pedro F Jacob, Thomas Walker, Oliver Barnstedt, Amelia J Edmondson-Stait, Markus W Pleijzier, Nils Otto, Philipp Schlegel, Nadiya Sharifi, Emmanuel Perisse, et al. "Integration of parallel opposing memories underlies memory extinction". In: *Cell* 175.3 (2018), pp. 709–722. DOI: /10.1016/j.cell.2018.08.021.

[69] Javier Perez-Orive, Ofer Mazor, Glenn C Turner, Stijn Cassenaer, Rachel I Wilson, and Gilles Laurent. "Oscillations and sparsening of odor representations in the mushroom body". In: *Science* 297.5580 (2002), pp. 359–365. DOI: 10.1126/science.1070502.

[70] Iori Ito, Rose Chik-Ying Ong, Baranidharan Raman, and Mark Stopfer. "Sparse odor representation and olfactory learning". In: *Nature neuroscience* 11.10 (2008), pp. 1177–1184. DOI: 10.1038/nn.2192.

[71] Eleftheria Vrontou, Lukas N Groschner, Susanne Szydlowski, Ruth Brain, Alina Krebbers, and Gero Miesenböck. "Response competition between neurons and antineurons in the mushroom body". In: *Current Biology* 31.22 (2021), pp. 4911–4922. DOI: 10.1016/j.cub.2021.09.008.

[72] Werner Krauth. *Statistical mechanics : algorithms and computations*. Oxford master series in statistical, computational, and theoretical physics. Oxford: Oxford University Press, 2006.

Model 1
$$R_{DAN}(t) = Ip_{ext}(t) + \overrightarrow{R}_{KC}(t) \cdot \overrightarrow{w}_{KC>DAN}(t) \quad (7)$$

$$w_{KC>KC}^{ij}(t+1) = \begin{cases} w_{KC>KC}^{ij}(t) + LR_{KC>KC}, & \text{if } R_{KC}^i(t) > KC_{baseline} \text{ and} \\ & R_{KC}^i(t) > KC_{baseline} \\ w_{KC>KC}^{ij}(t), & \text{otherwise} \end{cases} \quad (8)$$

Model 2
$$R_{DAN}(t) = Ip_{ext}(t) + \overrightarrow{R}_{KC}(t) \cdot \overrightarrow{w}_{KC>DAN}(t) \quad (9)$$

$$w_{KC>DAN}^{i}(t+1) = \begin{cases} w_{KC>DAN}^{i}(t) + LR_{KC>DAN} \cdot R_{DAN}(t), & \text{if } R_{KC}^i(t) > 0 \\ w_{KC>DAN}^{i}(t), & \text{otherwise} \end{cases} \quad (10)$$

Model 3
$$R_{DAN}(t) = Ip_{ext}(t) + R_{MBON+}(t) \cdot w_{MBON+>DAN}(t) - R_{MBON-}(t) \cdot w_{MBON->DAN}(t) \quad (11)$$

Model 4
$$R_{DAN}(t) = Ip_{ext}(t) + R_{DANbaseline} - R_{MBON-}(t) \cdot w_{MBON->DAN}(t) \quad (12)$$

Model 5
$$R_{DAN}(t) = Ip_{ext}(t) + R_{MBON+}(t) \cdot \overrightarrow{w}_{MBON+>DAN}(t) \quad (13)$$

$$\overrightarrow{w}_{MBON+>DAN}(t+1) = \begin{cases} \overrightarrow{w}_{MBON+>DAN}(t) + LR_{MBON+>DAN} \cdot R_{DAN}(t), & \text{if } R_{MBON+}(t) > 0 \\ \overrightarrow{w}_{MBON+>DAN}(t), & \text{otherwise} \end{cases} \quad (14)$$

Table 1: **Equations underlying the different models.**

19
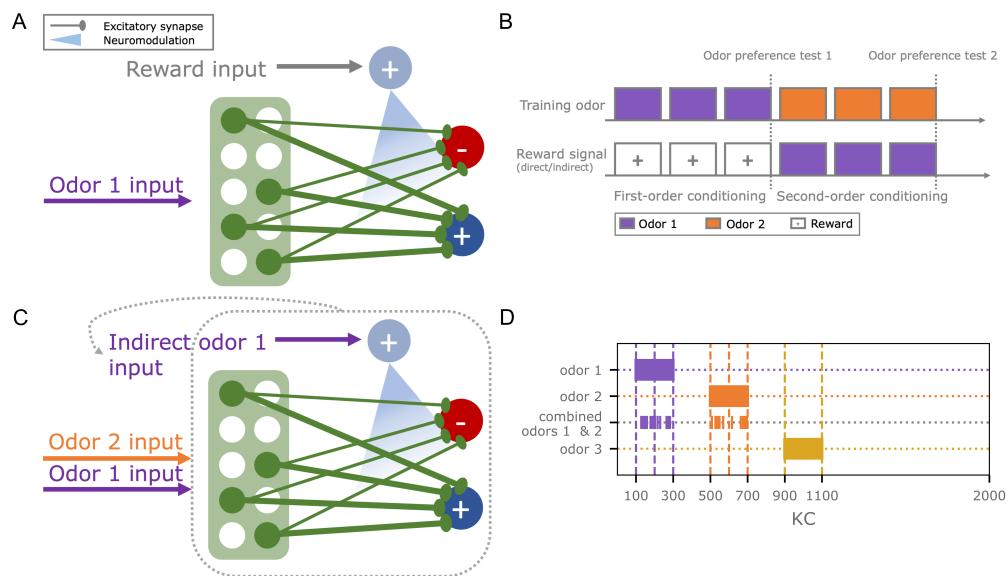
Figure 1: **Experimental design for testing second-order conditioning.** (A) Basic circuit motif for first-order conditioning, consisting of 2000 Kenyon cells (green), two output neurons (dark blue, red), and a single dopaminergic neuron (light blue). The co-occurrence of odor and reward input elicits plasticity at the MB output synapses. (B) The experimental paradigm consists of two phases (first and second-order conditioning). During first-order conditioning, odor 1 is paired with a reward. Subsequently, a novel odor 2 is paired with odor 1 during second-order conditioning. Odor valences are tested after first and second-order conditioning. (C) During second-order conditioning, the dopaminergic neuron (light blue) is indirectly activated by the previously trained odor 1 and paired with odor 2. We test different candidate mechanisms for this indirect activation of the dopaminergic neuron via the Kenyon cells (green) or the mushroom body output neurons (dark blue, red). (D) Initially, three non-overlapping odors were used in the experiments. Odors are encoded as Kenyon cell activity patterns. The joint presentation of odor 1 and odor 2 during second-order conditioning retains a randomly chosen 50% of the individual odor representations to maintain the same overall activation as with individually presented odors.

20

	Model circuit	FOC	SOC	Generalization
Model 1		✓		0.05
Model 2		✓	✓	0.03
Model 3		✓	✓	0.00
Model 4		✓	✓	0.52
Model 5		✓	✓	1.00

Figure 2: **Second-order conditioning in different circuit motifs.** Five different circuits were tested for their performance in first (FOC) and second-order conditioning (SOC) and the extent to which the odor-reward association generalizes to another novel odor. All circuits are constructed with 2000 Kenyon cells (green), two mushroom body output neurons (dark blue, red), and a single dopaminergic neuron (light blue), targeting the synapses between Kenyon cells and mushroom body output neurons. Additional feed-forward connections from the Kenyon cells (model 1, model 2) or feedback connections from the mushroom body output neurons onto the dopaminergic neuron (model 3, model 4, model 5) are implemented in the different circuits.

21
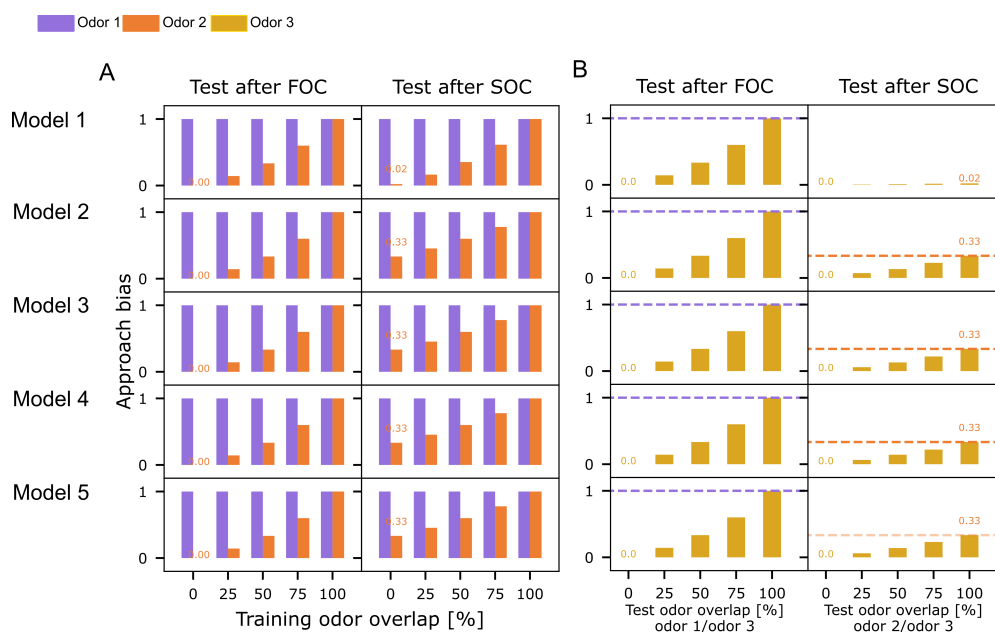
Figure 3: **Reward generalization for overlapping training or test odors.** (A) All five models were tested for their odor approach bias (Eqn. 5) to odors 1 (purple) and 2 (orange) after first (FOC) and second-order conditioning (SOC). The overlap between odors 1 and 2 was varied. (B) All models were tested for their approach bias (Eqn. 5) to odor 3, after training with non-overlapping odors 1 and 2. Overlap between odors 3 and 1 or 2 was varied, respectively. Orange and purple dashed lines indicate each model's FOC and SOC performance from an experiment without odor overlap as a reference (always the first bar).
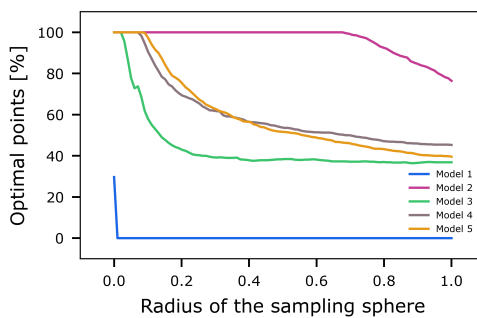
Figure 4: **Comparison of parameter robustness between models.** The five models were tested for the stability of their second-order learning performance (Eqn. 5) when their optimal parameter combinations were collectively shifted away from their optimum, which we used as the central starting point for a hypersphere with radius=0. We incrementally increased the radius of the hypersphere from 0 to 1 with 100 steps in a linear fashion and sampled 700 data points from its surface for each resulting radius.
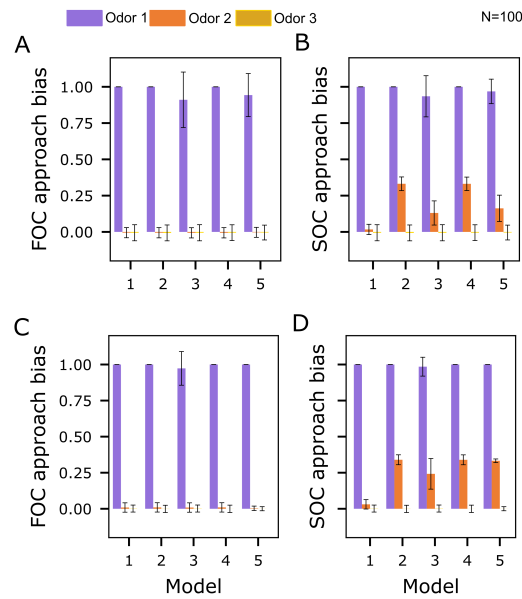
Figure 5: **Robustness of the learning performance across variations of the KC>MBON connectivity.** (A) Mean approach biases (error bars denote standard deviations, Eqn. 5) to odors 1, 2, and 3 after first order conditioning (FOC) for $N = 100$ model instances with a varying number of KC>MBON synapses, between 25% and 100%, drawn randomly from a uniform distribution. (B) Mean approach biases (Eqn. 5) to odors 1, 2, and 3 after second-order conditioning (SOC) for $N = 100$ model instances with a randomly chosen number of KC>MBON synapses, between 25% and 100%. (C) Mean approach biases (Eqn. 5) to odors 1, 2, and 3 after first order conditioning (FOC) for N=100 model instances with a randomly chosen initial KC>MBON weights (full connectivity), varying ±5% around the default value (Table 2). (D) Mean approach biases (Eqn. 5) to odors 1, 2, and 3 after second-order conditioning (SOC) for $N = 100$ model instances with a randomly chosen initial KC>MBON weights (full connectivity), varying ±5% around the default value (Table 2).

Parameter	Value
seed	999
trials_FOC	3
trials_SOC	3
num_KC	2000
$KC_{baseline}$	0 Hz
odor_activation	3 Hz
odor_FOC	odor 1
odor_SOC	odor 1_2
$init_{KC>MBON^+}$	0.083
$init_{KC>MBON^-}$	0.083

Table 2: **Fixed parameters shared between models.**

Parameter	min	max	optimum
Model 1			
$init_{KC>KC}$	0	0	0
$init_{KC>DAN}$	0.0	0.001	0.001
$LR_{KC>MBON^-}$	0.001	0.004	0.004
$LR_{KC>KC}$	0.0	0.004	0.000162
DAN_activation	1.0	10.0	7.272727
Model 2			
$init_{KC>DAN}$	0.0	0.001	0.000505
$LR_{KC>MBON^-}$	0.001	0.004	0.00333
$LR_{KC>DAN}$	0.0	0.001	0.000677
DAN_activation	1.0	10.0	5.727273
Model 3			
$init_{MBON^+>DAN}$	0.0	1.0	0.272727
$init_{MBON^->DAN}$	0.0	1.0	0.262626
$LR_{KC>MBON^-}$	0.001	0.004	0.003121
DAN_activation	1.0	10.0	5.00
Model 4			
$init_{MBON^->DAN}$	0.0	0.5	0.131313
$LR_{KC>MBON^-}$	0.001	0.004	0.003182
$R_{DANbaseline}$	0.0	30.0	11.212121
DAN_activation	1.0	10.0	3.727273
Model 5			
$init_{MBON^+>DAN}$	0.0	0.4	0.08008
$LR_{KC>MBON^-}$	0.001	0.004	0.003182
$LR_{MBON^+>DAN}$	0.001	0.01	0.003
DAN_activation	1.0	10.0	4.727273

Table 3: **Optimized model parameters.** For all models, 100 equally distributed values per parameter between the minimum and maximum values were used to construct a regular grid of parameter combinations.

4   Chapter: Adaptive behavior

4.1   A realistic locomotory model of Drosophila larva for behavioral simulations

A realistic locomotory model of *Drosophila* larva for behavioral simulations.

Panagiotis Sakagiannis[1,*]**, Anna-Maria Jürgensen**[1]**, and Martin Paul Nawrot**[1]

[1]Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Germany

The *Drosophila* **larva is extensively used as model species in experiments where behavior is recorded via tracking equipment and evaluated via population-level metrics. Although larva locomotion neuromechanics have been studied in detail, no comprehensive model has been proposed for realistic simulations of foraging experiments directly comparable to tracked recordings. Here we present a virtual larva for simulating autonomous behavior, fitting empirical observations of spatial and temporal kinematics. We propose a trilayer behavior-based control architecture for larva foraging, allowing to accommodate increasingly complex behaviors. At the basic level, forward crawling and lateral bending are generated via coupled, interfering oscillatory processes under the control of an intermittency module, alternating between crawling bouts and pauses. Next, navigation in olfactory environments is achieved via active sensing and top-down modulation of bending dynamics by concentration changes. Finally, adaptation at the highest level entails associative learning. We could accurately reproduce behavioral experiments on autonomous free exploration, chemotaxis, and odor preference testing. Inter-individual variability is preserved across virtual larva populations allowing for single animal and population studies. Our model is ideally suited to interface with neural circuit models of sensation, memory formation and retrieval, and spatial navigation.**

larva foraging | control architecture | exploratory behavior | chemotaxis | olfactory preference | autonomous agent

Introduction

Building virtual agents that behave indistinguishably from living organisms is an endeavour at the interface of behavioral and cognitive neuroscience, artificial intelligence and behavior-based robotics. The pursued level of similarity to real animal behavior defines the target level of model abstraction, which can vary from low-level detailed neuroscientific models to high-level abstract cognitive architectures. Recent advances in tracking equipment have allowed highly resolved recordings of larva populations under diverse foraging conditions (1–4). Here, we ask how to build virtual larvae behaving indistinguishably from the real ones as captured by established larva tracking protocols. To this end we suggest a layered behavioral control architecture (Fig. 1). At the basic layer, we propose a locomotory model that features a bisegmental virtual larva body, coupled crawling and bending oscillators, and crawling intermittency. Integrating previously suggested features and novel data-driven hypothesis we compare diverse configurations against empirical data in order to choose the optimal locomotory model. Expanding our model to the upper layers of the control architecture enables increasingly more complex foraging-related behaviors (Fig. 1). In virtual larva populations with realistic inter-individual variability we evaluate model performance in simulations of several established behavioral paradigms.

Drosophila larvae possess a fairly tractable behavioral reper-

toire that is consistent across the 4-5 days of the larval life stages (5) and controlled by a conserved neural circuit structure throughout development (6), making it a formidable system for studying behavioral control and decisions (7). Most of the larval time is dedicated to foraging the environment for suitable nutrients while avoiding danger. Foraging consists of a combination of more basic behaviors: crawling, turning, digging into the substrate (8), and feeding, the latter even including cannibalizing conspecifics in extreme cases (9). In the absence of available food resources, larvae have to engage in free exploration to locate food patches (10). This behavior is intermittent, meaning it consists of bouts of activity interspersed by brief pauses generated via cessation and re-initiation of crawling (11), a property also reported for adult fly behavior (12, 13). Salient olfactory cues can trigger chemotaxis during which larvae employ active sensing to navigate along chemical gradients (14). Finally, novel odorants coupled to food reward induce olfactory learning enabling long-term behavioral adaptations (15–19). After reaching critical mass for pupation, homeostatic signals switch behavior towards food aversion, hypermobility and collaborative burrowing (20), terminating the feeding state and leading to pupation and metamorphosis.

Statistical regularities that govern foraging behavior have been unveiled by analysis both at the microscale of body kinematics and at the macroscale of larva trajectories (2, 21, 22). Crawling and turning have been in the main focus of recent studies (10, 23, 24) whereas tracking studies of feeding behavior remain scarce (25). Both, crawling and feeding behavior are indisputably of oscillatory nature (23, 26, 27) controlled by central pattern generating circuits. With respect to turning, it is still debated whether individual turns should be considered as discrete reorientation events that are temporally non-overlapping with crawling bouts (10), or whether turning occurs in an oscillatory fashion generating turns both during crawling and during pauses (24, 28). The latter is supported by detailed eigenshape analysis confirming that larvae rarely crawl straight, rather forward locomotion is always accompanied by continuous small amplitude lateral bending (29). It follows that crawling does not exclude bending rather the two strictly co-occur. In contrast, both feeding and crawling movements require mouth-hook motion recruiting the same effector system, thus they can be considered competing, mutually exclusive behaviors (7, 30). Finally, it is unclear whether bending and feeding can overlap.

Modulation of exploratory behavior under salient olfactory

*To whom correspondence should be addressed.
e-mail : p.sakagiannis@uni-koeln.de
webpage : http://computational-systems-neuroscience.de/

input during positive and negative chemotaxis has been extensively studied (14, 16, 31, 32). There is general consensus that during appetitive chemotaxis increasing odor concentrations suppress pauses and turns leading to longer lasting bouts of activity, while turns are promoted during down gradient navigation. Turning is biased by the detection of minor concentration changes during lateral bending, a process described as active sensing (24, 28, 33). The opposite effect has been reported during aversive chemotactic behavior whether due to punishment or reward omission (34). Larva behavioral preference under conflicting olfactory stimuli has been established as a population-level metric in multiple settings from quantifying the formation of memory after associative learning (15–17, 35, 36) to detecting individual differences in genetically identical larva strains (37). This sensory-driven behavioral modulation does not seem to be affected by social cues, justifying the study of individual larva kinematics even during population-level experiments (38).

Results

Kinematic analysis of larva locomotion. We start out with the kinematic analysis of experimental larva trajectories and body postures in order to infer and parametrize several aspects of larva locomotion that will inform our modeling approach. Using diverse metrics that capture spatial and temporal dynamics we specifically assess the oscillation of forward velocity during individual peristaltic strides, the influence of this oscillation on lateral bending, the intermittent nature of crawling, and the inter-individual variability of a number of locomotion-related parameters across different larvae.

Color annotation of larva trajectories in Fig. 2 A illustrates nicely that forward locomotion consists of consecutive steps (strides) that are characterized by an alternating increase/decrease of the locomotion velocity $\hat{v}$. To characterize this oscillation we detected all strides performed by an individual animal and verified their stereotypical structure in terms of stride duration, resulting body displacement, and the phase-dependence of $\hat{v}$ (Fig. 2 D). Scaling both, displacement d_{str} and velocity to the individual larval body-length increases stride stereotypicality independent of larva size. This analysis justifies an oscillatory model process (crawler,) that generates forward-velocity in subsequent stride cycles (crawler, Fig. 1 B) is introduced for model comparison.

Our analysis of orientation velocity $\dot{\hat{\theta}}_{or}$ during crawling strides reveals that it is phasically coupled to the stride phase exhibiting an increase around $\frac{3\pi}{2}$ of the oscillatory cycle (Fig. 2 H). This implies phasic interference of the crawling and the lateral bending mechanism (Fig. 1 B), which is incorporated in our model (see Material and Methods). A plausible mechanistic explanation featuring bodily interference of crawling and bending is suggested in the discussion.

Larvae transiently pause crawling before re-assuming it resulting in sequences of concatenated strides (stridechains) intermitted by brief crawl-pauses (Fig.2 C). We analyzed the distributions of stridechain length and pause duration in the experimental dataset. The limited duration of the recordings (3') does not allow assessment of individual differences therefore the stride chain and pause bouts over a population of 200 larvae have been pooled together. Testing power-law, exponential and log-normal distributions revealed the highest quality of fit for the log-normal distribution for both parameters (see
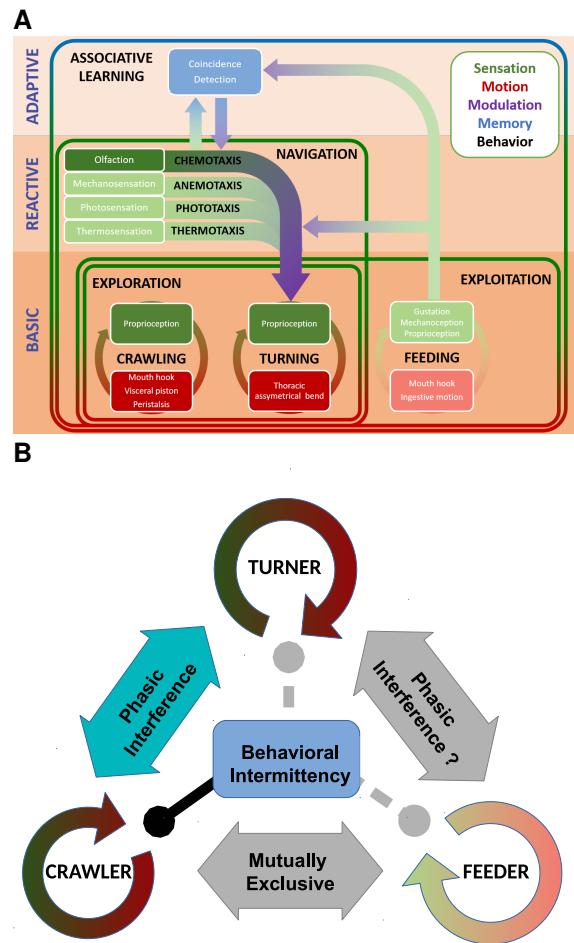
Fig. 1. Behavioral control architecture for larva foraging. A: In the trilayer control architecture the bottom layer consists of three basic sensorimotor effectors that constitute the locomotory model. The intermediate layer features innate reactive behavior in response to unexpected environmental stimuli. The top layer allows for behavioral adaptation through experience. Framed areas denote more complex behavioral modes that require subsumption of subordinate modes. Light-colored modules are plausible extensions described in Discussion. **B:** The locomotory model at the bottom layer of the architecture. Oscillatory behaviors are either phasically coupled or mutually exclusive. Initiation/cessation of oscillation is controlled by an intermittency module. Light-colored modules are plausible extensions described in Discussion.

Material and Methods). In our model we propose a behavioral intermittency module (Fig.1 B) that samples from the empirically fitted model distributions and controls cessation and re-initiation of crawling (see Material and Methods).

Behavioral architecture. Larval behavior is hierarchically structured in the sense that simple behavioral motifs such as crawling, bending and feeding motions can be integrated into more complex behavioral modes such as exploration, taxis and exploitation. It has been proposed that animal behavioral hierarchy is reflected by the underlying neuroanatomy as a hierarchy of nested sensorimotor loops (39). A functional modeling paradigm that exploits this idea regards the neural system
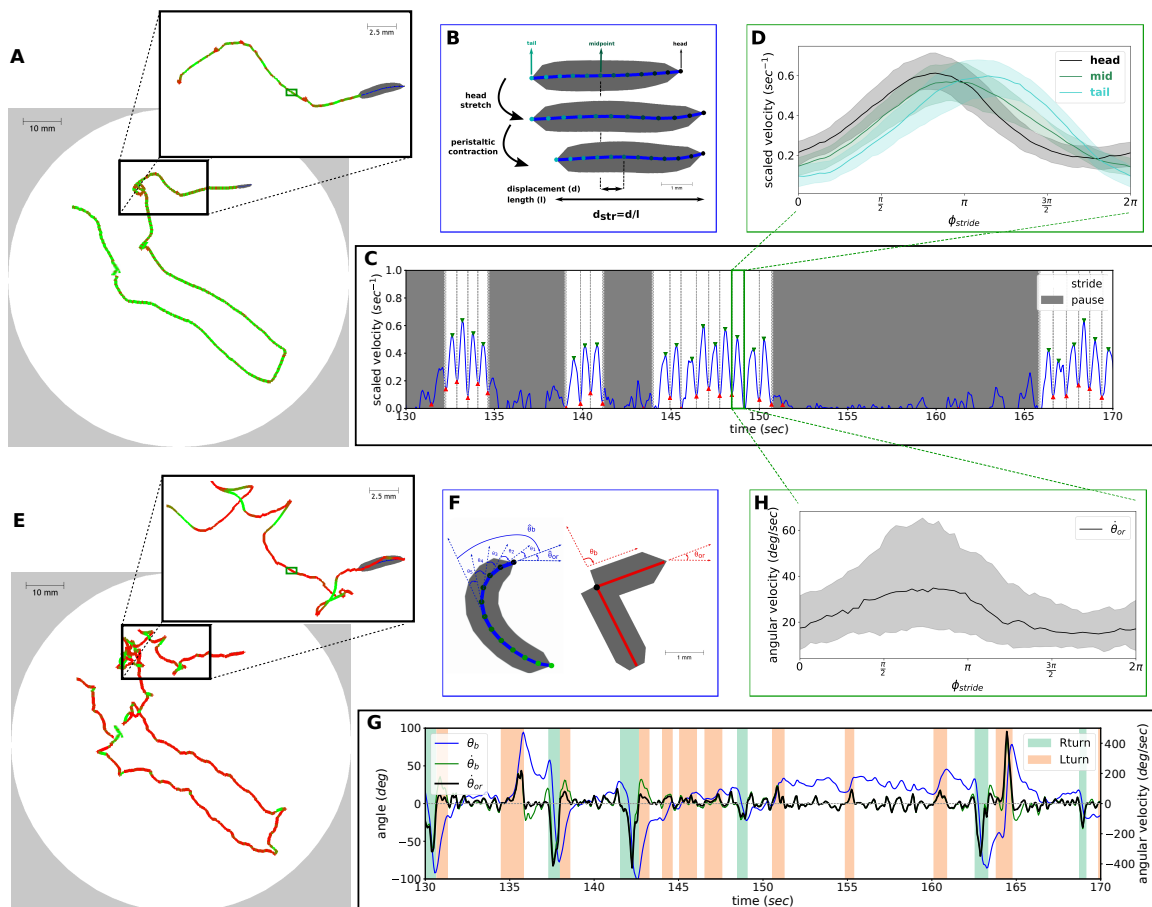
Fig. 2. Kinematic analysis of the *Drosophila* larva in locomotion. A Individual larva trajectory tracking a rear point along the midline of the animal. Trajectory color denotes the forward velocity $\hat{v}$ from 0 (red) to maximum (green). Inset focuses on the track slice analyzed in C and G. Dark green rectangle denotes a single stride described in B. **B:** Sketch of a single crawling stride. The larva first stretches its head forward, anchors it to the substrate and then drags its body forward via peristaltic contraction. Scaled stride displacement d_{str} is defined as the resulting displacement d scaled by the body-length l. **C:** Forward velocity $\hat{v}$ during the 40 s track slice selected in A, inset. Green and red markers denote the local minima and maxima used for stride annotation. Individual strides are circumscribed by vertical dashed lines. Successive strides constitute uninterrupted stridechains (white). Epochs that do not show any strides are annotated as crawl-pauses (gray). Velocity is scaled to the larva body-length. **D:** Forward velocity of head, midpoint and tail as a function of the stride cycle. All detected strides of the tracked larva have been interpolated to a stride oscillation cycle of period 2π. Solid lines denote the median, shaded areas the lower and upper quartiles. Velocity is scaled to the larva body-length. **E:** Same trajectory as in (A) now tracking the head segment. Color denotes the orientation angular velocity $\dot{\theta}_{or}$ from 0 (red) to maximum (green). **F:** Definition of bending angle θ_b and orientation angle θ_{or} for the original 12-segment (blue) and the simplified 2-segment (red) larvae. **G:** Three angular parameters during the same track slice shown in (C). Bending angle $\hat{\theta}_b$, bend and orientation angular velocities $\dot{\hat{\theta}}_b$, $\dot{\hat{\theta}}_{or}$ are shown. Background shadings denote left and right turning bouts. For illustration purposes only turns resulting in a change of orientation angle $\Delta\hat{\theta}_{or} > 20°$ are shown. **H:** Same as D but showing the average absolute orientation angular velocity $\hat{\theta}_{or}$ during the stride cycle.

as a layered control (subsumption) architecture (40) where low-level stereotyped reflexive and repetitive behaviors are autonomously generated by localized peripheral motor circuitry while multisynaptic loops involving more centralized neural circuits only act as top-down modulators on the local circuits in order to coordinate global and complex behavioral control. The central idea is that energy-efficient decentralized neural control is the rule, while higher centers are recruited only when more extensive integration is needed. Furthermore there are only limited degrees of freedom by which higher layers can influence local sensorimotor loops e.g. by starting/stopping or accelerating/decelerating their autonomous function (see Discussion). Layered control architectures have been used mainly

in behavior-based robotics allowing sequential calibration and modular integration of partial neuroscientific models under a common framework (39, 40).

We here propose a trilayer behavioral control architecture for *Drosophila* larva foraging as illustrated in Fig. 1 A. The bottom layer comprises three basic stereotyped behaviors: crawling, turning and feeding. Each is realized by a low-level local sensorimotor loop between motor effectors and sensory feedback - mainly proprioception and mechanoreception - and additionally gustatory input in the case of feeding. Integration of these basic behaviors gives rise to composite behaviors. Exploration in stimulus-free conditions requires crawling and turning while integration of all three basic elements allows

exploitation of a nutritious substrate. The intermediate layer introduces salient sensory stimulation of different modalities therefore allows reactive behavior in the face of presented risks and opportunities. Modulation of exploratory behavior under appetitive or aversive stimulation enables coherent navigation along sensory gradients. Here, we consider odor-driven chemotaxis as a process of active sensing in which the larva constantly samples the odor concentration during lateral bending motions enabling odorscape navigation and odor source discovery. At the top layer, associative learning during prior experience biases the evaluation of recurring sensory stimuli resulting in modulation of innate odor valence and providing experience-dependent adaptation of navigation.

The proposed behavioral architecture is naive to the underlying neural mechanisms that generate the respective behaviors. We hereby populate the architecture with diverse configurations of candidate mechanisms and explore autonomous and integrated control exerted by each subsequent layer in increasingly-complex behavioral simulations from exploration to chemotaxis to adaptive learning in odor preference experiments.

Locomotory model. We simplify the larva to a two-segment body (Fig. 2 F). This abstraction allows describing the body state at any point in time through only three parameters. These are (i) position of a selected midpoint (Fig. 9), (ii) absolute orientation of the front segment θ_{or}, and (iii) bending angle θ_b between the front and rear segments. This approach is in line with the common practice of quantifying larva bending via a single angle between one anterior and one posterior vector, although the chosen vectors vary between study groups (14, 34, 41). Body dynamics as analyzed in the empirical data depends on instantaneous linear (Fig. 2 A-C) and angular velocities (Fig. 2 E-G) generated through crawling and bending respectively. Kinematic analysis on the locomotion of these simplified real larvae allows us to realistically calibrate bisegmental virtual larvae resembling the real ones in multiple spatial and temporal parameters (see Materials and Methods). For a demonstration of the larva-body bisegmental simplification see Video 1. The individual trajectory depicted in Fig. 2 can be seen at its full length in Video 3, while the slice depicted in the inset is shown in Video 4.

Locomotion of the bisegmental body is achieved via dynamic coupling of forward crawling and lateral body bending. Crawling moves the midpoint along the orientation vector. Bending reorients the front segment by rotation around the midpoint. Forward displacement restores θ_b back to 0 by gradually aligning the rear segment to the orientation axis.

We adopt the hypothesis that attributes alternating left-right bending to an underlying oscillatory neural process and compare the lateral oscillator model described by Wystrach et al. (24) to a simple sinusoidal oscillator. This model (turner) applies torque oscillations on the two-segment body, which acts as a restorative torsional spring (see Materials and Methods). The turner can be coupled to a second oscillator producing crawling strides (crawler) in the form of linear velocity oscillations as suggested by our kinematic analysis. Under the constraints of the subsumption architecture paradigm these effectors can be influenced by higher-order circuits only in a limited number of ways. Frequency modulation and stride initiation/cessation are the only possible top-down modulations on the crawler. Likewise, the turner can receive olfactory input

affecting both the frequency and the amplitude of oscillation.

To quantify the contribution of each of the locomotory model's components to replicating real larva kinematics we perform a broad model comparison study across diverse model configurations (Fig. 3). The behavioral metrics used for evaluation are structured in five categories covering angular and forward motion, reorientation, spatial dispersion and stride-cycle structure. Overall, the bisegmental is superior to the single-segment body in capturing angular motion metrics, predominantly due to the decoupling of the bending velocity $\dot{\theta}_b$ from the orientation velocity $\dot{\theta}_{or}$ (in the single segment body these are considered identical). A realistic crawling oscillation allows much more accurate assessment of stride-cycle dependent metrics compared to non-oscillatory constant-speed forward motion (despite the lack of a velocity oscillation in the latter case, strides might still be detected due to noise). Finally, both oscillator coupling and crawl-intermittency contribute to better fitting of the empirical data. Our final model is represented by the far bottom right column, exhibiting the smallest error.

The detailed structure and optimal configuration for the basic layer of the behavioral architecture as suggested by model comparison is illustrated in Fig. 1 B, specifying the interplay between the oscillators. Two additional features are implemented based on our empirical analyses. First, the crawler and turner are phasically coupled such that turning is suppressed during a defined phase interval of the stride cycle, reflecting our finding in Fig. 2 H. Second, crawler-generated chains of concatenated strides (stridechains) are intermitted by brief pauses during which the crawler-induced interference is lifted, resulting in more acute turning events. The duration distribution of stridechains and pauses are fitted to the empirical data. The pipeline for model calibration is described in detail in Material and Methods. For a demonstration of the locomotory model in different configurations see Video 2. We note that in the current implementation effectors receive no sensory input from the environment. Sensory feedback is introduced only at the intermediate reactive layer (Fig. 1 A) via an olfactory sensor located at the front end of the body (head). All model parameters are shown in Table 1.

Inter-individual differences are crucial for achieving realistic population-level behavior. To capture variability across larva we computed three crawl-related parameters, body-length l, crawling frequency f_c, and scaled stride displacement d_{str} across a population of 200 experimental larvae. The univariate and bivariate empirical distributions are illustrated in blue in Fig. 4. When generating virtual larva populations for the behavioral simulations we sample all three parameters from a multivariate Gaussian distribution fitted to the empirical data. This preserves the linear correlation structure. A generated set of the three parameters completely defines the crawling oscillation of an individual virtual larva. The bivariate projections are shown in red in Fig. 4 with parameters given in Table 1 in blue.

Simulation of behavioral experiments. The layered behavioral architecture in Fig.1 is exploited in modeling as it justifies sequential calibration and evaluation of subsequent behavioral layers from the bottom to the top, as for any specific behavior to be successfully realized, only integration up to a certain layer is required. In this section we simulate increasingly more complex behavioral experiments using virtual larva popula-
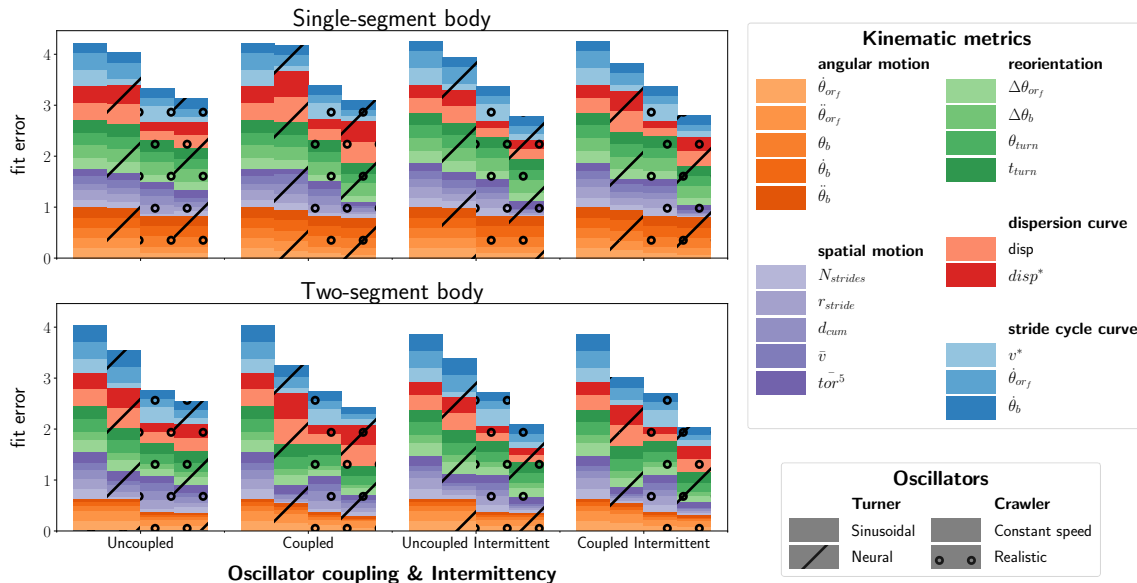
Fig. 3. Model comparison. A total of 32 model configurations have been tested in 3-minute free-exploration simulations of 50 virtual larvae. The controlled body either has a single (top) or two segments (bottom). The crawler and turner modules implemented in each model are illustrated in the bottom legend (simple sinusoidal vs neural turner oscillator and constant-speed vs analysis-fitted crawler oscillator). The oscillator coupling and crawling-intermittency mode is illustrated in the x axis (coupled vs uncoupled, intermittent vs uninterrupted). Noise levels are identical in all models as shown in Tab. 1. Derived metrics are compared to the experimentally measured in 3-minute free-exploration recordings of 200 real larvae. The evaluation metrics are grouped in 5 categories as shown in the top legend. For metrics in the angular, spatial motion and reorientation categories the KS distance between the simulated and the empirical distributions, pooled across individuals, has been computed while for metrics in the dispersion and stride cycle categories the least-squares distance of the mean simulated and empirical curves has been used. The maximum fit-error per category has been scaled to 1 based on the worst fit. The resulting global error for each model configuration is the sum of errors across the 5 evaluation categories and is reflected by the height of the respective column.

tions. Starting from stimulus-free exploration we advance to chemotactic navigation and finally to adaptive odor preference experiments. Individual virtual larvae behave independently of each other as they move through the spatial arena and odorscape (38).

Free exploration. Larvae explore a stimulus-free environment, dispersing in space from their starting positions. As there is no food nor any salient odor gradient, integration of crawling and bending behaviors is sufficient. We compared free exploration in populations of 200 virtual and real larvae respectively (Video 6). Statistical evaluation showed a good agreement of simulated and empirical data with respect to spatial dispersion of larvae from their initial position (Fig. 5 A,B), total distance traveled, time fraction allocated to crawling and number of performed strides (Fig. 5 C). For a demonstration of virtual and real larvae exploring a dish see Video 5 while for a comparative assessment of their dispersion dynamics see Video 6.

Chemotaxis. Chemotaxis describes the process of exploiting an odor gradient in space to locate an attractive or avoid a repelling odor source. An olfactory sensor (olfactor) placed at the front end of the virtual body enables active sensing during body bending and allows detection of concentration changes that modulate turning behavior accordingly (see Methods). To assess chemotactic behavior in our model we reconstruct the arena and odor landscape (odorscape, see Methods) of two behavioral experiments described in (14). In the first, larvae

are placed on the left side of the arena facing to the right. An appetitive odor source is placed on the right side. The virtual larvae navigate up the odor gradient approaching the source (Fig. 6 A), reproducing the experimental observation in Fig. 1 C in (14). In the second, both the odor source and the virtual larvae are placed at the center of the arena. The larvae perform localized exploration, generating trajectories across and around the odor source. (Fig. 6 B), again replicating the observation in Fig.1D in (14). Fig. 6 E and F show the average time-varying odor concentration encountered by the virtual larvae along their trajectories, replicating the estimations from real larval tracks in (14). Two sample simulations can be seen in Video 7.

Odor preference test. We simulate the odor preference paradigm as described in the Maggot Learning Manual (42). Larvae are placed at the center of a dish containing two odor sources in opposite sides and left to freely explore. The odor concentrations are Gaussian-shaped and overlapping, resulting in an odorscape of positive and/or negative opposing gradients. After 3 minutes the final situation is evaluated. The established population-level metric used is the olfactory preference index (PI), computed for the left odor as $PI_l = \frac{N_l - N_r}{N}$ where N_l and N_r is the number of larvae on the left and right side of the dish while N is the total number of larvae.

The extend of olfactory modulation on the turning behavior is determined by the odor-specific gain G (see Materials and Methods). As this is measured in arbitrary units, we first need to define a realistic value range that correlates with
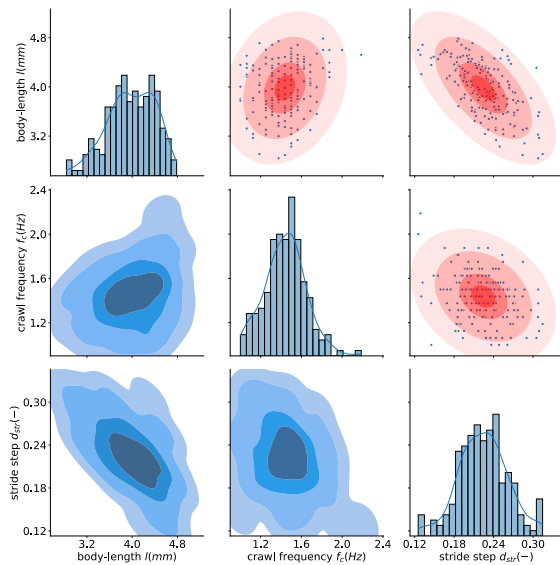
Fig. 4. Parameters of individuality: empirical and fitted distributions. Diagonal: Histogram and kernel density estimates (KDE) for body-length l, crawling frequency f_c and mean scaled displacement per stride d_{str} across a population of 200 larvae in the experimental dataset. Below: Bivariate projection of 3-dim. KDE outlined contours for each parameter pair. Above: Red ellipses represent the bivariate projections of the 3-dim. fitted Gaussian distributions at 0.5, 1, 2 and 3 standard deviations. In our model this Gaussian is used to sample a parameter set for each individual larva. The blue dots denote the empirically measured parameters.

Fig. 5. Simulation of free exploration. **A:** Dispersion of 200 larvae in experiment (left) and simulation (right) during 40 seconds. Individual tracks have been transposed to originate from the center of the arena. **B:** Median dispersion from origin. Shaded area denotes first and third quartiles. **C:** Histograms for total number of strides, time ratio allocated to crawling and pathlength. (arena dimensions = $500\text{x}500\,\text{mm}$, N = 200 larvae, experiment duration = 3 minutes)

the behaviorally measured PI. We perform a parameter-space search independently varying the gain for left and right odors and measuring the resulting PI in simulations of 30 larvae. The results for a total of 25^2 gain combinations within a suitable range of $G \in [-100, 100]$ are illustrated in Fig. 7 A. Simulation examples for one appetitive and one aversive odor are shown in Video 8.

In order to simulate larval group behavior in response to an associative learning paradigm we interface our behavioral simulation with the spiking mushroom body (MB) model introduced in (43) (Fig. 7 C). It implements a biologically realistic neural network model of the olfactory pathway according to detailed anatomical data using leaky integrate-and-fire neurons (44). The MB network undergoes associative plasticity at the synapses between the Kenyon cells and two MB output neurons as a result of concurrent stimulation with an odor and a reward signal. Both, odor and reward is simulated as spike train input to the receptor neurons and the reinforcement signalling dopaminergic neuron, respectively. The model employs two output neurons (MB$_+$, MB$_-$), representing a larger number of MB compartments associated with approach/avoidance learning respectively (45). The initially balanced firing rates between MB$_+$ and MB$_-$ are skewed after learning and encode the acquired odor valence (46, 47) here defined as

$$MB_{out} = \frac{MB_+ - MB_-}{MB_+ + MB_-} \quad \in [-1, 1].$$

We first trained the MB model via a classical conditioning experiment where, in each conditioning trial, it experiences an odor (conditioned stimulus, CS+) in combination with a

sugar reward during $5\,\text{min}$, following the standard training protocol in (42). 5 groups of 30 MB models undergo between 1 and 5 sequential conditioning trials (48). The resulting odor valence MB$_{out}$ from each MB model was converted to an odor gain G via a simple linear transformation and used to generate a virtual larva (Fig. 7 B). Each population of 30 larvae was then tested in an odor preference simulation. The larvae were placed on a dish in presence of the previously rewarded odor (CS+) and a neutral odor in opposite sides of the dish (Video 8), again following standard experimental procedures (42). To obtain robust results we replicated the experiment 100 times per population with a different random seed for a total of 600 simulations. The obtained preference indexes (PI) are illustrated in Fig. 7 D. The PI increase with increasing number of trials as well as its saturation resembles empirical observations (48). Note that the variability of the PI across the 100 simulations per condition is introduced solely by the behavioral simulation and resembles that seen across real experiments.

The current implementation only sequentially couples a trained MB model to be tested in a behavioral simulation. In the discussion we further elaborate on a possible extension featuring their closed-loop integration allowing for full behavioral simulations of both the training and the testing phase of the associative learning paradigm in a virtual environment.

Discussion

Neural pathways. The neural mechanisms enabling most of the basic behaviors have been extensively studied. Crawling occurs via fairly stereotypical repetitive strides. Head and tail segments initiate the stride concurrently with a 'visceral piston' mechanism followed by a laterally symmetric peristaltic wave traversing neighboring segments longitudinally from back to front (23) (Fig. 2 B). Segmental central pattern generators (CPGs) coupled via premotor-involving intersegmental short- and long-range connectivity motifs constitute the underlying neural circuitry (26, 49, 50). Lateral bending results from asymmetric contraction of body musculature initiated at the thoracic segments (51). Finally, feeding is
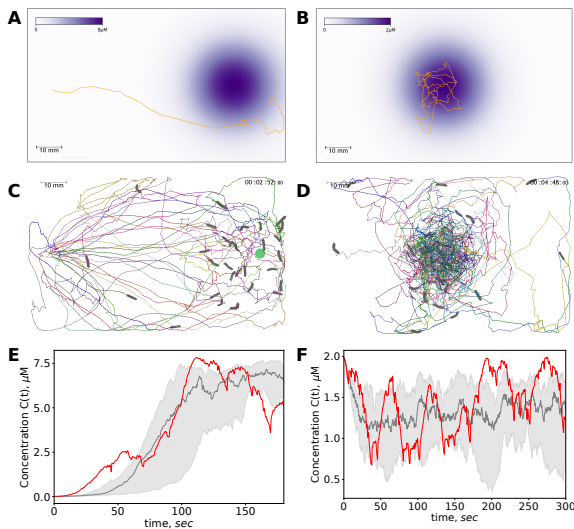
Fig. 6. Simulation of chemotaxis. **A:** Experiment 1: A single odor source of $8.9\,\mu$M peak concentration is placed on the right side of the rectangular arena creating a chemical gradient as indicated by the color scale. Larvae are placed on the left side facing to the right. Larvae are expected to navigate up the gradient approaching the source. A single larva trajectory is shown. This setup mimics the first experiment in (14). **D:** Experiment 2: A single odor source of $2.0\,\mu$M peak concentration is placed at the center of the rectangular arena. Larvae are placed in close proximity to the odor source. Larvae are expected to locally explore generating trajectories around and across the source. A single larva trajectory is shown. This setup mimics the second experiment in (14). **C,D:** The trajectories of 25 virtual larvae during the two experiments. **E,F:** The odor concentration encountered by the virtual larvae as a function of time. Red curves refer to the single larva in A and B. Gray denotes the mean and quartiles of all 25 larvae in C and D. The simulation results fit well to the experimental estimates of concentration sensing during larval chemotaxis in (14). (arena dimensions = 100×60 mm, N = 30 larvae, experiment duration = 3 and 5 minutes respectively)

generated via a network of mono- and multi-synaptic sensorimotor loops from enteric, pharyngeal and external sensory organs to motor neurons controlling mouth-hook movement, head-tilt and pharyngeal pumping (27).

All three basic behaviors are autonomously generated by local circuitry while higher brain centers modulate their activity via descending input, mediated by dopamine, serotonine, achetylocholine, octopamine and other neurotransmitters (27, 52–56). It has been suggested that the transition between exploration and exploitation (feeding) is acutely induced via dopaminergic signaling (57) while their long-term balance is regulated via hugin-mediated homeostatic neuromodulation (58). Identification of sensory pathways towards motor effector neuropiles further elucidates the role of interoception in behavioral modulation (59).

The neural mechanisms that underlie olfactory modulation of the basic locomotory behavior are also under intense investigation. Chemotactic approach and avoidance to innately valenced odors has been attributed to a predominantly innate pathway involving the antennal lobe (AL) and its direct projection to the lateral horn (LH), both in the larva (56) and in the adult fly (60–62). Modulation of learned odors strongly involves top-down control by the MB in juvenile (45, 55) and adult (31) stages. Both pathways have similar modulating effects on foraging behavior and are likely integrated in a

premotor network downstream of the AL (16, 63). In the sensorimotor loop, descending pathways involving the LH control cessation of crawling, possibly triggering sharper reorientation when navigating down-gradient, facilitating chemotaxis (64). Finally, the internal homeostatic state (e.g. starvation vs. satiation) regulates behavior via neuromodulatory transmitter release at multiple levels, including AL, LH and MB (56).

Existing computational models. Crawling and bending mechanisms have been successfully captured in previous computational models. CPG models of segmentally repeated paired excitatory and inhibitory (EI) neuronal rate units, standing for average EI population activity, can autonomously generate forward and backward crawling, possibly involving proprioceptive feedback (65, 66) and the contribution of the visceral-piston mechanism to the peristaltic cycle has been assessed in a biomechanical model (67). The hereby adopted, continuous oscillatory lateral bending process has been implemented as a pair of bilateral mutually inhibitory EI circuits (24). The idea has been elaborated in a neuromuscular model generating autonomous forward / backward crawling and turning as well as their interplay during free exploration in a 12-segment larva body by modeling segmental localized reflexes and substrate frictional forces and assuming empirically informed axial and transverse oscillatory frequencies (68). Chemotaxis has also been modeled computationally either in stochastic transition models assuming discrete behavioral states of crawling (runs) and turning (head-casts) (69) or by introducing olfactory modulatory input on the underlying locomotory circuit (24, 68).

We propose here that modeling locomotion as coupled intermittent oscillatory processes (Fig. 1 B) is adequate for generating realistic larva kinematics as these are captured via larva tracking. This locomotory model efficiently summarizes the underlying CPG activities into global linear and angular velocity oscillations (23, 26) and accurately reproduced a number of experimental observations. It is therefore well suited for the bottom layer of the proposed behavioral architecture (Fig. 1). The latter though allows for a modular approach, where the level of abstraction can be chosen independently for each individual module as exemplified by combining this locomotory model with a spiking neural network that captures plasticity in a central brain neuropile (MB) at the top layer. It follows that the proposed control architecture affords any substitution of the currently selected oscillator modules by one of the aforementioned more-detailed neuromechanical models if a higher degree of biological realism is pursued.

Crawl-bend interference. Crawling includes mouth hook motion. Specifically, the first phase of a crawling stride consists of concurrent forward motion of head and tail segments, aided by a 'visceral pistoning' mechanism that generates forward displacement of the gut. Subsequently, the mouth hooks anchor the head to the substrate so that the second phase of peristaltic motion can drag all other segments forward as well, completing the stride (23). Crawling and bending partially recruit the same effector neural circuitry and body musculature at least at the level of the thorax. Peristaltic motion during crawling includes sequential symmetric bilateral contraction of all segments while bending occurs due to asymmetric unilateral contraction of the thoracic segments. This partial effector overlap could result in interference between the two processes. Indeed here we report an increase of orientation

velocity during a specific phase of the stride cycle (Fig. 2 H), synchronous to an increase of head forward velocity (Fig. 2 D). The latter coincides with the stride phase when the head is not anchored to the substrate and therefore free to move laterally. When applying attenuation of lateral bending outside a phase interval $[\frac{\pi}{2}, \pi]$ of the stride cycle we managed to accurately reproduce the empirical relation (Fig. 8 A).

A reasonable hypothesis would then be that the asymmetric thoracic contraction generating lateral bending is only possible while the head is not anchored to the substrate therefore during a specific phase interval of the stride cycle. We suggest that crawling phasically interferes with lateral bending because of these bodily constraints. A consequence of the proposed hypothesis is that the amplitude of turns generated during crawl-pauses is larger in comparison to those generated during crawling because during pauses the crawling interference to lateral bending is lifted. We postulate that it is exactly this phenomenon that dominates the description of larva exploration as a Levy-walk with non-overlapping straight runs and reorientation events, where the minor orientation changes taking place during crawling are neglected (10, 70). We note that in the implemented model the turner neural oscillation is not inhibited at all during crawling strides as we consider this merely a bodily interference, although the resulting torque might eventually not be applied to the body depending on the stride-cycle phase.

Behavioral intermittency. Larval locomotion is intermittent meaning that crawling runs are transiently intermitted by brief pauses. The spatial dynamics of these alternating states have been studied in the context of motion ecology. During free exploration, power-law distributed runs, in line with Levy-walk theoretical models (10, 70) and diffusion-like kinematics have been reported (32) while the speed-curvature power-law relationship has been disputed (71–74). Regarding the temporal dynamics of intermittency, the duration distributions of activity and inactivity bouts captured via larva-tracking recordings have been studied. More specifically, the duration of inactivity bouts has been reported to follow a power-law while that of activity bouts a log-normal distribution (11), partly in line with findings in adult-fly studies (12, 13). In all these studies micro-movements like feeding and lateral bending could not be detected due to technical constraints. Thus, the reported inactivity and activity bouts can be regarded as crawl-pauses and crawl-runs respectively, the duration of the latter being equivalent to our discretized stridechain-length metric.

Our analysis reveals that both distributions are approximated best by log-normal distributions (Fig.8,B). The log-normal stridechain distribution is in line with previous findings, while the log-normal pause distribution diverges from a previously reported power-law (11). This might be attributed to the short, 3-minute duration of the recordings in the present study contrary to the long, up to one hour recordings used in the latter.

Computational models of behavioral intermittency are scarce. A recent study presented a simple binary-neuron model exhibiting state transitions between power-law and non
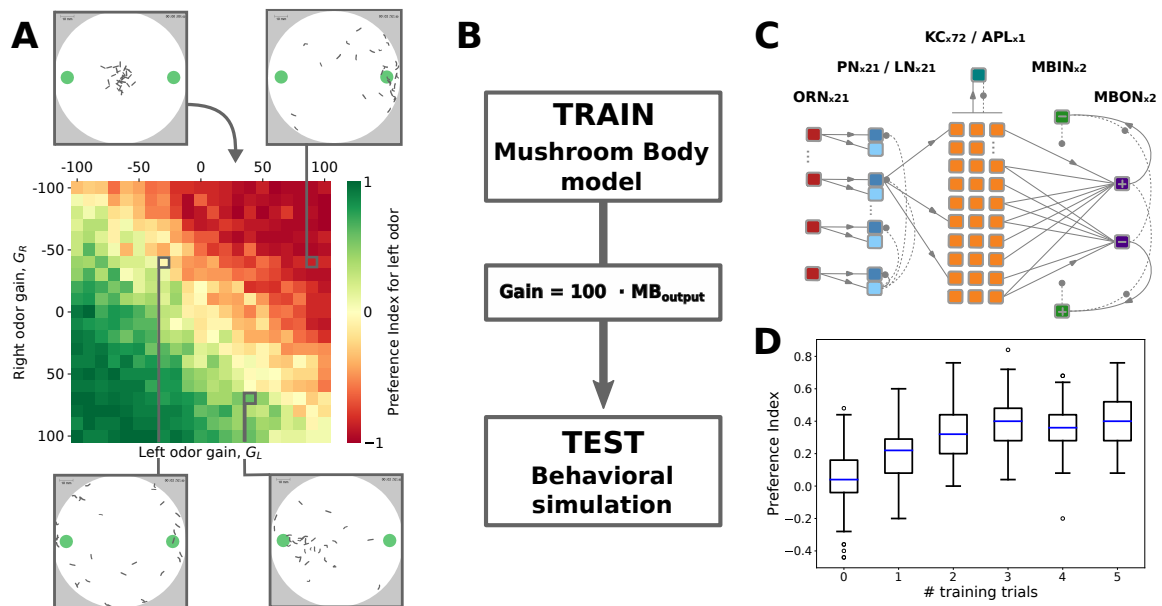
Fig. 7. Simulation of innate odor preference. **A:** A total of 25^2 simulations are shown with the resulting Preference Index for different gains of the left and right odor. On the top left the initial state is shown with the larvae randomly generated at the center of the dish. The final state of three additional simulations is depicted. **B:** The pipeline used for coupling the Mushroom Body (MB) model with the behavioral simulation. First a MB model is trained via a classical conditioning experiment where olfactory input is combined with reward. The resulting odor valence MB_{out} is then converted to odor gain G via a simple linear transformation and used to generate a virtual larva. Finally the odor preference of a virtual larva population is evaluated in a behavioral simulation. **C:** The spiking neural network comprising the MB model. The number of neurons comprising each layer is indicated. **D:** The resulting PIs for 100 simulations per number of training trials. In each of the 100 simulations per condition a population of 30 virtual larvae was generated and evaluated using a different random seed, always bearing the exact same 30 odor gains derived from the respective group of 30 trained MB models (arena dimensions = 100x100 mm, N = 30 larvae, experiment duration = 3 minutes)

Sakagiannis *et al.*

power-law regimes via self-limiting neuronal avalanches and proposed a plausible underlying mechanism that explains initiation/cessation of crawling (11). In the current model we remain agnostic to the intermittency generative process and simply sample pause duration and stidechain length from the empirically fitted distributions. Further neuroanatomical evidence of the underlying circuits is needed to elucidate the neural mechanism of intermittent crawling.

Architecture extensions. The here adopted subsumption architecture paradigm is mainly used in behavior-based robotics (75) and is supported as a theoretical framework for the description of the nervous system of living organisms (39). According to this, neural control of behavior consists of nested sensorimotor loops where more stereotyped reflexive behaviors are autonomously generated by localized neural circuitry at a faster timescale without recruiting more centralized resources. Once there is need for more extensive integration, e.g. in order to react suitably to unexpected sensory stimulation, slower multisynaptic loops enable higher centers to modulate local circuits achieving more coordinated global behavioral control. In other words, the nervous system is considered a multilayered control architecture within which higher layers subsume their subordinate layers into comprehensive behavioral modes. The main idea governing this paradigm is that there are only limited degrees of freedom in which higher layers can influence the lower ones e.g. by initiation/cessation or acceleration/deceleration of their autonomous function.

The modularity of the proposed architecture facilitates further extensions capturing novel aspects of the body, the nervous system, the metabolism or the environment as shown in Fig.1, allowing for more refined or more complex behavioral patterns. We here summarize some plausible extensions.

Feeding. Despite their central role in foraging, feeding mechanisms have not been modeled computationally nor has their integration with the exploratory circuitry in the context of substrate exploitation. Feeding behavior also consists of repetitive stereotypical movements of the head, mouth hook and ingestive muscles (27). Every cycle is an intertwined sensorimotor loop under instantaneous feedback from the environment and slower regulation by higher self-regulatory centers. In adult flies successive feeding movements are organized in intermittent feeding bouts, interspersed by locomotory or idle periods (76). Although detection of larval feeding motions in freeforaging conditions is technically difficult due to their small amplitude, it is reasonable to assume these are also structured in intermittent bouts. The frequency of this repetitive motion has been reported to vary at least from 1 to 2.5 Hz. Therefore feeding behavior can be implemented as a third oscillator (feeder). Top-down modulation affecting initiation/cessation and oscillation frequency can be assumed, similar to crawling while sensory feedback can be implemented as recurrent modulation depending on successful ingestion.

Considering oscillator coupling, crawling and feeding cycles partially compete for control of the same effectors as they both recruit the head and mouth muscles. It has been reported that the number of mouth hook motions over a given duration of foraging does not differ between rover and sitter larva phenotypes and is not correlated to the amount of food ingested, although rovers crawl more and feed less than sitters (30). Therefore an individual mouth hook motion can equally be part of either a crawling or a feeding cycle. The conclusion drawn is that individual crawling and feeding cycles are competing mutually exclusive behavioral motifs meaning crawling and feeding bouts can alternate but do not overlap. Conversely there is no empirical evidence on the potential coupling of feeding and bending motions. Integration of the feeder in the current locomotory model is illustrated in Fig.1 B.

Multimodal sensory feedback. Sensory feedback from the environment can be extended to other modalities beyond olfaction. Mechanosensation can be implemented via additional touch sensors around the body contour. This will allow detection and behavioral modulation by conspecific contact (collisions) (77) and external mechanosensory stimulation driving startle/evasion (hunch/bend) (78) or navigation along wind gradients (anemotaxis) (79). Accordingly, temperature and light sensors could allow thermotaxis (80–82) and phototaxis (83). The respective environmental sensory gradients can be Gaussian as in the case of odorscapes or linear along a certain arena axis. Behavioral modulation can be either summed up across modalities or separately applied to the bending, crawling and feeding effectors. Integration of additional sensory modalities is illustrated in Fig.1 A.

Olfactory learning in closed loop behavioral simulations. We have shown an open-loop simulation of the classical conditioning paradigm (Fig. 7 B-D) reproducing a basic experimental result in the fruit fly larva (36). This modeling approach can be extended in multiple ways. First, the larva demonstrates a number of interesting learning abilities/features that require synaptic and circuit mechanisms such as differential conditioning (15, 16, 36), extinction learning (84), and relief learning (35, 45, 48, 85). Interfacing neural network simulations of these mechanisms with our behavioral model allows to directly compare virtual and empirical behavioral experiments, both for the typical group assays and for individual animals. Second, while information about odor concentration is provided via olfactory sensing, direct input from a feeder module could provide the reward stimulus that activates the dopaminergic pahtway required for synaptic plasticity the mushroom body (45, 86). This would further allow realistic foraging scenarios with food depletion and competition. Closing the loop from active sensing to associative memory formation and behavioral control requires to synchronize a (spiking) neural network at the adaptive layer with the sensory (reactive layer) and locomotor modules (basic layer). This will enable the simulation of virtual larva experiencing spatial and temporal dynamics in a virtual environment or on a robotic platform (87). Such model approaches will allow to test model hypotheses on sensorymotor integration and to infer predictions for experimental interventions such as optogenetic stimulation (45) or genetic manipulations (88–91).

Materials and Methods

Dataset description. The larva-tracking dataset was obtained by M.Schleyer and J. Thoener at the Leipzig Institute of Neurobiology (https://doi.org/10.12751/g-node.5e1ifd). It consists of 200 third-instar larvae tracked at a framerate of 16 Hz for 3 minutes while exploring a non-nutritious substrate. 12 points are detected along the longitudinal axis of each larva. Detected collisions have been excluded and the data has been filtered with a first-order butterworth low-pass filter with a

cutoff frequency of 2 Hz in order to decrease tracking-related noise but retain the behaviorally relevant crawling frequency of ≈ 1.5 Hz. The effect of inadequate and excessive filtering is illustrated in Video 9.

All data processing and all simulations were performed using the *Larvaworld* behavioral analysis and simulation platform programmed in python. We made *Larvaworld* freely available at https://github.com/nawrotlab/larvaworld. In *Larvaworld*, simulated and empirical data are treated indistinguishably, meaning that the exact same analysis pipeline and behavioral metrics are applied in both.

	Parameter		Symbol	Value	Unit
PHYSICS	linear damping		z_l	1	-
	angular damping		z_a	2.5	-
BODY	length		l	4.325*	mm
	front/rear ratio		$r_{2\text{-seg}}$	5/6	
	spring constant		k	0.02	
	bend correction coefficient		b_c	1.4	-
TURNER	initial activation		A_t	20	-
	activation range		A_{range}	[10,40]	-
	noise$_{input}$		n_A	0.5	-
	torque coefficient		c_t	0.4	-
	noise$_{output}$		n_t	0.15	-
CRAWLER	frequency		f_c	1.428*	Hz
	stride displacement $_{scaled}$		d_{str}	0.225*	-
	noise$_{output}$		n_c	0.1	-
INTERFERENCE	interference-free phase interval		P_{range}	[π/2, π]	-
	attenuation ratio		c_{at}	0.1	-
INTERMITTENCY	pause duration	range	p_{range}	[0.1 - 2]	sec
		distribution	lognormal(m_p, s_p)	-1.2, 0.7	-
	stridechain length	range	s_{range}	[1 - 41]	# strides
		distribution	lognormal(m_c, s_c)	1.4, 1	-

Table 1. Model parameters. Red values: hard coded; green values: fitted from empirical data; blue values: collectively sampled from an empirically fitted multivariate Gaussian distribution (Fig. 4) for which mean values are shown here.

Model definition and calibration.

Crawler oscillator. Crawling behavior is modeled as an oscillatory process. Each oscillation generates a cycle of forward velocity v increase-decrease resulting in displacement of the larva along the axis of its front-segment, simplistically modeling the result of exactly one peristaltic stride. An analytically tractable curve is fitted to the average empirical velocity curve measured during strides (Fig. 2 D):

$$v^* = d_{str} \cdot f_c \cdot (0.6 \cos(\Phi - \Phi_{max}) + 1) \qquad [1]$$

where d_{str} is the displacement per stride, scaled to the larval body-length l, f_c being the crawling frequency and Φ the instantaneous phase of the oscillation iterating from 0 to 2π during an oscillatory cycle. $\Phi_{max} = \pi$ is the phase where the maximum velocity occurs. The equation ensures a constant d_{str} despite f_c changes. Gaussian noise is applied so that $v = N(v^*, n_C \cdot \overline{v}^*)$. n_C is selected so that the variance across strides fits the empirical observations (Fig. 8 A,top). The 3 parameters (l, d_{str} and f_c) defining the velocity curve are sampled from empirical data, retaining their paired correlations (Fig. 4). The oscillatory process can only be halted/initiated at the start of a cycle meaning that once a stride is initiated it will be completed.

To calibrate Φ_{max} in the stride cycle, a sliding window analysis was performed (data not shown). d_{str} was computed for all possible Φ_{max} in windows of the reference stride duration

f_c^{-1}, selecting the one displaying the minimum variance for displacement, which results in more stereotypical strides.

Turner oscillator. The lateral oscillator model described in (24) assumes an underlying oscillatory process driving alternating bending to the left and right side. The oscillator consists of two mutually inhibiting components (L vs R) that quickly settle in antiphase, while adaptation ensures that periodic transitions occur. The system is driven by external activation A_t. The baseline activation $\overline{A}_t = 20$ was held constant as in the original implementation (24), resulting in an average 0.3 Hz oscillation frequency. Perturbations of this external drive cause transient changes in both amplitude and frequency, up to transient loss of oscillation. This feature is exploited during olfactory modulation (see Olfactory sensor). Turner activity is the instantaneous difference in the firing rates $\Delta f = (f_L - f_R)$ and is scaled by a coefficient c_t to generate the oscillating torque $T = c_t \cdot \Delta f$. This is applied to the body which is modeled as a torsional spring of restorative spring constant k, causing lateral bending θ_b. The angular velocity is attenuated by angular damping ratio z. It holds that :

$$\ddot{\theta}_{or} = T - z \cdot \dot{\theta}_{or} - k \cdot \theta_b \qquad [2]$$

Gaussian noise is applied to both the input activation $A_t = N(A_t^*, n_A \cdot \overline{A}_t^*)$ and the output turner activity $\Delta f = N(\Delta f^*, n_T \cdot \overline{\Delta f}^*)$.

In its original implementation the model deliberately neglects two aspects of the real turning behavior of the larva. First there is no distinction between bending $\dot{\theta}_b$ and orientation $\dot{\theta}_{or}$ angular velocities. Second there is no correction of the bending angle θ_b due to forward motion. We tackle the first via the bisegmental body so that $\dot{\theta}_b$ between the front and rear vector is distinct from the front vector's $\dot{\theta}_{or}$. Regarding the second, we introduce a simple linear bending-angle correction as the rear vector is aligned to the front vector's orientation during forward motion, according to the equation :

$$\theta_b' = \begin{cases} \theta_b \cdot (1 - d/d_{max}), & \text{if } 0 \le d < d_{max} \\ 0, & \text{if } d \ge d_{max} \end{cases}$$

$$d_{max} = \frac{l}{2b_c} \qquad [3]$$

where d is the linear displacement during a timestep, b and b' are the original and corrected bending angles, l is the body length and b_c is the bend correction coefficient which has been fitted from the empirical data. In Fig. 8 E the resulting $\Delta\theta_b$ correction during individual strides is plotted in relation to θ_b at the beginning of the stride. The fitted linear curves show close matching of empirical and simulated findings.

To calibrate the turner oscillator we need to determine 5 parameters, namely the torque coefficient c_t, the angular damping ratio z, the restorative spring constant k and the noise ratios n_A, n_T. To this end we analyze the angular dynamics of a population of 200 larvae and compute the distribution of 3 empirical metrics, namely bending angle $\hat{\theta}_b$, bend and orientation angular velocities $\hat{\dot{\theta}}_b$, $\hat{\dot{\theta}}_{or}$. In order to find optimal parameter combinations providing the best fit of simulated and empirical distributions a two-step parameter space search is performed :

- We first evaluate the function of the isolated turner oscillator, in autonomous runs. We constrain the parameter
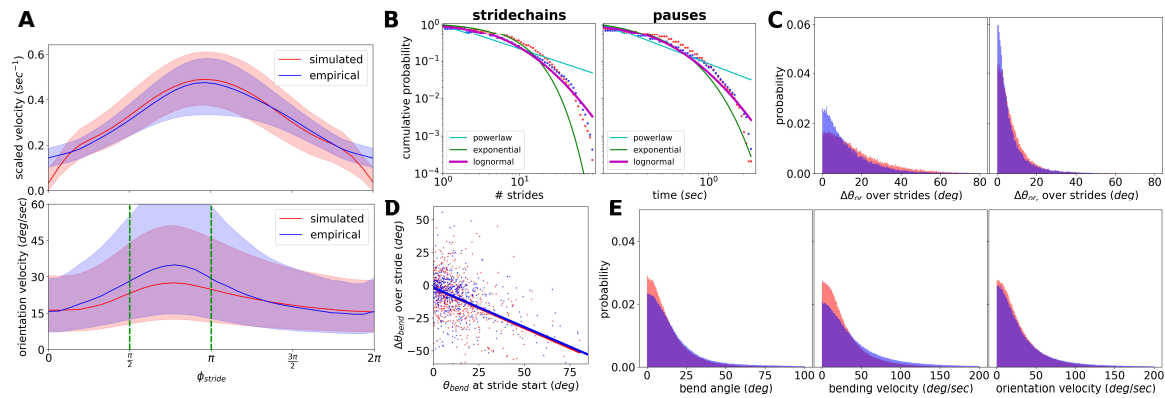
Fig. 8. Model calibration A variety of behavioral metrics are used to optimize the locomotory model's fit to empirical data. Here the final calibrated model is assessed in a 3-minute simulation of free exploration using a population of 200 virtual larvae. The simulated data (red) are compared to empirical data (blue) from a 3-minute free exploration experiment of 200 real larvae. Data are pooled across individuals in each population. **A:** Crawl-bend interference. Mean scaled forward velocity (top) and orientation velocity $\dot{\theta}_{or}$ (bottom) during strides. All detected strides across the entire populations have been interpolated to a common stride cycle from 0 to 2π. Shaded areas denote first and third quartiles. Dashed green lines show the cycle interval boundaries of unconstrained angular motion. Outside the boundaries angular motion is attenuated due to crawler interference. **B:** Bout distribution. Dots describe the cumulative probability density over logarithmic bins for the length of stridechains and the duration of crawl-pauses. Lines indicate the distribution with the lowest Kolmogorov-Smirnov distance among the best fitting power-law, exponential and log-normal distributions. Stridechain length and pause duration are best approximated by log-normal distributions. **C:** Distribution of front and rear segment orientation change $\Delta\theta_{or_f}$, $\Delta\theta_{or_r}$ during individual strides. **D:** Bend correction due to forward motion : Scatterplot of the bending angle change $\Delta\theta_b$ over individual strides relevant to θ_b at the beginning of the stride. Lines denote the fitted linear curve for real and virtual larvae. **E:** Histograms of angular metrics using the optimal set of turner-relevant parameters as determined by the calibration process. The three metrics shown are bending angle θ_b, bending and orientation angular velocities $\dot{\theta}_b$, $\dot{\theta}_{or}$.

space by selecting parameter combinations that fit the range of the empirical metrics (first and fourth quartiles within the empirical lower and higher halfs) without θ_b overshooting the maximum allowed π (data not shown).

- We then evaluate the turner-crawler integration in free exploration simulations of virtual larva populations optimizing for best fit to the empirical distributions. The results of the optimal parameter set is shown in Fig. 8 D.

Crawler-turner coupling. In order to reproduce bending behavior during strides we define a phase interval P_{range} during which the turner is free to exhibit its full effect on the body. Outside this interval, the torque T generated by the turner is attenuated by a scalar coefficient c_{at} before being applied to the body, according to the equation :

$$T = \begin{cases} T, & \text{if } p_c \text{ in } P_{range} \\ c_{at} \cdot T, & \text{otherwise} \end{cases}$$

where p_c is the instantaneous phase of the crawling stride. We calibrate the two parameters P_{range} and c_{at} so that the average orientation angular velocity $\dot{\theta}_{or}$ during the stride cycle fits the empirically measured. The results for $P_{range} = [\frac{\pi}{2}, \pi]$ and $c_{at} = 0.1$ are shown in Fig. 8 A,bottom. Furthermore the orientation change during individual strides for the front $\Delta\theta_{or}$ and rear $\Delta\theta_{or_r}$ segments fits well to empirical measurements (Fig. 8 C).

Olfactory sensor. Olfaction is introduced in the second layer of the control architecture allowing chemotactic behavior. The olfactory sensor is located at the front end of the virtual larva therefore any reorientation and/or displacement influences sensory input. As in (14) we assume that olfactory perception A_o relates to changes in odor concentration C according to the

Weber-Fechner law, meaning that $\Delta A_o \sim \ln \Delta C$. We further add a decay term that slowly resets A_o back to 0. The rate of change is given by the equation :

$$\dot{A}_o = -A_o \cdot c_o + \sum_i G_i \cdot \frac{\dot{C}_i}{C_i} \qquad \text{with} \quad -1 \leq A_o \leq 1 \quad [4]$$

where $c_o = 1$ is the olfactory decay coefficient, G_i is the gain for odor i and C_i the respective odor concentration. Perceived olfactory stimulation A_o modulates the turner activation A_t from its baseline value $\overline{A}_t = 20$ within a suitable range $A_{range} = [A_t{}^{min}, A_t{}^{max}] = [10, 40]$:

$$A_t = \bar{A}_t + A_o \left(A_t{}^{lim} - \bar{A}_t \right) \qquad [5]$$

$$\text{where} \quad A_t{}^{lim} = \begin{cases} A_t{}^{max} & \text{if } A_o \geq 0 \\ A_t{}^{min} & \text{if } A_o \leq 0 \end{cases} \qquad [6]$$

Parameter definition.

Segmentation and angular metrics. To specify the body segmentation providing the most suitable contact/rotation point for the definition of the correlated empirical bending $\dot{\theta}_b$ and orientation $\dot{\theta}_{or}$ angular velocities we analyse their relationship in a subset of 40 larvae. Tracking of 12 midline points allows computation of the absolute orientation of 11 body-segments and the respective 10 angles θ_{1-10} between successive body segments (Fig. 2 F). We define $\dot{\theta}_{or}$ as the head-segment orientation because this defines the movement orientation of the animal. We ask how $\dot{\theta}_{or}$ results from the bending of the body as this is captured by the 10 angular velocities $\dot{\theta}_{1-10}$. The regression analysis depicted in Fig. 9 B shows as expected that $\dot{\theta}_{or}$ depends primarily on the front angular velocities
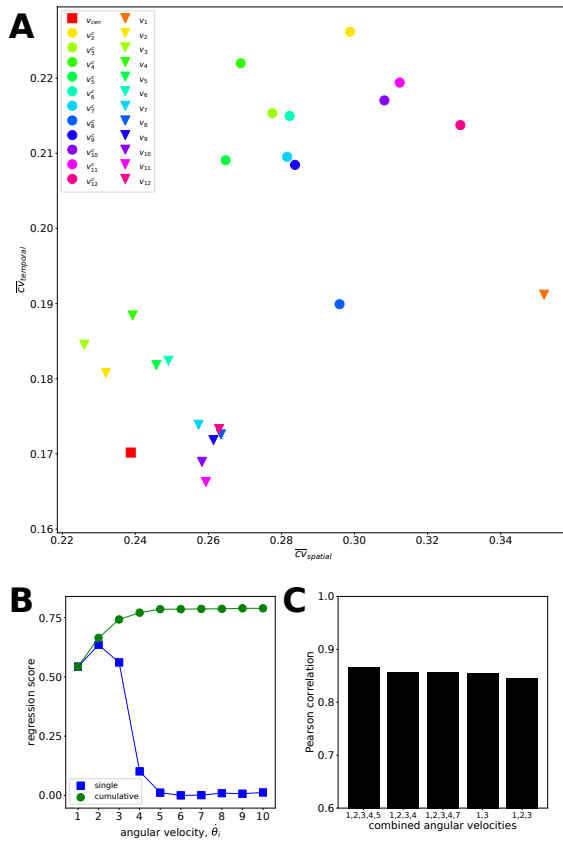
Fig. 9. Segmentation and velocity definition. A: Forward velocity definition. 24 candidate velocity metrics are compared for use in stride annotation of 3-minute tracks of a population of 20 larvae. For each candidate the mean coefficient of variation of temporal duration $\overline{cv}_t$ and spatial displacement $\overline{cv}_s$ of the annotated strides is shown. The centroid velocity v_{cen} provides the most temporally and spatially stereotypical strides, therefore it is selected as the reference forward velocity for stride annotation and model fitting. v_{cen} : centroid velocity, $v_1 - v_{12}$: 1^{st}-12^{th} point's velocity, $v_2^c - v_{12}^c$: 2^{nd}-12^{th} point's component velocity parallel to the front segment's orientation vector. **B:** Regression analysis of individual and cumulative angular velocities $\dot{\theta}_{i=1-10}$ to orientation angular velocity $\dot{\theta}_{or}$. When considered individually, $\dot{\theta}_2$ best predicts reorientation with the $\dot{\theta}_1$ and $\dot{\theta}_3$ following. When considered cumulatively the anterior 5 $\dot{\theta}_i$ allow optimal prediction of reorientation velocity. **C:** Correlation analysis of the sum of all possible $\dot{\theta}_i$ combinations to $\dot{\hat{\theta}}_{or}$. The sum $\sum_{i=1}^{5} \dot{\theta}_i$ shows the highest correlation therefore we define $\hat{\theta}_b = \sum_{i=1}^{5} \theta_i$ as shown in A. For illustration purposes only the 5 highest correlations are shown.

while this dependence decays as we move towards the rear segments, in line with previous studies (41). Timeshift analysis also shows that the front 3 angles change concurrently while angles further down the midline are increasingly lagging behind (data not shown). The correlation analysis depicted in Fig. 9 C shows that the sum of the front 5 angular velocities best correlates to $\dot{\hat{\theta}}_{or}$. In other words the cumulative body bend of the front 5 segments best predicts head reorientation. Therefore we define the reorientation-relevant bending angle as $\hat{\theta}_b = \sum_{i=1}^{5} \theta_i$ (Fig. 2 F). The remaining 5 angles between

the rear body-segments can safely be neglected as they do not contribute to reorientation. This analysis results in a segmentation of the body in a front and a rear segments of length ratio 5:6. The segmentation process is demonstrated in Video 1.

Forward velocity. To define forward velocity we need to choose which midline-point is most suitable to track and which velocity metric to use for defining the start and end of a stride. To this end we perform stride annotation of 3-minute tracks of a population of 20 larvae using each of 24 candidate instantaneous velocity metrics, namely the velocities of the 12 points, the component velocities of the rearest 11 points parallel to their front segment's absolute orientation and finally the centroid velocity. To compare the candidate metrics we compute the spatial cv_s and temporal cv_t coefficient of variation of the annotated strides for each larva to assess how variant their time duration and displacement is. We finally compute the mean $\overline{cv}_s$ and $\overline{cv}_t$ across individuals. In Fig. 9 A the spatiotemporal stride variance is shown for each candidate metric. We choose the metric that provides the minimal spatial and temporal stride variance, assuming that strides of an individual larva are more or less stereotypical in both duration and displacement (23). Our study reveals that the centroid velocity is the most suitable metric for stride annotation. All spatial metrics are therefore computed via this point's displacement.

The instantaneous body-length of an individual larva fluctuates due to subsequent stretching and contraction during crawling. Its histogram is well fitted by a Gaussian distribution (data not shown). Therefore individual larva length l was computed as the median of the midline length across time (total length of the line connecting all 12 midline points). All spatial parameters, including displacement and velocity, are scaled to this body-length.

To analyze the temporal dynamics of crawling we perform spectogram analysis of the linear velocity. The dominant frequency f_c across a plausible range of 0.75 to 2.5 Hz is defined as the crawling frequency. We use the inverse of this frequency f_c^{-1} as a reference for the expected stride duration during stride annotation.

Epoch annotation. Strides are annotated using the forward velocity v timeseries, under a number of constraints (Fig. 2 C):

- Each stride epoch is contained between two v local minima.

- The v local maxima contained in the epoch needs to be higher than the mean.

- The duration of the epoch needs to be between 0.6 and 1.6 times the reference stride duration $t_{str} = f_c^{-1}$, where f_c the crawling frequency. This allows individual strides to temporally vary without overlapping so that adjacent strides can be concatenated in stridechains.

Stridechains are defined as uninterrupted sequences of successive strides (Fig. 2 C). Stridechain length equals the number of concatenated strides and is a discrete metric equivalent to crawl-run duration. Pauses are defined as epochs containing no strides. After stride annotation the resulting displacement due to each individual stride is computed for each larva and divided by the larva's body-length (Fig.2 A-C). The individual

distributions are well fitted by Gaussians (data not shown). Therefore stride displacement d_{str} is defined as the average scaled displacement per stride for each larva.

Turn epochs are contained between pairs of successive sign changes of orientation angular velocity $\dot{\theta}_{or}$. For each epoch the turning angle θ_{turn} is defined as the absolute total change of orientation angle $\Delta\theta_{or}$ (Fig. 2 G).

Supplementary videos. Here we provide the supplementary videos cited in the text. Videos can be played by clicking on the image frame.
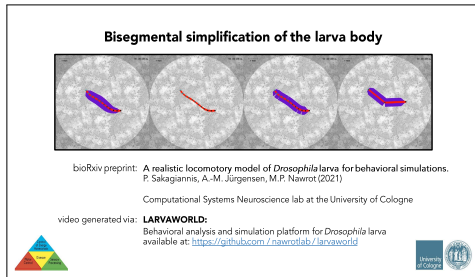
Video 1. Bisegmental larva-body simplification The first video shows the original larva body as recorded by the tracker. 12 points are tracked along its longitudinal axis defining 11 segments while 22 points constitute the body contour. In the second video the body contour is dropped. In the third video an artificial rectangular contour is added for each body segment. In the last video the body-midline is segmented into 2 segments. The absolute head orientation angle $\hat{\theta}_{or}$ is preserved while the single bending angle between the 2 segmentw is defined as $\hat{\theta}_b = \sum_{i=1}^{5} \theta_i$
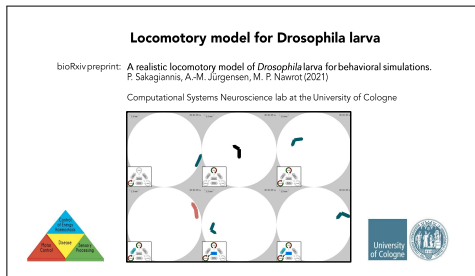
Video 2. Locomotory model for Drosophila larva The function of the locomotory model in Fig. 1 B is illustrated by gradually integrating its 4 modules (crawler, turner, oscillator-coupling, crawling-intermittency). In each of the 6 videos the implemented modules are shown in the inset.
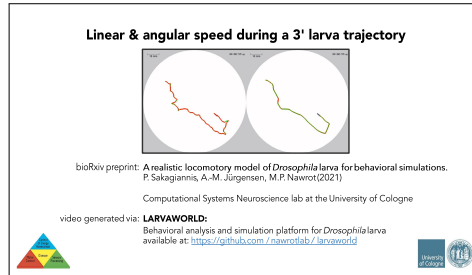
Video 3. Full-length larva trajectory The trajectory depicted in Fig. 2 is shown in its full length. Trajectory color tracks instantaneous forward v and angular $\dot{\theta}_{or}$ velocities while tracking a midline point or the head respectively. Color code ranges from red (0 velocity) to green (maximum velocity).
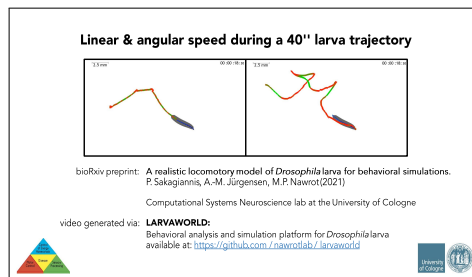
Video 4. Larva trajectory slice The trajectory slice depicted in the inset of Fig. 2 is shown. Trajectory color tracks instantaneous forward v and angular $\dot{\theta}_{or}$ velocities while tracking a midline point or the head respectively. Color code ranges from red (0 velocity) to green (maximum velocity).
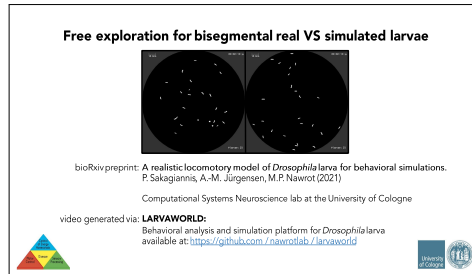
Video 5. Free exploration simulation A population of 25 real (left) or virtual (right) larvae is placed on a dish and left to freely explore. The body of the real larvae has been simplified into 2 segments as shown in Video 2 (see Model, Bisegmental body).

1. T Ohyama, et al., High-throughput analysis of stimulus-evoked behaviors in drosophila larva reveals multiple modality-specific escape strategies. PLOS ONE **8**, 1–21 (2013).
2. B Risse, D Berh, N Otto, C Klämbt, X Jiang, FIMTrack: An open source tracking and locomotion analysis software for small animals. PLoS Comput. Biol. **13**, 1–15 (2017).
3. I Schumann, T Triphan, The pedtracker: An automatic staging approach for drosophila melanogaster larvae. Front Behav Neurosci. **14** (2020).
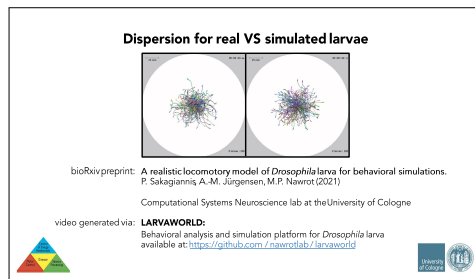
Video 6. Dispersion simulation The temporal course of dispersion from origin is shown for a population of 200 real (left) and virtual (right) larvae (Fig. 5 A). Trajectories have been transposed to initiate from the same origin at the center of the screen.
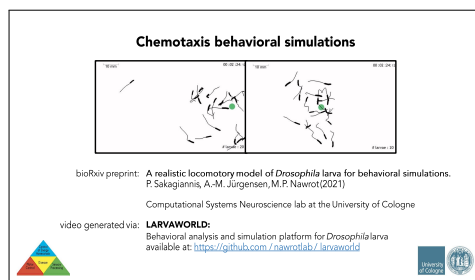
Video 7. Chemotaxis simulation The 2 simulations depicted in Fig. 6 are shown. In the first, larvae are placed on the right side of the arena and are expected to approach an odor source located on the left side. In the second, larvae are placed in the proximity of an odor source located at the center of the arena and are expected to perform localized exploration, generating trajectories across and around the odor source.
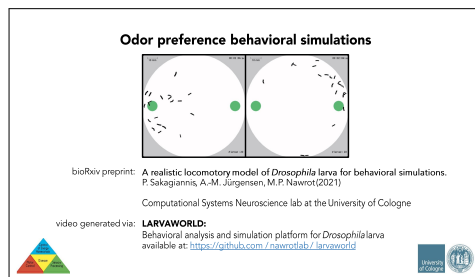
Video 8. Odor preference simulation Two sample odor preference simulations are shown. Either an appettitive or an aversive odor source is placed on the left side of the dish while a non-valenced odor source is placed on the right. Larvae are left to explore and potentially navigate along the respective salient concentration gradient.

Video 9. Filter selection The effect of inadequate or excessive filtering of the empirical larva recordings is illustrated. The left video shows the jittery original recording while the effect of lowpass filtering at cuttoff frequencies of 4 Hz, 2 Hz and 0.5 Hz is shown on the rest. Selection of an intermediate 2 Hz cutoff frequency eliminates the unrealistic jitter while preserving the behaviorally relevant crawling frequency.

melanogaster larvae. Nat. Commun. **4**, 1788–1789 (2013).

10. DW Sims, NE Humphries, N Hu, V Medan, J Berni, Optimal searching behaviour generated intrinsically by the central pattern generator for locomotion. Elife **8:e50316** (2019).

11. P Sakagiannis, M Aguilera, MP Nawrot, A plausible mechanism for drosophila larva intermittent behavior in Biomimetic and Biohybrid Systems. (Springer International Publishing, Cham), pp. 288–299 (2020).

12. T Ueno, N Masuda, S Kume, K Kume, Dopamine Modulates the Rest Period Length without Perturbation of Its Power Law Distribution in Drosophila melanogaster. PLoS One **7** (2012).

13. AM Reynolds, Evidence for a pervasive 'idling-mode' activity template in flying and pedestrian insects. R. Soc. Open Sci. **2** (2015).

14. A Gomez-Marin, GJ Stephens, M Louis, Active sampling and decision making in Drosophila chemotaxis. Nat. Commun. **2**, 410–441 (2011).

15. M Schleyer, et al., A behavior-based circuit model of how outcome expectations organize learned behavior in larval Drosophila. Learn. Mem. **18**, 639–653 (2011).

16. M Schleyer, et al., The impact of odor-reward memory on chemotaxis in larval Drosophila. Learn. Mem. **22**, 267–277 (2015).

17. B Gerber, RF Stocker, The drosophila larva as a model for studying chemosensation and chemosensory learning: A review. Chem. Senses **32**, 65–89 (2007).

18. S Diegelmann, B Klagges, B Michels, M Schleyer, B Gerber, Maggot learning and synapsin function. J. Exp. Biol. **216**, 939–951 (2013).

19. A Widmann, K Eichler, M Selcho, AS Thum, D Pauls, Odor-taste learning in Drosophila larvae. J. Insect Physiol. **106**, 47–54 (2018).

20. Q Wu, et al., Developmental control of foraging and social behavior by the Drosophila neuropeptide Y-like system. Neuron **39**, 147–161 (2003).

21. G Denisov, T Ohyama, T Jovanic, M Zlatic, Model-based detection and analysis of animal behaviors using signals extracted by automated tracking. Proc. Int. Conf. on Bio-inspired Syst. Signal Process., 175–181 (2013).

22. D Karagyozov, M Mihovilovic Skanata, A Lesar, M Gershow, Recording neural activity in unrestrained animals with three-dimensional tracking two-photon microscopy. Cell Reports **25**, 1371–1383.e10 (2018).

23. ES Heckscher, SR Lockery, CQ Doe, Characterization of Drosophila Larval Crawling at the Level of Organism, Segment, and Somatic Body Wall Musculature. J. Neurosci. **32**, 12460–12471 (2012).

24. A Wystrach, K Lagogiannis, B Webb, Continuous lateral oscillations as a core mechanism for taxis in Drosophila larvae. Elife **5** (2016).

25. G Ruiz-Dubreuil, B Burnet, K Connolly, P Furness, Larval foraging behaviour and competition in Drosophila melanogaster. Hered. (Edinb). **76**, 55–64 (1996).

26. C Mantziaris, T Bockemühl, A Büschges, Central Pattern Generating Networks in Insect Locomotion. Dev. Neurobiol. **00**, 1–15 (2020).

27. A Miroschnikow, et al., Convergence of monosynaptic and polysynaptic sensory paths onto common motor outputs in a Drosophila feeding connectome. Elife **7**, 1–23 (2018).

28. M Thane, V Viswanathan, TC Meyer, E Paisios, M Schleyer, Modulations of microbehaviour by associative memory strength in Drosophila larvae. PLoS One **14**, e0224154 (2019).

29. B Szigeti, A Deogade, B Webb, Searching for motifs in the behaviour of larval Drosophila melanogaster and Caenorhabditis elegans reveals continuity between behavioural states. J. R. Soc. Interface **12** (2015).

30. KR Kaun, et al., Natural variation in food acquisition mediated via a Drosophila cGMP-dependent protein kinase. J. Exp. Biol. **210**, 3547–3558 (2007).

31. G Slater, P Levy, KL Andrew Chan, C Larsen, A central neural pathway controlling odor tracking in drosophila. J. Neurosci. **35**, 1831–1848 (2015).

32. M Klein, et al., Exploratory search during directed navigation in C. elegans and Drosophila larva. Elife **6**, 1–14 (2017).

33. A Gomez-Marin, M Louis, Active sensation during orientation behavior in the Drosophila larva: More sense than luck. Curr. Opin. Neurobiol. **22**, 208–215 (2012).

34. E Paisios, A Rjosk, E Pamir, M Schleyer, Common microbehavioral "footprint" of two distinct classes of conditioned aversion. Learn. Mem. **24**, 191–198 (2017).

35. B Gerber, et al., Pain-relief learning in flies, rats, and man: Basic research and applied perspectives. Learn. Mem. **21**, 232–252 (2014).

4. D Tadres, M Louis, PiVR: An affordable and versatile closed-loop platform to study unrestrained sensorimotor behavior. PLoS Biol. **18**, 1–25 (2020).

5. MJ Almeida-Carvalho, et al., The Ol1mpiad: Concordance of behavioural faculties of stage 1 and stage 3 Drosophila larvae. J. Exp. Biol. **220**, 2452–2475 (2017).

6. S Gerhard, I Andrade, RD Fetter, A Cardona, CM Schneider-Mizell, Conserved neural circuit structure across drosophila larval development revealed by comparative connectomics. Elife **6**, 1–17 (2017).

7. T Jovanic, Studying neural circuits of decision-making in Drosophila larva. J. Neurogenet. **34**, 162–170 (2020).

8. D Kim, M Alvarez, LM Lechuga, M Louis, Species-specific modulation of food-search behavior by respiration and chemosensation in Drosophila larvae. Elife **6**, 1–23 (2017).

9. RK Vijendravarma, S Narasimha, TJ Kawecki, Predatory cannibalism in Drosophila

36. M Schleyer, M Fendt, S Schuller, B Gerber, Associative learning of stimuli paired and unpaired with reinforcement: evaluating evidence from maggots, flies, bees, and rats. Front. psychology **9**, 1494 (2018).

37. N Koseki, et al., Individual differences in sensory responses influence decision making by Drosophila melanogaster larvae on exposure to contradictory cues. J. Neurogenet. **30**, 288–296 (2016).

38. T Niewalda, I Jeske, B Michels, B Gerber, 'Peer pressure' in larval Drosophila? Biol. Open **3**, 575–582 (2014).

39. TJ Prescott, P Redgrave, K Gurney, Layered control architectures in robots and vertebrates. Adapt. Behav. **7**, 99–127 (1999).

40. R Brooks, A Robust Layered Control System For A Mobile Robot. IEEE J. Robot. Autom. **RA-2**, 14–23 (1986).

41. S Lahiri, et al., Two alternating motor programs drive navigation in Drosophila larva. PLoS One **6** (2011).

42. B Michels, et al., Pavlovian conditioning of larval Drosophila: An illustrated, multilingual, hands-on manual for odor-taste associative learning in maggots. Front. Behav. Neurosci. **11**, 1–6 (2017).

43. AM Jürgensen, A Khalili, MP Nawrot, Reinforcement-mediated plasticity in a spiking model of the drosophila larva olfactory system. BMC Neurosci. 2019 **20(Suppl 1): P225**, 56 (2019).

44. AM Jürgensen, A Khalili, E Chicca, G Indiveri, MP Nawrot, A neuromorphic model of olfactory processing and sparse coding in the drosophila larva brain. bioRxiv (2021).

45. T Saumweber, et al., Functional architecture of reward learning in mushroom body extrinsic neurons of larval drosophila. Nat. communications **9**, 1–19 (2018).

46. D Owald, S Waddell, Olfactory learning skews mushroom body output pathways to steer behavioral choice in drosophila. Curr. opinion neurobiology **35**, 178–184 (2015).

47. D Owald, et al., Activity of defined mushroom body output neurons underlies learned olfactory behavior in drosophila. Neuron **86**, 417–427 (2015).

48. A Weiglein, F Gerstner, N Mancini, M Schleyer, B Gerber, One-trial learning in larval Drosophila. Learn. Mem. **26**, 109–120 (2019).

49. H Kohsaka, et al., Regulation of forward and backward locomotion through intersegmental feedback circuits in Drosophila larvae. Nat. Commun. **10**, 2654 (2019).

50. AA Zarin, B Mark, A Cardona, A Litwin-Kumar, CQ Doe, A multilayer circuit architecture for the generation of distinct locomotor behaviors in Drosophila. Elife **8** (2019).

51. J Berni, Genetic dissection of a regionally differentiated network for exploratory behavior in drosophila larvae. Curr. Biol. **25**, 1319–1326 (2015).

52. J Berni, SR Pulver, LC Griffith, M Bate, Autonomous circuitry for substrate exploration in freely moving drosophila larvae. Curr. Biol. **22**, 1861–1870 (2012).

53. T Zhang, A Branch, P Shen, Octopamine-mediated circuit mechanism underlying controlled appetite for palatable food in Drosophila. Proc. Natl. Acad. Sci. **110**, 15431–15436 (2013).

54. CA Malloy, et al., Pharmacological identification of cholinergic receptor subtypes: modulation of locomotion and neural circuit excitability in Drosophila larvae. Neuroscience **411**, 47–64 (2019).

55. C Eschbach, et al., Recurrent architecture for adaptive regulation of learning in the insect brain. Nat. Neurosci. **23**, 544–555 (2020).

56. K Vogt, et al., Internal state configures olfactory behavior and early sensory processing in Drosophila larvae. Sci. Adv. **7**, 1, eabd6900 (2021).

57. M Schleyer, et al., Identification of dopaminergic neurons that can both establish associative memory and acutely terminate its behavioral expression. J. Neurosci. **40**, 5990–6006 (2020).

58. A Schoofs, et al., Selection of Motor Programs for Suppressing Food Intake and Inducing Locomotion in the Drosophila Brain. PLoS Biol. **12** (2014).

59. CS Qian, M Kaplow, JK Lee, WB Grueber, Diversity of internal sensory neuron axon projection patterns is controlled by the POU-domain protein pdm3 in drosophila larvae. J. Neurosci. **38**, 2081–2093 (2018).

60. JS de Belle, M Heisenberg, Associative Odor Learning in Drosophila Abolished by Chemical Ablation of Mushroom Bodies. Sci. (80-.). **263**, 692–695y (1994).

61. A Strutz, et al., Decoding odor quality and intensity in the Drosophila brain. Elife **3**, e04147 (2014).

62. MJ Dolan, et al., Communication from Learned to Innate Olfactory Processing Centers Is Required for Memory Retrieval in Drosophila. Neuron **100**, 651–668.e8 (2018).

63. C Eschbach, et al., Circuits for integrating learnt and innate valences in the fly brain. Bioarxiv, 1–72 (2020).

64. I Tastekin, et al., Sensorimotor pathway controlling stopping behavior during chemotaxis in the Drosophila melanogaster larva. Elife **7**, 1–38 (2018).

65. J Gjorgjieva, J Berni, JF Evers, SJ Eglen, Neural circuits for peristaltic wave propagation in crawling drosophila larvae: Analysis and modeling. Front. Comput. Neurosci. **7**, 1–19 (2013).

66. C Pehlevan, P Paoletti, L Mahadevan, Integrative neuromechanics of crawling in D. melanogaster larvae. Elife **5**, 1–23 (2016).

67. D Ross, K Lagogiannis, B Webb, A model of larval biomechanics reveals exploitable passive properties for efficient locomotion. Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics) **9222**, 1–12 (2015).

68. J Loveless, K Lagogiannis, B Webb, Modelling the neuromechanics of exploration and taxis in larval Drosophila. PLoS Comput. Biol. **15**, 7 (2019).

69. A Davies, M Louis, B Webb, A Model of Drosophila Larva Chemotaxis. PLoS Comput. Biol. **11**, 1–24 (2015).

70. MN Günther, G Nettesheim, GT Shubeita, Quantifying and predicting Drosophila larvae crawling phenotypes. Sci. Rep. **6**, 1–10 (2016).

71. M Zago, F Lacquaniti, A Gomez-Marin, The speed– curvature power law in Drosophila larval locomotion. Biol. Lett. **12**, 3–6 (2016).

72. RS Marken, DM Shaffer, The power law of movement : an example of a behavioral illusion. Exp. Brain Res. **235**, 1835–1842 (2017).

73. M Zago, A Matic, T Flash, A Gomez-Marin, F Lacquaniti, The speed - curvature power law of movements: a reappraisal. Exp. Brain Res. **236**, 69–82 (2017).

74. RS Marken, DM Shaffer, The power law as behavioral illusion: reappraising the reappraisals. Exp. Brain Res. **236**, 1537–1544 (2018).

75. E Bicho, Ph.D. thesis (University of Minho) (1999).

76. PM Itskov, et al., Automated monitoring and quantitative analysis of feeding behaviour in Drosophila. Nat. Commun. **5** (2014).

77. N Otto, et al., Interactions among Drosophila larvae before and during collision. Sci. Rep. **6**, 1–11 (2016).

78. JB Masson, et al., Identifying neural substrates of competitive interactions and sequence transitions during mechanosensory responses in drosophila. PLOS Genet. **16**, 1–40 (2020).

79. T Jovanic, et al., Neural Substrates of Drosophila Larval Anemotaxis. Curr. Biol. **29**, 554–566.e4 (2019).

80. T Sokabe, HC Chen, J Luo, C Montell, A Switch in Thermal Preference in Drosophila Larvae Depends on Multiple Rhodopsins. Cell Rep. **17**, 336–344 (2016).

81. M Klein, et al., Sensory determinants of behavioral dynamics in Drosophila thermotaxis. Proc. Natl. Acad. Sci. U. S. A. **112**, E220–E229 (2015).

82. L Luo, et al., Navigational decision making in Drosophila thermotaxis. J. Neurosci. **30**, 4261–72 (2010).

83. ML Zhu, KJ Herrera, K Vogt, A Bahl, Navigational strategies underlying temporal phototaxis in Drosophila larvae. bioRxiv (2020).

84. J Felsenberg, et al., Integration of Parallel Opposing Memories Underlies Memory Extinction. Cell **175**, 709–722.e15 (2018).

85. C König, et al., Reinforcement signaling of punishment versus relief in fruit flies. Learn. Mem. **25**, 247–257 (2018).

86. AS Thum, B Gerber, Connectomics and function of a memory network: the mushroom body of larval drosophila. Curr. opinion neurobiology **54**, 146–154 (2019).

87. T Landgraf, et al., NeuroCopter: Neuromorphic computation of 6D ego-motion of a quad-copter in Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics). Vol. 8064 LNAI, pp. 143–153 (2013).

88. T Saumweber, et al., Behavioral and synaptic plasticity are impaired upon lack of the synaptic protein SAP47. J. Neurosci. **31**, 3508–3518 (2011).

89. B Michels, et al., Cellular site and molecular mode of synapsin action in associative learning. Learn. Mem. **18**, 332–344 (2011).

90. A Widmann, et al., Genetic Dissection of Aversive Associative Olfactory Learning and Memory in Drosophila Larvae. PLoS Genet. **12**, 1–32 (2016).

91. M Springer, MP Nawrot, A mechanistic model for reward prediction and extinction learning in the fruit fly. eNeuro (2021).

5   Discussion

This thesis comprises four manuscripts in three chapters that examine different aspects of the Trias of stimulus representation (Chapter 2.1), the formation of associations (Chapters 3.1, 3.2) and the initiation of anticipatory behavior in the insect MB (Chapters 3.1, 4.1), that we studied in computational models of the *Drosophila* olfactory pathway and MB with varying levels of abstraction and on different substrates (Chapters 2.1, 3.1, 3.2). Additionally, we used a locomotory model of the *Drosophila* larva body as a virtual agent in simulated experiments (Chapters 3.1, 4.1).

In a spiking network of the larval olfactory pathway and MB, we showed that both population and temporal sparseness can be achieved via feedback inhibition within the MB. Temporal sparseness can also be increased by spike frequency adaptation. We then successfully verified the robustness of these mechanisms by additionally implementing them on neuromorphic hardware. The *Drosophila* MB has been investigated extensively as a model for associative learning (larva: [42, 43, 47, 48, 94, 106], adult: [49–51, 53, 56, 87, 100, 107]), and population sparseness has been shown to increase pattern separation by decreasing overlap [22–24], which can be beneficial for learning [108–110]. In implementing mechanisms for increasing the sparseness of odor representations suggested playing a role in the insect antennal lobe and MB (lateral inhibition: [31, 111–114], feedback inhibition: [40, 115], spike-frequency adaptation: [116, 117]) and demonstrating that they enhance odor separability in this system, we investigated an important feature underlying efficient learning in the MB (Chapter 2.1). Additionally, we showed that increased distance between odor patterns prevents the *Drosophila* larva MB from generalizing between odors in a reward-learning experiment using a spiking network (Chapter 3.1).

The formation of associations between sensory inputs in the insect MB depends on detecting their coincidence [11–14]. In a spiking network of the *Drosophila* larva MB, we showed that this concept of coincidence detection at the synapses between KCs and MBONs can be extended in time by implementing an eligibility trace at the plastic synapses, creating a time window for the integration of sensory input and reinforcement (Chapter 3.1). Experimental evidence from the larval MB hints at the presence of such a trace [47, 118, 119]. We further investigated the temporal dynamics of associative learning, given prolonged simultaneous sensory input and reinforcement encoded by DANs. DANs play an essential role in the formation of associations in the MB and have been argued to respond to reinforcement, depending on the learning history in vertebrates [99, 120–123]. A similar role has been discussed in the *Drosophila* literature [42, 89, 91]. We implemented a feedback mechanism giving the reinforcement-mediating DANs access to the learning history and demonstrated that this affects the acquisition of associations in the MB by decreasing the driving force in the learning process; as the pairing of an odor with reinforcement continues, a mechanism referred to as prediction error (Chapter 3.1). This mechanism could explain the temporal dynamics observed in larval learning experiments (Chapter 3.1). Once associations of sensory stimuli with reinforcement are formed, the sensory stimuli can also serve as indirect reinforcement, a phenomenon referred to as second-order conditioning, which has also been discussed to be mediated by the DANs [124–126]. We explored mechanisms for DAN-mediated second-order conditioning in the MB using minimal rate models inspired by the adult *Drosophila* MB and suggest multiple biologically plausible feedback and feed-forward computational motifs for second-order conditioning (Chapter 3.2).

While computational models of the insect MB have been used successfully by us and many others to investigate the different mechanisms of learning in the insect MB [42, 108, 127–135] the difficulty of comparing MB readout with animal behavior in experiments remains an issue. We used a locomotory model of the larval *Drosophila* body to bridge this gap and allow direct comparisons between simulated and animal experiments with high temporal resolution (Chapters 3.1, 4.1). The virtual agent representing a single animal consists of a two-segmental body that combines forward movement with bending behavior, which characterizes active sensing. We showed that one of the influences on the agent's behavior is the MB readout, which represents the result of learning processes, directly affecting their movements towards or away from an odor source (Chapters 3.1, 4.1).

5.1   Mushroom body readout

In the learning MB models in this thesis (Chapters 3.1, 3.2), the difference between the output spikes or rates, respectively, of two MBONs was used to compute a behavioral tendency for approach or avoidance. Traditionally, most *Drosophila* MBONs have been viewed as approach or avoidance signaling, depending on their compartmental localization with different types of modulatory neurons (larva: [43, 66], adult:[44, 51, 53–55, 65, 88]). Approach and avoidance of stimuli is the observed and quantified measure of learning in the experiments that inspired the modeling work in this thesis. While serving its purpose in these modeling experiments, viewing the MB output as merely signaling either one of these opposite behaviors might be overly simplified, considering the complex integration and association processes occurring in the MB based on sensory input of different modalities that are integrated into the MB, as has been shown in adult [44–46, 136] and larval *Drosophila* [42, 43], as well as other insect species [137–139].

Throughout the literature, the labels *valence* and *preference* are used fairly inconsistently. Here, I will use them according to the following definitions:

Valence is the specific value of a stimulus to an animal that can either be positive, negative, or neutral. Valence can refer to an innate or learned property. Innate valence is based on the natural relevance of the stimulus to an animal (e.g. food)

Preference is an animal's tendency to bias its actions when confronted with a sensory stimulus/cue. It is an often behaviorally expressed indicator of the valence of the respective stimulus.

5.1.1   Mushroom body projections to pre-motor areas

Following the integration of multisensory external and potentially internal input and learning, the output of the MB has been shown to be involved in initiating behavior in *Drosophila* [51–53, 56] and the cockroach [140]. Additionally, Hancock et al. [52] showed that blocking γ1 MBON output abolishes specifically learned behavior. MBONs project to pre-motor areas, as demonstrated in *Drosophila* (larva: [42, 48], adult: [44, 49]) and the cockroach [141, 142]. Plasticity at the synapses between the KCs and MBONs induces a shift in the activity of MBONs in response to a learned

sensory cue (larva: [48], adult: [50–53]) and it has been suggested that this would also affect the activity of MBON downstream partners [53, 56]. MBONs in the adult *Drosophila* have been shown to be glutamatergic, GABAergic or cholinergic [44, 53] and innervate different compartments, depending on the transmitter they employ [44]. The general principle in the connectivity downstream of the MBONs seems to be convergence onto common targets [44], while some few MBONs have been demonstrated to innervate other MBONs [49]. Most MBON downstream partners are located in the same few regions [44]. Concerning motor control and the initiation of behavior, the central complex and the lateral horn (LH) are especially relevant MBON targets [44]. The central complex is innervated by MBONs directly or via interneurons [44] and is a highly relevant region for directed locomotion in *Drosophila* [143].

5.2   Using computational models to study learning in the mushroom body

Neuroscience research investigates, among other things, which structure or circuit in the brain is responsible for which function and potentially the underlying mechanism. One of the tools used in this is models. The modeling approach can be driven either by a bottom-up or top-down perspective. Bottom-up approaches are inspired by knowledge about structures, such as the connectome of a brain region. They allow us to test hypotheses about the localization of a particular function or even its exact mechanism. On the other hand, top-down models seek to explore potential circuits or mechanisms behind observable phenomena. They allow for a more explorative perspective and can generate hypotheses for experimental testing.

Another way of categorizing types of models can be their level of abstraction. This can range from purely mathematical models designed to capture relationships between variables to circuit models of varying degrees of richness of details. An example of the first type would be the Rescorla-Wagner model [144], an old yet influential model of associative learning. The advantage of such models is that they can offer direction and a framework that overarches species and scientific disciplines. They can contribute to identifying the biological implementation of their conceptual components (Chapter 3.1). More detailed models of neurons and circuits (Chapters 3.1, 3.2) can range from rate models that reduce each neuronal unit to its output rate to including the morphology of individual neurons (for an overview, see [145]). The level choice depends on the specific question targeted, as a more naturalistic model is likely more difficult to interpret, a consideration of the computational resources required, and the knowledge about the neurons/circuits in question. For example, a very detailed type of model that considers the characteristics of different types of ion channels requires a lot more detailed prior knowledge about the neuron characteristics to yield meaningful results. It stands to reason that sometimes, a more abstract and simple model could be the better choice.

When using a type of model that includes the process of integration of input into neurons, the timing of spikes can be extracted. This is possible whenever the membrane voltage is recorded and threshold crossings are interpreted as spike events. This applies to many different spiking neuron models, from complex Hodgkin-Huxley type models [145] to more simple integrate-and-fire type neurons [145]. The evolution of the membrane potential, as a response to external or network input and (to a varying extent) neuron-internal changes, can be read out. This allows the capture of a spike train at high temporal resolution. Rate models, on the other hand, bypass this integration

process and rely on neuron units that receive and output only a rate. Often, they can have great explanatory power for effects regarding the circuit or external input. Since the information about the timing of spikes is not captured, they don't allow the study of effects with high fluctuation. They are also, to some extent, blind to phenomena of spike-timing, like adaptation or bursting.

Overall, different types of models are very valuable tools in neuroscience, making different circuit elements easily accessible for readout to generate or test hypotheses. Both bottom-up and top-down approaches can contribute helpful insights. An additional advantage is that computational models allow for very clean and controlled experiments. To fully exploit their potential, they have to be combined with experimental approaches that provide models with relevant underlying anatomical and functional data or test the model predictions generated in an artificially reduced and constrained circuit in the complete animal brain.

5.2.1   Behavioral measures of Drosophila learning

One of the challenges of transferring knowledge between modeling and animal experiments is the different types of learning measures used. Adult and larval *Drosophila* in their normal life outside a lab are confronted with different levels of challenges and environmental complexity. While the adult fly, as a flying insect, lives in a higher dimensional world than the larva, it is also faced with a broader range of tasks. It needs to locate food, mates, and egg-laying sites and perform a range of behaviors at these locations [61]. The larva, on the other hand, spends this entire phase of the *Drosohila* life span on its food source (the egg-laying site), and its main focus is on feeding [61]. In learning experiments, flies are usually trained to learn the positive or negative valence of a stimulus. Thus, behavioral measures of this valence focus on a resulting behavioral preference towards approach or avoidance of these stimuli.

In most olfactory learning experiments conducted with larvae, the animals are exposed to odors inside a covered petri dish layered with a rewarding food [146] or punishing salt [147] substrate. During a test phase, they are allowed to move around in another dish freely, given a choice between a trained vs. untrained odor on each side in the standard protocol [146, 147] or the trained vs. no odor [148]. After a given amount of time, the number of larvae on each side is compared. These protocols assume that the animals tend to orient their heads towards a previously rewarded odor source and then move towards it and vice versa for punished odors. When an attractive odor is detectable in the air, foraging consists of a combination of crawling and turning [149]. Larvae employ active sensing to navigate along chemical gradients [150] towards an odor source. Their forward locomotion is accompanied by lateral bending [151] for active sensing and tracking the gradient. Moving towards an appetitive odor source, accompanied by an increase in odor concentration, suppresses bending and turning [150, 152, 153]. This leads to a straighter, faster movement toward the odor sources and increases the number of animals on the respective side of the dish. The opposite is true for an aversive odor source.

In most experiments conducted with adult flies, electric shock was used as punishment. Tully and Quinn first developed the punishing protocol [154] and it has been used continuously since then, even though the mechanical implementations have been updated [155]. The protocols feature odors in an airstream in combination with pulses of electric shock. Appetitive conditioning protocols use water in thirsty animals [95] or sugar, following a protocol as introduced by Tempel et al. [156]. Preference is usually

quantified in a T-maze setup, presenting a reinforced training odor on one side and another odor, never paired with reinforcement on the other side. The T-maze, as a setup for testing preference, easily allows quantification of this preference by counting the number of animals on either side. Aside from locomotory behavior, the extension (or retraction) of the proboscis, which is a part of the animal's feeding behavior, was established as an indicator of odor preference [157] but seems rarely used.

5.2.2   Comparing measures of preference between model and animal experiments

Preference, as measured in behavioral experiments, is thus not directly comparable with the typically used MBON output reported in modeling studies of insect MBs [42, 108, 127–134]. While some *Drosophila* lab experiments report a measure of preference directly at the MB output by quantifying the activity of MBONs (larva: [48], adult: [50–53]), often the model MB readout has to be compared to animal behavior. Chapters 3.1 and 4.1 present one method for doing that by linking the simulated MB output with the locomotory control architecture of a virtual larval agent that accounts for the odor preference, as read out from the MBONs, when initiating locomotory behavior in a virtual petri-dish that resembles the test environment in larval lab experiments [146, 147]. Similar simulations of behavior have been used by others to enhance comparability with experiments of *Drosophila* behavior in learning experiments (larva: [132, 158], adult: [135]). The larva, with its robust body, also lends itself to implementation in robotics beyond just virtual agents [159]. Adding the component of a physical substrate to agent behavior can make the integration of environmental stimuli and the learning process even more realistic and enforce the synchronization of the MB learning and agent behavior temporal resolution. If the sensors and the materials are well chosen, robots can add a more realistic level of sensory feedback about both the environment and their own body [159, 160] compared to a virtual agent. Additionally, passive dynamics can add another layer of complexity to the larval body and create more robust and naturalistic behavior in the physical world [161].

In the following, I will, among other things, discuss three additional components relevant to the framework of stimulus representation (Chapter 2.1), learning (Chapters 3.1, 3.2), and the initiation of anticipatory behavior (Chapters 3.1, 4.1) that are not covered in these chapters. These components are highlighted in red in Figure 2 and refer to the influence of innate valence of stimuli and that of internal states on the three components of the Trias.
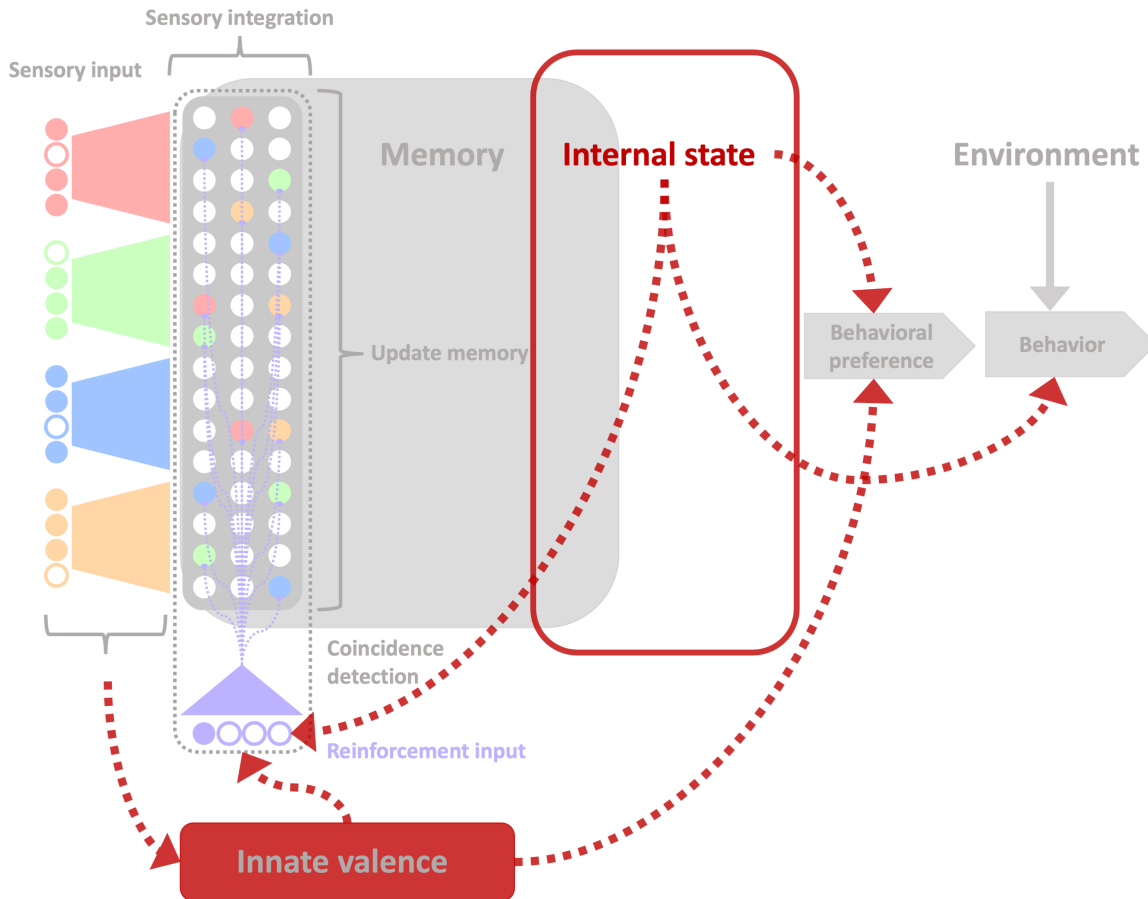
Figure 3: **The role of innate valence and internal state on the proposed relationship between sensory integration, learning, and the initiation of behavior.** Dense multisensory input is integrated into higher-order processing centers, where coincidences among them and with reinforcement are detected, and associations are formed between concurrently active sensory inputs, internal states, and reactivated prior knowledge (learned or innate). The initiation of behavior depends on the integrated output of the higher-order processing center, signaling learned valence, compared against behavioral preferences based on potential innate valences. The execution of a behavioral preference can be influenced by an animal's abilities in the respective situation and by environmental factors. Colored in red are the motif components not included in this thesis's three chapters.

5.3   The influence of internal states

Internal states can modify sensory integration and learning in the MB and thus also influence the MB output [107, 162–164]. Internal states such as hunger or fear represent strong motivational forces in animal behavior. In the γ-lobe of the MB in adult *Drosophila*, hunger has been shown to modify the synapses between the KCs and the MBONs in olfactory learning by altering the KC odor response in the presence of sugar after fasting [162] and potentiating synaptic transmission [162] through an increased salience of a sugar reward after fasting. Siju et al. [163] varied the hunger state of adult *Dsosophila* between two starved and a single fed condition and observed modulation of the sugar response of DANs in some compartments. The state variation between mated and virgin also modulates learning at the synapses between the KCs and MBONs [107, 163]. Mating increased the response of DANs innervating the MB to pheromones [107, 163], as well as enabling pheromone-triggered activation of the β'1-MBONs. Overall, there is evidence that both behavioral and intrinsic motivational states influence sensory integration and learning in the MB on multiple levels.

5.4   The integration of innate and learned valences

Both in larval [165–168] and adult [169] *Drosophila* many odors appear to have some innate positive or negative valence. Such innate valences lead to preferences that have been shown to initiate behavior (larva: [48, 167], adult: [169, 170]). This could potentially affect learning processes and behavior initiation on different levels. Firstly, the formation of an association between reinforcement and an odor that has an innate valence might differ from the formation of an association with a neutral odor. It seems important to understand how the innate valence of a stimulus might affect the dynamics of the acquisition of a reinforcement association with it, as well as how associative learning modifies the innate valence of a given stimulus. Secondly, an innate odor valence might influence the behavioral preference generated by the MB downstream of the associative learning process and thus affect the actual initiation of the respective behavior. Determining the site of the integration of an innate and learned valence of a stimulus is thus a relevant question that targets whether or not innate and learned valence is encoded in separate pathways and, if so, where they converge.

5.4.1   Encoding of innate valence in the lateral horn

It has been argued that innate valences of stimuli might be encoded in the LH [169, 171–173] and suggested that the integration of innate and learned valences takes place downstream of the MB at a point of convergence between the MB and the LH pathways (larva:[48], adult: [169, 173]). Both types of valences are likely encoded in different pathways in larval [43, 48, 61, 63] and adult [174, 175] *Drosophila*. In the olfactory system, projection neurons project to both the MB and the LH in a divergent manner (larva: [43, 61], adult: [76, 174, 176, 177]). Amin et al. [176] described a potential motif for representing innate valence among the LH neurons. Some types of receptors are very specific to biologically highly relevant stimuli, such as pheromones, as shown in adult *Drosophila* [175, 178, 179]. They might lay grounds for particular hard-wired connections to the LH, allowing the LH to encode their innate valence. Additionally, glomeruli in the antennal lobe are distinctly activated by either aversive or attractive odors in adult *Drosophila* [180]. While some LH neurons are broadly tuned, some

respond specifically to a particular type of projection neuron from the antennal lobe, potentially transferring the encoded valence to the LH [173]. These authors show that the representations of innately positively and negatively valent odors are clearly segregated among the LH neurons [173]. The valence encoded in the activation pattern of the antennal lobe glomeruli would not persist in the MB due to the random connectivity between the antennal lobe projection neurons and the MB intrinsic neurons [46]. However, recently, the postulated strict roles of the MB and the LH in encoding learned vs. innate valence, respectively, have been challenged by the observation that context-dependent memory might be mediated by neurons in the LH, independently of the MB [181]. Furthermore, the MB has been shown to be involved in some specific innate behavioral preferences in adult *Drosophila* [164, 182, 183].

5.4.2   Sites of integration

Eschbach et al. [48] studied the MB output region in the larva and found a convergence of the MB and LH pathway onto neurons downstream of the MB output neurons, which they, therefore, term convergence neurons. They investigate different MB output neurons that either promote or suppress turning behavior and find inhibitory interactions between those encoding different behavioral preferences, as well as synapses of these MBONs, onto LH neurons and vice versa. Additionally, they more closely examined a particular convergence neuron and found that it receives both functionally excitatory and inhibitory feedback from MB output neurons, encoding approach and avoidance, respectively, and input from the LH [48]. They also found that an innately positive odor causes activation of this particular convergence neuron, even in the absence of MB input [48]. They argued that these convergence neurons might be the location of innate and learned valence integration and demonstrated that their optogenetic activation could suppress turning and foster crawling behavior [48]. In the adult, a convergence pathway has been found in some LH output neurons [44, 169]. These neurons receive input from projection neurons in the antennal lobe, as well as from particular MBONs [44, 169]. The authors demonstrate that these LH output neurons are relevant for encoding innate valence and necessary for memory retrieval and behavior after training [169].

The involvement of DANs in the integration of innate and learned valences has recently been investigated in larval [48] and adult [184] *Drosophila*. Huang et al. [184] showed that DANs that are activated by punishment are also activated by odors that elicit innate avoidance behavior in their experiments. Inspired by this observation, they demonstrate that these dopaminergic neurons integrate innate and learned valences of odors. The dopaminergic teaching signal for plasticity at the synapses between MB intrinsic and output neurons is thus modified by the intrinsic valence of odors used in learning experiments [184]. In the larva, connections between convergence neurons and dopaminergic neurons have been found. Some of these connections are indirect via interneurons [48].

5.4.3   The bi-directional interaction between innate and learned valence

Many odors are innately valent in *Drosophila* (larva: [165, 166, 185], adult: [180, 185, 186]), and an innate valence might affect learning about associations with that stimulus, as well as learning might alter the innate valence of it. Saumweber et al. [167] showed that the strength of the innate behavioral preference in larvae can vary between odors.

When comparing the preference for different odors against the same control order after learning, they found that for some odors, the strength of the innate preference influenced the preference after learning. For these odors, higher innate preference caused a higher preference after learning [167]. In the case of another combination of odors with different innate valences, the larvae showed a similar preference after training for both odors compared to the same control odor, respectively [167]. Their experiments also revealed that while odors differ in the degree of innate preference they elicit, this preference also depends on odor concentration, which can shift the relative innate preference between two odors [167]. Few learning experiments in *Drosophila* directly targeted the change of behaviorally expressed stimulus preference over the time course of the experiment. In the larva, it has been shown that pairing with reinforcement can enhance the innate valence of odors [167], while pairing an innately attractive odor with punishment could reverse the odor valence in adult *Drosophila* [154]. While direct tests of this are scarce, other studies show differences in MBON output before and after training, indicating a change in the preference during learning [50–52, 56]. Behaviorally expressed preferences in the larva can also be enhanced [187–190] or shifted [189, 190] concerning what we can assume about their innate preferences [165–168]. The direction depends on the type of reinforcement used.

5.5 Conclusion

In the manuscripts that comprise this thesis, we covered different aspects of the Trias of sensory representation, learning, and anticipatory behavior. Using the *Drosophila* olfactory pathway and the MB, we contributed to understanding the mechanisms and relevance of sparse sensory representations (Chapter 2.1), key features of associative learning, such as coincidence detection and prediction error (Chapters 3.1, 3.2), as well as the link between MB output, behavioral preference, and adaptive behavior (Chapters 3.1, 4.1). The discussion covered two main areas in which the models used in these manuscripts fall short of capturing relevant components of generating a behavioral preference: The integration of internal states and innate valences.

The exclusion of innate stimulus valences from the processing and learning framework causes difficulties in the comparability with behavioral experiments, especially when studying the temporal dynamics of the acquisition of associations, and has produced some initially unexpected results because the animals seem not to be naive at the beginning of the experiments (Chapter 3.1).

Likewise, internal states are missing here but have been shown to influence learning in the MB [107, 162–164]. It is rare for models of MB learning to include internal states [132]. Jiang et al. [132] included fluctuations in the activity of DANs in their network that represent transitions between different internal states [162, 163]. They show that this allows for associations between such states and sensory stimuli that can co-exist with other learned associations between stimuli and reinforcement. This suggests that accounting for varying states might be a valuable addition to models to reproduce more realistic animal behavior.

5.6   Towards comprehensive circuit models for parallel valence representation and integration

Stimuli of innate valence (larva: [48, 167], adult: [169, 170]) or that acquired through associative learning (larva: [43, 66], adult:[44, 51, 53–55, 65, 88]) elicit a range of approach or avoidance behaviors in *Drosophila*. Associative learning comprises various processes, such as acquisition, maintenance, forgetting on different timescales, and extinction learning, which requires associations to be retained for potential recovery. Some of these phenomena, such as the loss of an association [144, 191–194] either due to extinction or forgetting, higher order conditioning without a decline in the association [124, 195], or the association of a single stimulus with positive and negative reinforcement [93, 196] appear contradictory and would likely be mutually exclusive if encoded within a single valance representation. The coexistence of these phenomena hints at the existence of multiple, parallel innate and learned valence representations.

Consequently, the LH and MB output circuitry must retain multiple valences in parallel. Their integration is required momentarily to initiate behavior. The temporary blend of sensory stimuli and contextual information at retrieval presumably determines the exact combination of valence representations. The highly parallel, compartmentalized output circuitry of the MB enables the different compartments to encode different valences [48, 50–53]. The combined output across multiple compartments yields an additive effect on behaviorally expressed valence [53], attributed to the combination of excitatory and inhibitory neurotransmitters employed by MB output neurons encoding opposing valences [197, 198]. In combination with interactions between MBONs of different compartments [44, 51, 88] and inter-compartment interactions between MBONs and modulatory neurons [42, 44, 49], several mechanisms emerge that might contribute to an integration of different learned valences at the output of the MB. A candidate mechanism for integrating innate and learned valence is the convergence of the LH and MB pathways [44, 48, 169] as the activity of their convergence neurons has been shown to initiate behavior [48, 169]. Additionally, dopaminergic neurons are activated by reinforcement and odors that elicit innate avoidance behavior [184] and connections between convergence neurons and dopaminergic neurons have been found in the larva [48].

To date, models of learning in the MB have mostly simplified the underlying circuit structure and viewed its output as a unified valence, resulting in a preference for approach or avoidance[104, 128–131, 134]. Few have explored more diverse output circuits [132, 135] to enable parallel representations in a more variable and interconnected output circuitry but ignore the LH pathway and the role of innate valence. Building on this, the development of more comprehensive models that consider the parallel MB and LH output circuitry and their convergence emerge as a promising continuation of prior research to investigate further the mechanisms underlying the initiation of behavior. Using computational circuit models, the different components of the MB and LH output circuitry could be queried continuously throughout a learning process to provide simultaneous insights into parallel valence representations and their interactions and integration at retrieval.

6   References

[1] Hermann Ebbinghaus. *Memory: A Contribution to Experimental Psychology.* Dover Publications, 1964. DOI: 10.5214/ans.0972.7531.200408.

[2] Alan Baddeley, Michael W. Eysenck, and Michael C. Anderson. *Memory.* Psychology Press, 2015. DOI: 10.4324/9781315749860.

[3] Eric R. Kandel, Yadin Dudai, and Mark R. Mayford. "The molecular and systems biology of memory". In: *Cell* 157.1 (2014), pp. 163–186. DOI: 10.1016/j.cell.2014.03.001.

[4] Paul W. Frankland, Sheena A. Josselyn, and Stefan Köhler. "The neurobiological foundation of memory retrieval". In: *Nature neuroscience* 22.10 (2019), pp. 1576–1585. DOI: 10.1038/s41593-019-0493-1.

[5] Androulla Ioannou and Xenia Anastassiou-Hadjicharalambous. "Non-associative Learning". In: *Encyclopedia of Evolutionary Psychological Science.* Ed. by Todd K. Shackelford and Viviana A. Weekes-Shackelford. Springer International Publishing, 2018, pp. 1–13. DOI: 10.1007/978-3-319-16999-6_1027-1.

[6] Chi-Sang Poon and Susanne Schmid. "Nonassociative Learning". In: *Encyclopedia of the Sciences of Learning.* Ed. by Norbert M. Seel. Springer US, 2012, pp. 2475–2477. DOI: 10.1007/978-1-4419-1428-6_1849.

[7] Jeremie Jozefowiez. "Associative Learning". In: *Encyclopedia of the Sciences of Learning.* Ed. by Norbert M. Seel. Springer US, 2012, pp. 330–334. DOI: 10.1007/978-1-4419-1428-6_703.

[8] Joseph E. Steinmetz, Kim Jeansok, and Richard F. Thompson. In: *Handbook of psychology.* Ed. by M. Gallagher and R. Nelson. Wiley, 2003.

[9] Burrhus F. Skinner. "The experimental analysis of operant behavior: A history." In: *Psychology: Theoretical-historical perspectives.* Ed. by R. Rieber and K. Salzinger. American Psychological Association, 1998, pp. 289–299. DOI: 10.1037/10276-011.

[10] Ivan Pavlov. *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex.* Oxford University Press, 1972.

[11] Stopfer. Mark. "Central processing in the mushroom bodies". In: *Current Opinion in Insect science* 6.301 (2014), pp. 99–103. DOI: 10.1016/j.cois.2014.10.009.

[12] Martin Heisenberg. "Mushroom body memoir: from maps to models". In: *Nature Reviews Neuroscience* 4.4 (2003), pp. 266–275. DOI: 10.1038/nrn1074.

[13] Ronald L. Davis. "Mushroom bodies and drosophila learning". In: *Neuron* 11.1 (1993), pp. 1–14. DOI: 10.1016/0896-6273(93)90266-T.

[14] Josh Dubnau. "Ode to the Mushroom Bodies". In: *Science* 335.6069 (2012), pp. 664–665. DOI: 10.1126/science.1218171.

[15] Jennifer L. Raymond and Stephen G. Lisberger. "Neural learning rules for the vestibulo-ocular reflex". In: *Journal of Neuroscience* 18.21 (1998), pp. 9112–9129.

[16]   Dmitry V. Sarkisov and Samuel S.-H. Wang. "Order-dependent coincidence detection in cerebellar Purkinje neurons at the inositol trisphosphate receptor". In: *Journal of Neuroscience* 28.1 (2008), pp. 133–142. DOI: 10.1523/JNEUROSCI.1729-07.2008.

[17]   Shogo Ohmae and Javier F. Medina. "Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice". In: *Nature neuroscience* 18.12 (2015), pp. 1798–1803. DOI: 10.1038/nn.4167.

[18]   Cleveland P. Hickman. "Zoologie". In: Pearson Deutschland GmbH, 2008. Chap. Nervöse Steuerung, pp. 1067–1100.

[19]   Eric R. Kandel, James H. Schwartz, Thomas M. Jessell, et al. *Principles of neural science*. Vol. 4. McGraw-hill New York, 2000. DOI: 10.1086/670559.

[20]   Simon B. Laughlin and Terrence J. Sejnowski. "Communication in neuronal networks". In: *Science* 301.5641 (2003), pp. 1870–1874. DOI: 10.1126/science.1089662.

[21]   Bruno A. Olshausen and David J. Field. "Sparse coding of sensory inputs". In: *Current Opinion in Neurobiology* 14.4 (2004), pp. 481–487. DOI: 10.1016/j.conb.2004.07.007.

[22]   Horace B. Barlow. "Sensory mechanisms, the reduction of redundancy, and intelligence". In: *Mechanisation of thought processes*. Her Majesty's Stationery Office, 1959, pp. 535–539.

[23]   David Marr and W. Thomas Thach. "A Theory of Cerebellar Cortex". In: *From the Retina to the Neocortex*. Ed. by Lucia Vaina. Birkhäuser Boston, 1991, pp. 11–50. DOI: 10.1007/978-1-4684-6775-8_3.

[24]   James S. Albus. "A theory of cerebellar function". In: *Mathematical Biosciences* 10 (1971), pp. 25–61. DOI: 10.1016/0025-5564(71)90051-4.

[25]   Iori Ito, Rose Chik-Ying Ong, Baranidharan Raman, et al. "Sparse odor representation and olfactory learning". In: *Nature neuroscience* 11.10 (2008), pp. 1177–1184. DOI: 10.1038/nn.2192.

[26]   Roger Herikstad, Jonathan Baker, Jean-Philippe Lachaux, et al. "Natural movies evoke spike trains with low spike time variability in cat primary visual cortex". In: *Journal of Neuroscience* 31.44 (2011), pp. 15844–15860. DOI: 10.1523/JNEUROSCI.5153-10.2011.

[27]   Bilal Haider, Matthew R. Krause, Alvaro Duque, et al. "Synaptic and Network Mechanisms of Sparse and Reliable Visual Cortical Activity during Nonclassical Receptive Field Stimulation". In: *Neuron* 65.1 (2010), pp. 107–121. DOI: 10.1016/j.neuron.2009.12.005.

[28]   Chris Häusler, Alex Susemihl, and Martin P. Nawrot. "Natural image sequences constrain dynamic receptive fields and imply a sparse code". In: *Brain research* 1536 (2013), pp. 53–67. DOI: 10.1016/j.brainres.2013.07.056.

[29]   Hannes Rapp and Martin P. Nawrot. "A spiking neural program for sensorimotor control during foraging in flying insects". In: *Proceedings of the National Academy of Sciences* 117.45 (Nov. 2020), pp. 28412–28421. DOI: 10.1073/pnas.2009821117.

[30] Farzad Farkhooi, Anja Froese, Eilif Muller, et al. "Cellular adaptation facilitates sparse and reliable coding in sensory pathways". In: *PLoS computational biology* 9.10 (2013), e1003251. DOI: `10.1371/journal.pcbi.1003251`.

[31] Javier Perez-Orive, Ofer Mazor, Glenn C. Turner, et al. "Oscillations and sparsening of odor representations in the mushroom body". In: *Science* 297.5580 (2002), pp. 359–365. DOI: `10.1126/science.1070502`.

[32] Ron A. Jortner, Sarah Farivar, and Gilles Laurent. "A simple connectivity scheme for sparse coding in an olfactory system". In: *The Journal of neuroscience* 27.7 (2007), pp. 1659–1669. DOI: `10.1523/JNEUROSCI.4171-06.2007`.

[33] Luca A. Finelli, Seth Haney, Maxim Bazhenov, et al. "Synaptic Learning Rules and Sparse Coding in a Model Sensory System". In: *PLOS Computational Biology* 4.4 (2008), e1000062. DOI: `10.1371/journal.pcbi.1000062`.

[34] Peter Kloppenburg and Martin P. Nawrot. "Neural coding: Sparse but on time". In: *Current biology* 24.19 (2014), R957–R959. DOI: `10.1016/j.cub.2014.08.041`.

[35] Cindy Poo and Jeffry S. Isaacson. "Odor Representations in Olfactory Cortex: "Sparse" Coding, Global Inhibition, and Oscillations". In: *Neuron* 62.6 (2009), pp. 850–861. DOI: `10.1016/j.neuron.2009.05.022`.

[36] Jason Wolfe, Arthur R. Houweling, and Michael Brecht. "Sparse and powerful cortical spikes". In: *Current opinion in neurobiology* 20.3 (2010), pp. 306–312. DOI: `10.1016/j.conb.2010.03.006`.

[37] Jeffry S. Isaacson. "Odor representations in mammalian cortical circuits". In: *Current opinion in neurobiology* 20.3 (2010), pp. 328–331. DOI: `10.1016/j.conb.2010.02.004`.

[38] Tomáš Hromádka, Michael R. DeWeese, and Anthony M. Zador. "Sparse representation of sounds in the unanesthetized auditory cortex". In: *PLoS Biology* 6.1 (2008), e16. DOI: `10.1371/journal.pbio.0060016`.

[39] N. Schweighofer, K. Doya, and F. Lay. "Unsupervised learning of granule cell sparse codes enhances cerebellar adaptive control". In: *Neuroscience* 103.1 (2001), pp. 35–50. DOI: `10.1016/S0306-4522(00)00548-0`.

[40] Andrew C. Lin, Alexei M. Bygrave, Alix De Calignon, et al. "Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination". In: *Nature neuroscience* 17.4 (2014), pp. 559–568. DOI: `10.1038/nn.3660`.

[41] Ashok Litwin-Kumar, Kameron Decker Harris, Richard Axel, et al. "Optimal degrees of synaptic connectivity". In: *Neuron* 93.5 (2017), pp. 1153–1164. DOI: `10.1016/j.neuron.2017.01.030`.

[42] Claire Eschbach, Akira Fushiki, Michael Winding, et al. "Recurrent architecture for adaptive regulation of learning in the insect brain". In: *Nature Neuroscience* 23.4 (2020), pp. 544–555. DOI: `10.1038/s41593-020-0607-9`.

[43] Katharina Eichler, Feng Li, Ashok Litwin-Kumar, et al. "The complete connectome of a learning and memory centre in an insect brain". In: *Nature* 548.7666 (2017), pp. 175–182. DOI: `10.1038/nature23455`.

[44]   Feng Li, Jack W. Lindsey, Elizabeth C. Marin, et al. "The connectome of the adult Drosophila mushroom body provides insights into function". In: *Elife* 9 (2020), e62576. DOI: 10.7554/eLife.62576.

[45]   Colleen Kirkhart and Kristin Scott. "Gustatory learning and processing in the Drosophila mushroom bodies". In: *Journal of neuroscience* 35.15 (2015), pp. 5950–5958. DOI: 10.1523/JNEUROSCI.3930-14.2015.

[46]   Sophie J.C. Caron, Vanessa Ruta, Larry F. Abbott, et al. "Random convergence of olfactory inputs in the Drosophila mushroom body". In: *Nature* 497.7447 (2013), pp. 113–117. DOI: 10.1038/nature12063.

[47]   Timo Saumweber, Astrid Rohwedder, Michael Schleyer, et al. "Functional architecture of reward learning in mushroom body extrinsic neurons of larval Drosophila". In: *Nature communications* 9.1 (2018), p. 1104. DOI: 10.1038/s41467-018-03130-1.

[48]   Claire Eschbach, Akira Fushiki, Michael Winding, et al. "Circuits for integrating learned and innate valences in the insect brain". In: *eLife* 10 (2021), e62567. DOI: 10.7554/eLife.62567.

[49]   Yoshinori Aso, Daisuke Hattori, Yang Yu, et al. "The neuronal architecture of the mushroom body provides a logic for associative learning". In: *eLife* 3 (2014), e04577. DOI: 10.7554/eLife.04577.

[50]   Toshihide Hige, Yoshinori Aso, Mehrab N. Modi, et al. "Heterosynaptic plasticity underlies aversive olfactory learning in Drosophila". In: *Neuron* 88.5 (2015), pp. 985–998. DOI: /10.1016/j.neuron.2015.11.003.

[51]   David Owald and Scott Waddell. "Olfactory learning skews mushroom body output pathways to steer behavioral choice in Drosophila". In: *Current opinion in neurobiology* 35 (2015), pp. 178–184. DOI: /10.1016/j.conb.2015.10.002.

[52]   Clare E. Hancock, Vahid Rostami, El Yazid Rachad, et al. "Visualization of learning-induced synaptic plasticity in output neurons of the Drosophila mushroom body γ-lobe". In: *Scientific Reports* 12.1 (2022), p. 10421. DOI: 10.1038/s41598-022-14413-5.

[53]   Yoshinori Aso, Divya Sitaraman, Toshiharu Ichinose, et al. "Mushroom body output neurons encode valence and guide memory-based action selection in Drosophila". In: *elife* 3 (2014), e04580. DOI: 10.7554/eLife.04580.

[54]   Toshihide Hige, Yoshinori Aso, Gerald M. Rubin, et al. "Plasticity-driven individualization of olfactory coding in mushroom body output neurons". In: *Nature* 526.7572 (2015), pp. 258–262. DOI: 10.1038/nature15396.

[55]   Julien Séjourné, Pierre-Yves Plaçais, Yoshinori Aso, et al. "Mushroom body efferent neurons responsible for aversive olfactory memory retrieval in Drosophila". In: *Nature neuroscience* 14.7 (2011), pp. 903–910. DOI: /10.1038/nn.2846.

[56]   David Owald, Johannes Felsenberg, Clifford B. Talbot, et al. "Activity of defined mushroom body output neurons underlies learned olfactory behavior in Drosophila". In: *Neuron* 86.2 (2015), pp. 417–427. DOI: 10.1016/j.neuron.2015.03.025.

[57]   Hugo J. Bellen, Chao Tong, and Hiroshi Tsuda. "100 years of Drosophila research and its impact on vertebrate neuroscience: a history lesson for the future". In: *Nature Reviews Neuroscience* 11.7 (2010), pp. 514–522. DOI: 10.1038/nrn2839.

[58] Hiroshi Kohsaka and Akinao Nose. "Optogenetics in Drosophila". In: *Optogenetics: Light-Sensing Proteins and Their Applications in Neuroscience and Beyond*. Ed. by Hiromu Yawo, Hideki Kandori, Amane Koizumi, et al. Springer Singapore, 2021, pp. 309–320. DOI: 10.1007/978-981-15-8763-4_19.

[59] Robert J. Riemensperger Thomasand Kittel and André Fiala. "Optogenetics in Drosophila Neuroscience". In: *Optogenetics: Methods and Protocols*. Ed. by Arash Kianianmomeni. Springer New York, 2016, pp. 167–175. DOI: 10.1007/978-1-4939-3512-3_11.

[60] Jacob G. Bernstein, Paul A. Garrity, and Edward S. Boyden. "Optogenetics and thermogenetics: technologies for controlling the activity of targeted cells within intact neural circuits". In: *Current Opinion in Neurobiology* 22.1 (2012), pp. 61–71. DOI: 10.1016/j.conb.2011.10.023.

[61] Bertram Gerber and Reinhard F. Stocker. "The Drosophila larva as a model for studying chemosensation and chemosensory learning: a review". In: *Chemical senses* 32.1 (2007), pp. 65–89. DOI: 10.1093/chemse/bjl030.

[62] Lily Kahsai and Troy Zars. "Learning and Memory in Drosophila: Behavior, Genetics, and Neural Systems". In: *International Review of Neurobiology* 99 (2011), pp. 139–167. DOI: 10.1016/B978-0-12-387003-2.00006-9.

[63] Michael Winding, Benjamin D. Pedigo, Christopher L. Barnes, et al. "The connectome of an insect brain". In: *Science* 379.6636 (2023), eadd9330. DOI: 10.1126/science.add9330.

[64] Louis K. Scheffer, Shan Xu, Michal Januszewski, et al. "A connectome and analysis of the adult *Drosophila* central brain". In: *eLife* 9 (2020), e57443. DOI: 10.7554/eLife.57443.

[65] Shin-ya Takemura, Yoshinori Aso, Toshihide Hige, et al. "A connectome of a learning and memory center in the adult *Drosophila* brain". In: *eLife* 6 (2017), e26975. DOI: 10.7554/eLife.26975.

[66] Andreas S. Thum and Bertram Gerber. "Connectomics and function of a memory network: the mushroom body of larval Drosophila". In: *Current opinion in neurobiology* 54 (2019), pp. 146–154. DOI: 10.1016/j.conb.2018.10.007.

[67] Guangwei Si, Jessleen K. Kanwal, Yu Hu, et al. "Structured Odorant Response Patterns across a Complete Olfactory Receptor Neuron Population". In: *Neuron* 101.5 (2019), 950–962.e7. DOI: 10.1016/J.NEURON.2018.12.030.

[68] Scott A. Kreher, Jae Young Kwon, and John R. Carlson. "The molecular basis of odor coding in the Drosophila larva". In: *Neuron* 46.3 (2005), pp. 445–456. DOI: 10.1016/j.neuron.2005.04.007.

[69] Li-Hui Cao, Bi-Yang Jing, Dong Yang, et al. "Distinct signaling of Drosophila chemoreceptors in olfactory sensory neurons". In: *Proceedings of the National Academy of Sciences* 113.7 (2016), E902–E911. DOI: 10.1073/pnas.1518329113.

[70] Rachel I. Wilson. "Early olfactory processing in Drosophila: mechanisms and principles". In: *Annual review of neuroscience* 36 (2013), pp. 217–241. DOI: 10.1146/annurev-neuro-062111-150533.

[71]   Reinhard F. Stocker. "Drosophila as a focus in olfactory research: mapping of olfactory sensilla by fine structure, odor specificity, odorant receptor expression, and central connectivity". In: *Microscopy research and technique* 55.5 (2001), pp. 284–296. DOI: 10.1002/jemt.1178.

[72]   Matthew E. Berck, Avinash Khandelwal, Lindsey Claus, et al. "The wiring diagram of a glomerular olfactory system". In: *Elife* 5 (2016), e14859. DOI: 10.7554/eLife.14859.

[73]   Jane Anne Horne, Carlie Langille, Sari McLin, et al. "A resource for the Drosophila antennal lobe provided by the connectome of glomerulus VA1v". In: *eLife* 7 (2018). Ed. by Liqun Luo, Eve Marder, and Moritz Helmstaedter, e37550. DOI: 10.7554/eLife.37550.

[74]   Rachel I. Wilson and Gilles Laurent. "Role of GABAergic inhibition in shaping odor-evoked spatiotemporal patterns in the Drosophila antennal lobe". In: *Journal of Neuroscience* 25.40 (2005), pp. 9069–9079. DOI: 10.1523/JNEUROSCI.2070-05.2005.

[75]   Nobuaki K. Tanaka, Hiromu Tanimoto, and Kei Ito. "Neuronal assemblies of the Drosophila mushroom body". In: *Journal of Comparative Neurology* 508.5 (2008), pp. 711–755.

[76]   Alexander S. Bates, Philipp Schlegel, Ruairi J.V. Roberts, et al. "Complete Connectomic Reconstruction of Olfactory Projection Neurons in the Fly Brain". In: *Current Biology* 30.16 (2020), 3183–3199.e6. DOI: 10.1016/j.cub.2020.06.042.

[77]   Mala Murthy, Ila Fiete, and Gilles Laurent. "Testing odor response stereotypy in the Drosophila mushroom body". In: *Neuron* 59.6 (2008), pp. 1009–1023. DOI: 10.1016/j.neuron.2008.07.040.

[78]   Glenn C. Turner, Maxim Bazhenov, and Gilles Laurent. "Olfactory representations by Drosophila mushroom body neurons". In: *Journal of neurophysiology* 99.2 (2008), pp. 734–746. DOI: 10.1152/jn.01283.2007.

[79]   Kyle S. Honegger, Robert A.A. Campbell, and Glenn C. Turner. "Cellular-resolution population imaging reveals robust sparse coding in the Drosophila mushroom body". In: *Journal of neuroscience* 31.33 (2011), pp. 11772–11785. DOI: 10.1523/JNEUROSCI.1099-11.2011.

[80]   Jill R. Crittenden, Efthimios M.C. Skoulakis, Kyung-An Han, et al. "Tripartite mushroom body architecture revealed by antigenic markers". In: *Learning & memory* 5.1 (1998), pp. 38–51. DOI: 10.1101/lm.5.1.38.

[81]   Troy Zars. "Behavioral functions of the insect mushroom bodies". In: *Current Opinion in Neurobiology* 10.6 (2000), pp. 790–795. DOI: 10.1016/S0959-4388(00)00147-1.

[82]   Michael J. Krashes, Alex C. Keene, Benjamin Leung, et al. "Sequential use of mushroom body neuron subsets during Drosophila odor memory processing". In: *Neuron* 53.1 (2007), pp. 103–115. DOI: 10.1016/j.neuron.2006.11.021.

[83]   Liria M. Masuda-Nakagawa, Nobuaki K. Tanaka, and Cahir J. O'Kane. "Stereotypic and random patterns of connectivity in the larval mushroom body calyx of Drosophila". In: *Proceedings of the National Academy of Sciences* 102.52 (2005), pp. 19027–19032. DOI: 10.1073/pnas.0509643102.

[84] James W. Truman, Jacquelyn Price, Rosa L. Miyares, et al. "Metamorphosis of memory circuits in Drosophila reveals a strategy for evolving a larval brain". In: *Elife* 12 (2023), e80594. DOI: `10.7554/eLife.80594`.

[85] Suewei Lin. "The making of the Drosophila mushroom body". In: *Frontiers in Physiology* 14 (2023), p. 5. DOI: `10.3389/fphys.2023.1091248`.

[86] André Fiala and Thomas Riemensperger. "Localization of a memory trace: Aversive associative olfactory learning and short-term memory in Drosophila". In: *Reference Module in Neuroscience and Biobehavioral Psychology*. Elsevier, 2017. DOI: `10.1016/B978-0-12-809324-5.04212-7`.

[87] Yoshinori Aso, Igor Siwanowicz, Lasse Bräcker, et al. "Specific dopaminergic neurons for the formation of labile aversive memory". In: *Current biology* 20.16 (2010), pp. 1445–1451. DOI: `10.1016/j.cub.2010.06.048`.

[88] Yoshinori Aso, Andrea Herb, Maite Ogueta, et al. "Three Dopamine Pathways Induce Aversive Odor Memories with Different Stability". In: *PLOS Genetics* 8.7 (2012), pp. 1–17. DOI: `10.1371/journal.pgen.1002768`.

[89] Scott Waddell. "Reinforcement signalling in Drosophila; dopamine does it all after all". In: *Current Opinion in Neurobiology* 23.3 (2013). Social and emotional neuroscience, pp. 324–329. DOI: `10.1016/j.conb.2013.01.005`.

[90] Christopher J. Burke, Wolf Huetteroth, David Owald, et al. "Layered reward signalling through octopamine and dopamine in Drosophila". In: *Nature* 492.7429 (2012), pp. 433–437. DOI: `10.1038/nature11614`.

[91] Thomas Riemensperger, Thomas Völler, Patrick Stock, et al. "Punishment prediction by dopaminergic neurons in Drosophila". In: *Current biology* 15.21 (2005), pp. 1953–1960. DOI: `10.1016/j.cub.2005.09.042`.

[92] Martin Schwaerzel, Maria Monastirioti, Henrike Scholz, et al. "Dopamine and octopamine differentiate between aversive and appetitive olfactory memories in Drosophila". In: *Journal of Neuroscience* 23.33 (2003), pp. 10495–10502. DOI: `10.1523/JNEUROSCI.23-33-10495.2003`.

[93] Yoshinori Aso and Gerald M. Rubin. "Dopaminergic neurons write and update memories with cell-type-specific rules". In: *Elife* 5 (2016), e16135. DOI: `/10.7554/eLife.16135`.

[94] Christian Schroll, Thomas Riemensperger, Daniel Bucher, et al. "Light-induced activation of distinct modulatory neurons triggers appetitive or aversive learning in Drosophila larvae". In: *Current biology* 16.17 (2006), pp. 1741–1747. DOI: `10.1016/j.cub.2006.07.023`.

[95] Suewei Lin, David Owald, Vikram Chandra, et al. "Neural correlates of water reward in thirsty Drosophila". In: *Nature neuroscience* 17.11 (2014), pp. 1536–1542. DOI: `10.1038/nn.3827`.

[96] Wolfram Schultz. "Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey". In: *Journal of neurophysiology* 56.5 (1986), pp. 1439–1461. DOI: `10.1152/jn.1986.56.5.1439`.

[97] Genela Morris, David Arkadir, Alon Nevet, et al. "Coincident but Distinct Messages of Midbrain Dopamine and Striatal Tonically Active Neurons". In: *Neuron* 43.1 (2004), pp. 133–143. DOI: `10.1016/J.NEURON.2004.06.012`.

[98] Takemasa Satoh, Sadamu Nakai, Tatsuo Sato, et al. "Correlated Coding of Motivation and Outcome of Decision by Dopamine Neurons". In: *Journal of Neuroscience* 23.30 (2003), pp. 9913–9923. DOI: 10.1523/JNEUROSCI.23-30-09913.2003.

[99] Jeremiah Y. Cohen, Sebastian Haesler, Linh Vong, et al. "Neuron-type-specific signals for reward and punishment in the ventral tegmental area". In: *Nature* 482.7383 (2012), pp. 85–88. DOI: 10.1038/nature10754.

[100] Seth M. Tomchik and Ronald L. Davis. "Dynamics of Learning-Related cAMP Signaling and Stimulus Integration in the Drosophila Olfactory Pathway". In: *Neuron* 64.4 (2009), pp. 510–521. DOI: 10.1016/j.neuron.2009.09.029.

[101] Young-Cho Kim, Hyun-Gwan Lee, and Kyung-An Han. "D1 dopamine receptor dDA1 is required in the mushroom body neurons for aversive and appetitive learning in Drosophila". In: *Journal of Neuroscience* 27.29 (2007), pp. 7640–7647. DOI: 10.1523/JNEUROSCI.1167-07.2007.

[102] Ronald L. Davis. "Olfactory memory formation in Drosophila: from molecular to systems neuroscience". In: *Annu. Rev. Neurosci.* 28 (2005), pp. 275–302. DOI: 10.1146/annurev.neuro.28.061604.135651.

[103] Emmanuel Perisse, David Owald, Oliver Barnstedt, et al. "Aversive learning and appetitive motivation toggle feed-forward inhibition in the Drosophila mushroom body". In: *Neuron* 90.5 (2016), pp. 1086–1099. DOI: 10.1016/j.neuron.2016.04.034.

[104] Anna-Maria Jürgensen, Panagiotis Sakagiannis, Michael Schleyer, et al. "Prediction error drives associative olfactory learning and conditioned behavior in a spiking model of Drosophila larva". In: *bioRxiv* (2022). DOI: 10.1101/2022.12.21.521372.

[105] Anna-Maria Jürgensen, Felix J. Schmitt, and Martin P. Nawrot. "Minimal circuit motifs for second-order conditioning in the insect mushroom body". In: *bioRxiv* (2023). DOI: 10.1101/2023.09.11.557174.

[106] Michael Schleyer, Alиće Weiglein, Juliane Thoener, et al. "Identification of dopaminergic neurons that can both establish associative memory and acutely terminate its behavioral expression". In: *Journal of Neuroscience* 40.31 (2020), pp. 5990–6006. DOI: 10.1523/JNEUROSCI.0290-20.2020.

[107] Ariane C. Boehm, Anja B. Friedrich, Sydney Hunt, et al. "A dopamine-gated learning circuit underpins reproductive state-dependent odor preference in Drosophila females". In: *eLife* 11 (2022), e77643. DOI: 10.7554/eLife.77643.

[108] Ramón Huerta and Thomas Nowotny. "Fast and robust learning by reinforcement signals: Explorations in the insect brain". In: *Neural computation* 21.8 (2009), pp. 2123–2151. DOI: 10.1162/neco.2009.03-08-733.

[109] Rinaldo Betkiewicz, Benjamin Lindner, and Martin P. Nawrot. "Circuit and cellular mechanisms facilitate the transformation from dense to sparse coding in the insect olfactory system". In: *Eneuro* 7.2 (2020), O.0305–18. DOI: 10.1523/ENEURO.0305-18.2020.

[110] Günther Palm. "Neural associative memories and sparse coding". In: *Neural Networks* 37 (2013), pp. 165–171. DOI: 10.1016/j.neunet.2012.08.013.

[111] Joshua P. Martin, Aaron Beyerlein, Andrew M. Dacks, et al. "The neurobiology of insect olfaction: sensory processing in a comparative context". In: *Progress in neurobiology* 95.3 (2011), pp. 427–447. DOI: 10.1016/j.pneurobio.2011.09.007.

[112] Sabine Krofczik, Randolf Menzel, and Martin P. Nawrot. "Rapid odor processing in the honeybee antennal lobe network". In: *Frontiers in computational neuroscience* 2 (2009), p. 9. DOI: 10.3389/neuro.10.009.2008.

[113] Debora Fuscà and Peter Kloppenburg. "Odor processing in the cockroach antennal lobe—the network components". In: *Cell and Tissue Research* 383 (2021), pp. 1–15. DOI: 10.1007/s00441-020-03387-3.

[114] Shawn R. Olsen and Rachel I. Wilson. "Lateral presynaptic inhibition mediates gain control in an olfactory circuit". In: *Nature* 452.7190 (2008), pp. 956–960. DOI: 10.1038/nature06864.

[115] Zhengchang Lei, Ke Chen, Hao Li, et al. "The GABA system regulates the sparse coding of odors in the mushroom bodies of Drosophila". In: *Biochemical and Biophysical Research Communications* 436.1 (June 2013), pp. 35–40. DOI: 10.1016/j.bbrc.2013.05.036.

[116] Heike Demmer and Peter Kloppenburg. "Intrinsic Membrane Properties and Inhibitory Synaptic Input of Kenyon Cells as Mechanisms for Sparse Coding?" In: *Journal of Neurophysiology* 102.3 (2009), pp. 1538–1550. DOI: 10.1152/jn.00183.2009.

[117] Jan Kropf and Wolfgang Rössler. "In-situ recording of ionic currents in projection neurons and Kenyon cells in the olfactory pathway of the honeybee". In: *PloS one* 13.1 (2018), e0191425. DOI: 10.1371/journal.pone.0191425.

[118] Aliće Weiglein, Juliane Thoener, Irina Feldbruegge, et al. "Aversive teaching signals from individual dopamine neurons in larval Drosophila show qualitative differences in their temporal "fingerprint"". In: *Journal of Comparative Neurology* 529.7 (2021), pp. 1553–1570. DOI: /10.1002/cne.25037.

[119] Juliane Thoener, Aliće Weiglein, Bertram Gerber, et al. "Optogenetically induced reward and 'frustration'memory in larval Drosophila melanogaster". In: *Journal of Experimental Biology* 225.16 (2022), jeb244565. DOI: 10.1242/jeb.244565.

[120] Wolfram Schultz, Paul Apicella, and Tomas Ljungberg. "Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task". In: *Journal of neuroscience* 13.3 (1993), pp. 900–913. DOI: 10.1523/JNEUROSCI.13-03-00900.1993.

[121] Wolfram Schultz, Peter Dayan, and P. Read Montague. "Getting formal with dopamine and reward". In: *Science* 275.5306 (1997), pp. 1593–1599. DOI: 10.1126/science.275.5306.1593.

[122] Kareem A. Zaghloul, Justin A. Blanco, Christoph T. Weidemann, et al. "Human substantia nigra neurons encode unexpected financial rewards". In: *Science* 323.5920 (2009), pp. 1496–1499. DOI: 10.1126/science.1167342.

[123] Wei-Xing Pan, Robert Schmidt, Jeffery R. Wickens, et al. "Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network". In: *Journal of Neuroscience* 25.26 (2005), pp. 6235–6242. DOI: /10.1523/JNEUROSCI.1478-05.2005.

[124] Daichi Yamada, Daniel Bushey, Feng Li, et al. "Hierarchical architecture of dopaminergic circuits enables second-order conditioning in Drosophila". In: *Elife* 12 (2023), e79042. DOI: 10.7554/eLife.79042.

[125] Kristina V. Dylla, Georg Raiser, C. Giovanni Galizia, et al. "Trace conditioning in Drosophila induces associative plasticity in mushroom body Kenyon cells and dopaminergic neurons". In: *Frontiers in neural circuits* 11 (2017), p. 42. DOI: 10.3389/fncir.2017.00042.

[126] El Yazid Rachad. "Neural circuit plasticity underlying learning and memory in Drosophila melanogaster: from synaptic connections to behavior". PhD thesis. University of Göttingen, 2023. DOI: 10.53846/goediss-9845.

[127] Paolo Arena, Luca Patané, Vincenzo Stornanti, et al. "Modeling the insect mushroom bodies: Application to a delayed match-to-sample task". In: *Neural Networks* 41 (2013), pp. 202–211. DOI: 10.1016/j.neunet.2012.11.013.

[128] Faramarz Faghihi, Ahmed A. Moustafa, Ralf Heinrich, et al. "A computational model of conditioning inspired by Drosophila olfactory system". In: *Neural Networks* 87 (2017), pp. 96–108. DOI: 10.1016/j.neunet.2016.11.002.

[129] Magdalena Springer and Martin P. Nawrot. "A mechanistic model for reward prediction and extinction learning in the fruit fly". In: *Eneuro* 8.3 (2021). DOI: 10.1523/ENEURO.0549-20.2021.

[130] Jan Wessnitzer, Joanna M. Young, J. Douglas Armstrong, et al. "A model of non-elemental olfactory learning in Drosophila". In: *Journal of computational neuroscience* 32 (2012), pp. 197–212. DOI: 10.1007/s10827-011-0348-6.

[131] Fei Peng and Lars Chittka. "A simple computational model of the bee mushroom body can explain seemingly complex forms of olfactory learning and memory". In: *Current Biology* 27.2 (2017), pp. 224–230. DOI: 10.1016/j.cub.2016.10.054.

[132] Linnie Jiang and Ashok Litwin-Kumar. "Models of heterogeneous dopamine signaling in an insect learning and memory center". In: *PLoS Computational Biology* 17.8 (2021), e1009205. DOI: 10.1371/journal.pcbi.1009205.

[133] Chang Zhao, Yves F. Widmer, Sören Diegelmann, et al. "Predictive olfactory learning in Drosophila". In: *Scientific reports* 11.1 (2021), pp. 1–17. DOI: 10.1038/s41598-021-85841-y.

[134] James E.M. Bennett, Andrew Philippides, and Thomas Nowotny. "Learning with reinforcement prediction errors in a model of the Drosophila mushroom body". In: *Nature Communications* 12.1 (2021). DOI: 10.1038/s41467-021-22592-4.

[135] Evripidis Gkanias, Li Y. McCurdy, Michael N. Nitabach, et al. "An incentive circuit for memory dynamics in the mushroom body of Drosophila melanogaster". In: *Elife* 11 (2022), e75611. DOI: 10.7554/eLife.75611.

[136] Elizabeth C. Marin, Laurin Büld, Maria Theiss, et al. "Connectomics analysis reveals first-, second-, and third-order thermosensory and hygrosensory neurons in the adult Drosophila brain". In: *Current Biology* 30.16 (2020), pp. 3167–3182. DOI: `10.1016/j.cub.2020.06.028`.

[137] Nicholas J. Strausfeld and Yongsheng Li. "Organization of olfactory and multimodal afferent neurons supplying the calyx and pedunculus of the cockroach mushroom bodies". In: *Journal of Comparative Neurology* 409.4 (1999), pp. 603–625. DOI: `10.1002/(SICI)1096-9861(19990712)409:4<603::AID-CNE7>3.0.CO;2-P`.

[138] C. Giovanni Galizia and Wolfgang Rössler. "Parallel Olfactory Systems in Insects: Anatomy and Function". In: *Annual Review of Entomology* 55.1 (2010), pp. 399–420. DOI: `10.1146/annurev-ento-112408-085442`.

[139] Hiroshi Nishino, Shingo Yamashita, Yoshiyuki Yamazaki, et al. "Projection neurons originating from thermo-and hygrosensory glomeruli in the antennal lobe of the cockroach". In: *Journal of Comparative Neurology* 455.1 (2003), pp. 40–55. DOI: `10.1002/cne.10450`.

[140] Cansu Arican, Felix J. Schmitt, Wolfgang Rössler, et al. "The mushroom body output encodes behavioral decision during sensory-motor transformation". In: *bioRxiv* (2022), pp. 2022–09. DOI: `10.1101/2022.09.14.507924`.

[141] Yongsheng Li and Nicholas J. Strausfeld. "Morphology and sensory modality of mushroom body extrinsic neurons in the brain of the cockroach, Periplaneta americana". In: *Journal of Comparative Neurology* 387.4 (1997), pp. 631–650. DOI: `10.1002/(SICI)1096-9861(19971103)387:4<631::AID-CNE9>3.0.CO;2-3`.

[142] Ryuichi Okada, Midori Sakura, and Makoto Mizunami. "Distribution of dendrites of descending neurons and its implications for the basic organization of the cockroach brain". In: *Journal of Comparative Neurology* 458.2 (2003), pp. 158–174. DOI: `10.1002/cne.10580`.

[143] Brad K. Hulse, Hannah Haberkern, Romain Franconville, et al. "A connectome of the Drosophila central complex reveals network motifs suitable for flexible navigation and context-dependent action selection". In: *eLife* 10 (2021), e66039. DOI: `10.7554/eLife.66039`.

[144] Robert A. Rescorla and Allan R. Wagner. "A theory of Pavlovian conditioning : Variations in the effectiveness of reinforcement and non- reinforcement". In: *Classical Conditioning 2: Current Theory and Research.* Ed. by Abraham H. Black and William F. Prokasy. Appelton-century-Crofts, 1972, pp. 64–99.

[145] David Sterratt, Bruce Graham, Andrew Gillies, et al. *Principles of computational modelling in neuroscience.* Cambridge University Press, 2011. DOI: `10.1017/CBO9780511975899.009`.

[146] Bertram Gerber, Roland Biernacki, and Jeannette Thum. "Odor-taste learning assays in Drosophila larvae". In: *Cold Spring Harbor Protocols* 2013.3 (2013), pdb–prot071639. DOI: `10.1101/pdb.prot071639`.

[147] Denise Weber, Vincent Richter, Astrid Rohwedder, et al. "The analysis of aversive olfactory-taste learning and memory in Drosophila larvae". In: *Cold Spring Harbor Protocols* (2022). DOI: `10.1101/pdb.prot108050`.

[148] Aliće Weiglein, Florian Gerstner, Nino Mancini, et al. "One-trial learning in larval Drosophila". In: *Learning & Memory* 26.4 (2019), pp. 109–120. DOI: 10.1101/lm.049106.118.

[149] Daeyeon Kim, Mar Alvarez, Laura M. Lechuga, et al. "Species-specific modulation of food-search behavior by respiration and chemosensation in Drosophila larvae". In: *Elife* 6 (2017), pp. 1–23. DOI: 10.7554/eLife.27057.

[150] "Active sampling and decision making in Drosophila chemotaxis". In: *Nat. Commun.* 2.1 (2011), pp. 410–441. DOI: 10.1038/ncomms1455.

[151] Balazs Szigeti, Ajinkya Deogade, and Barbara Webb. "Searching for motifs in the behaviour of larval Drosophila melanogaster and Caenorhabditis elegans reveals continuity between behavioural states". In: *J. R. Soc. Interface* 12.113 (2015). DOI: 10.1098/rsif.2015.0899.

[152] Gemma Slater, Peter Levy, Andrew Chan, et al. "A central neural pathway controlling odor tracking in drosophila". In: *J. Neurosci.* 35.5 (2015), pp. 1831–1848. DOI: 10.1523/JNEUROSCI.2331-14.2015.

[153] Michael Schleyer, Daisuke Miura, Teiichi Tanimura, et al. "Learning the specific quality of taste reinforcement in larval Drosophila". In: *Elife* 4 (2015). DOI: 10.7554/eLife.04711.

[154] Tim Tully and William G. Quinn. "Classical conditioning and retention in normal and mutant Drosophila melanogaster". In: *Journal of Comparative Physiology A* 157.2 (1985), pp. 263–277. DOI: 10.1007/BF01350033.

[155] Satoshi Murakami, Chuntao Dan, Brendan Zagaeski, et al. "Optimizing Drosophila olfactory learning with a semi-automated training device". In: *Journal of Neuroscience Methods* 188.2 (2010), pp. 195–204. DOI: 10.1016/j.jneumeth.2010.02.007.

[156] Bruce L. Tempel, Nancy Bonini, Douglas R. Dawson, et al. "Reward learning in normal and mutant Drosophila". In: *Proceedings of the National Academy of Sciences* 80.5 (1983), pp. 1482–1486. DOI: 10.1073/pnas.80.5.1482.

[157] Marie-Ange Chabaud, Jean-Marc Devaud, Minh-Hà Pham-Delègue, et al. "Olfactory conditioning of proboscis activity in Drosophila melanogaster". In: *Journal of Comparative Physiology A* 192 (2006), pp. 1335–1348. DOI: 10.1007/s00359-006-0160-3.

[158] Alex Davies, Matthieu Louis, and Barbara Webb. "A model of Drosophila larva chemotaxis". In: *PLoS computational biology* 11.11 (2015), e1004606. DOI: 10.1371/journal.pcbi.1004606.

[159] Tianqi Wei, Adam Stokes, and Barbara Webb. "A Soft Pneumatic Maggot Robot". In: *Conference on Biomimetic and Biohybrid Systems*. Ed. by Nathan F. Lepora, Anna Mura, Michael Mangan, et al. Springer International Publishing, 2016, pp. 375–386.

[160] Jan Wessnitzer and Barbara Webb. "Multimodal sensory integration in insects—towards insect brain control architectures". In: *Bioinspiration & biomimetics* 1.3 (2006), p. 63. DOI: 10.1088/1748-3182/1/3/001.

[161] Dylan Ross, Konstantinos Lagogiannis, and Barbara Webb. "A model of larval biomechanics reveals exploitable passive properties for efficient locomotion". In: *Biomimetic and Biohybrid Systems: 4th International Conference, Living Machines.* Springer. 2015, pp. 1–12. DOI: `10.1007/978-3-319-22979-9_1`.

[162] Raphael Cohn, Ianessa Morantte, and Vanessa Ruta. "Coordinated and Compartmentalized Neuromodulation Shapes Sensory Processing in Drosophila". In: *Cell* 163.7 (2015), pp. 1742–1755. DOI: `10.1016/j.cell.2015.11.019`.

[163] KP Siju, Vilim Štih, Sophie Aimon, et al. "Valence and State-Dependent Population Coding in Dopaminergic Neurons in the Fly Mushroom Body". In: *Current Biology* 30.11 (2020), 2104–2115.e4. DOI: `10.1016/j.cub.2020.04.037`.

[164] Chang-Hui Tsao, Chien-Chun Chen, Chen-Han Lin, et al. "Drosophila mushroom bodies integrate hunger and satiety signals to control innate food-seeking behavior". In: *eLife* 7 (2018), e35264. DOI: `10.7554/eLife.35264`.

[165] Dennis Mathew, Carlotta Martelli, Elizabeth Kelley-Swift, et al. "Functional diversity among sensory receptors in a Drosophila olfactory circuit". In: *Proceedings of the National Academy of Sciences* 110.23 (2013), E2134–E2143. DOI: `10.1073/pnas.130697611`.

[166] Scott A. Kreher, Dennis Mathew, Junhyong Kim, et al. "Translation of sensory input into behavioral output via an olfactory system". In: *Neuron* 59.1 (2008), pp. 110–124. DOI: `10.1016/j.neuron.2008.06.010`.

[167] Timo Saumweber, Jana Husse, and Bertram Gerber. "Innate attractiveness and associative learnability of odors can be dissociated in larval Drosophila". In: *Chemical senses* 36.3 (2011), pp. 223–235. DOI: `10.1093/chemse/bjq128`.

[168] Matthew Cobb. "What and how do maggots smell?" In: *Biological Reviews* 74.4 (1999), pp. 425–459. DOI: `10.1017/S0006323199005393`.

[169] Michael-John Dolan, Ghislain Belliart-Guérin, Alexander Shakeel Bates, et al. "Communication from Learned to Innate Olfactory Processing Centers Is Required for Memory Retrieval in Drosophila". In: *Neuron* 100.3 (2018), 651–668.e8. DOI: `10.1016/j.neuron.2018.08.037`.

[170] Greg S.B. Suh, Allan M. Wong, Anne C. Hergarden, et al. "A single population of olfactory sensory neurons mediates an innate avoidance behaviour in Drosophila". In: *Nature* 431.7010 (2004), pp. 854–859. DOI: `10.1038/nature02980`.

[171] Gertrud Heimbeck, Véronique Bugnon, Nanaë Gendre, et al. "A central neural circuit for experience-independent olfactory and courtship behavior in Drosophila melanogaster". In: *Proceedings of the National Academy of Sciences* 98.26 (2001), pp. 15336–15341. DOI: `10.1073/pnas.011314898`.

[172] Moshe Parnas, Andrew C. Lin, Wolf Huetteroth, et al. "Odor Discrimination in Drosophila: From Neural Population Codes to Behavior". In: *Neuron* 79.5 (2013), pp. 932–944. DOI: `10.1016/j.neuron.2013.08.006`.

[173] Sudeshna Das Chakraborty and Silke Sachse. "Olfactory processing in the lateral horn of Drosophila". In: *Cell and tissue research* 383 (2021), pp. 113–123. DOI: `10.1007/s00441-020-03392-6`.

[174] Elizabeth C. Marin, Gregory S.X.E. Jefferis, Takaki Komiyama, et al. "Representation of the glomerular olfactory map in the Drosophila brain". In: *Cell* 109.2 (2002), pp. 243–255. DOI: `10.1016/S0092-8674(02)00700-6`.

[175] Gregory S.X.E. Jefferis, Christopher J. Potter, Alexander M. Chan, et al. "Comprehensive maps of Drosophila higher olfactory centers: spatially segregated fruit and pheromone representation". In: *Cell* 128.6 (2007), pp. 1187–1203. DOI: 10.1016/j.cell.2007.01.040.

[176] Andrew C. Amin Hoger adn Lin. "Neuronal mechanisms underlying innate and learned olfactory processing in Drosophila". In: *Current Opinion in Insect Science* 36 (2019), pp. 9–17. DOI: 10.1016/j.cois.2019.06.003.

[177] Nobuaki K. Tanaka, Takeshi Awasaki, Takashi Shimada, et al. "Integration of chemosensory pathways in the Drosophila second-order olfactory centers". In: *Current biology* 14.6 (2004), pp. 449–457. DOI: 10.1016/j.cub.2004.03.006.

[178] Peter Clyne, Alan Grant, Robert O'connell, et al. "Odorant response of individual sensilla on the Drosophila antenna". In: *Invertebrate Neuroscience* 3 (1997), pp. 127–135. DOI: 10.1007/BF02480367.

[179] Amina Kurtovic, Alexandre Widmer, and Barry J. Dickson. "A single class of olfactory neurons mediates behavioural responses to a Drosophila sex pheromone". In: *Nature* 446.7135 (2007), pp. 542–546. DOI: 10.1038/nature05672.

[180] Markus Knaden, Antonia Strutz, Jawaid Ahsan, et al. "Spatial representation of odorant valence in an insect brain". In: *Cell reports* 1.4 (2012), pp. 392–399. DOI: 10.1016/j.celrep.2012.03.002.

[181] Bohan Zhao, Jiameng Sun, Xuchen Zhang, et al. "Long-term memory is formed immediately without the need for protein synthesis-dependent consolidation in Drosophila". In: *Nature communications* 10.1 (2019), p. 4550. DOI: 10.1038/s41467-019-12436-7.

[182] Lasse B. Bräcker, K.P. Siju, Nélia Varela, et al. "Essential Role of the Mushroom Body in Context-Dependent CO_2 Avoidance in Drosophila". In: *Current Biology* 23.13 (2013), pp. 1228–1234. DOI: 10.1016/j.cub.2013.05.029.

[183] Laurence P.C. Lewis, K.P. Siju, Yoshinori Aso, et al. "A higher brain circuit for immediate integration of conflicting sensory information in Drosophila". In: *Current biology* 25.17 (2015), pp. 2203–2214. DOI: 10.1016/j.cub.2015.07.015.

[184] Cheng Huang, Junjie Luo, Seung Je Woo, et al. "Dopamine signals integrate innate and learnt valences to regulate memory dynamics". In: *Research Square* (2022). DOI: 10.21203/rs.3.rs-1915648/v1.

[185] Qian Li and Stephen D. Liberles. "Aversion and Attraction through Olfaction". In: *Current Biology* 25.3 (2015), R120–R129. DOI: 10.1016/j.cub.2014.11.044.

[186] Elane Fishilevich and Leslie B. Vosshall. "Genetic and functional subdivision of the Drosophila antennal lobe". In: *Current Biology* 15.17 (2005), pp. 1548–1553. DOI: 10.1016/j.cub.2005.07.066.

[187] Sabine Scherer, Reinhard F. Stocker, and Bertram Gerber. "Olfactory learning in individually assayed Drosophila larvae". In: *Learning & Memory* 10.3 (2003), pp. 217–225. DOI: 10.1101/lm.57903.

[188] Michael Schleyer, Samuel F. Reid, Evren Pamir, et al. "The impact of odor–reward memory on chemotaxis in larval Drosophila". In: *Learning & Memory* 22.5 (2015), pp. 267–277. DOI: 10.1101/lm.037978.114.

[189] Michael Schleyer, Timo Saumweber, Wiebke Nahrendorf, et al. "A behavior-based circuit model of how outcome expectations organize learned behavior in larval Drosophila". In: *Learning & Memory* 18.10 (2011), pp. 639–653. DOI: 10.1101/lm.2163411.

[190] Emmanouil Paisios, Annabell Rjosk, Evren Pamir, et al. "Common microbehavioral "footprint" of two distinct classes of conditioned aversion". In: *Learning & Memory* 24.5 (2017), pp. 191–198. DOI: 10.1101/lm.045062.117.

[191] Johannes Felsenberg, Oliver Barnstedt, Paola Cognigni, et al. "Re-evaluation of learned information in Drosophila". In: *Nature* 544.7649 (2017), pp. 240–244. DOI: 10.1038/nature21716.

[192] Johannes Felsenberg, Pedro F. Jacob, Thomas Walker, et al. "Integration of Parallel Opposing Memories Underlies Memory Extinction". In: 175 (2018), 709–722.e15. DOI: 10.1016/J.CELL.2018.08.021.

[193] Lingling Wang, Qi Yang, Binyan Lu, et al. "A behavioral paradigm to study the persistence of reward memory extinction in Drosophila". In: *Journal of genetics and genomics* 46.12 (2019), pp. 599–601. DOI: 10.1016/j.jgg.2019.11.001.

[194] Yukinori Hirano, Kunio Ihara, Tomoko Masuda, et al. "Shifting transcriptional machinery is required for long-term memory maintenance and modification in Drosophila mushroom bodies". In: *Nature communications* 7.1 (2016), p. 13471. DOI: 10.1038/ncomms13471.

[195] Christopher J. Tabone and J. Steven de Belle. "Second-order conditioning in Drosophila". In: *Learning & Memory* 18.4 (2011), pp. 250–253. DOI: 10.1101/lm.2035411.

[196] Gaurav Das, Martın Klappenbach, Eleftheria Vrontou, et al. "Drosophila learn opposing components of a compound food stimulus". In: *Current biology* 24.15 (2014), pp. 1723–1730. DOI: 10.1016/j.cub.2014.05.078.

[197] Michael-John Dolan, Shahar Frechter, Alexander Shakeel Bates, et al. "Neurogenetic dissection of the Drosophila lateral horn reveals major outputs, diverse behavioural functions, and interactions with the mushroom body". In: *Elife* 8 (2019), e43079. DOI: 10.7554/eLife.43079.

[198] Mehrab N. Modi, Yichun Shuai, and Glenn C. Turner. "The Drosophila mushroom body: from architecture to algorithm in a learning circuit". In: *Annual review of neuroscience* 43 (2020), pp. 465–484. DOI: 10.1146/annurev-neuro-080317-0621333.

7   Appendix

Contribution Statement

The introduction and discussion of this thesis were written by Anna-Maria Jürgensen (AMJ). All data in chapters 2.1, 3.1, and 3.2 was collected by AMJ, except the neuromorphic chip experiment, which Afshin Khalili conducted. The locomotory simulation in chapter 3.1 and 4.1 were conducted by Panagiotis Sakagiannis. The MB simulation in chapter 4.1 was done by AMJ. All data analyses in chapters chapter 2.1, 3.1 and 3.2 were performed by AMJ. The hypersphere sampling in chapter 3.2 was done by Felix Schmitt and AMJ. The grid search in chapter 3.2 was largely performed by Felix Schmitt. All figures in chapters chapter 2.1, chapter 3.1, and 3.2 were created by AMJ. The MB learning figure in chapter 4.1 was created by AMJ. Most text in chapters chapter 2.1, 3.1, and 3.2 was written by AMJ. Martin Nawrot, Elisabetta Chicca, and Giacomo Indiveri wrote parts of the text in chapter 2.1. The manuscript was revised by Afshin Khalili, Martin Nawrot, Elisabetta Chicca, and Giacomo Indiveri. Martin Nawrot wrote parts of the text in chapter 3.1. The manuscript was revised by Michael Schleyer, Bertram Gerber, and Martin Nawrot. The manuscript in chapter 3.1 was revised by Felix Schmitt and Martin Nawrot. Most text in chapter 4.1 was written by Panagiotis Sakagiannis. AMJ and Martin Nawrot wrote parts of the text. The manuscript was revised by Panagiotis Sakagiannis, Martin Nawrot, and AMJ.

Data accessibility statement

All data is stored on internal lab servers. Data belonging to published manuscripts is available upon request to the corresponding author.

Acknowledgements