



A spiking neural program for sensorimotor control during foraging in flying insects

Hannes Rapp^a and Martin Paul Nawrot^{a,1}

^aComputational Systems Neuroscience, Institute of Zoology, University of Cologne, Cologne 50674, Germany

Edited by L. B. Vosshall, The Rockefeller University, New York, NY, and approved September 22, 2020 (received for review May 16, 2020)

Foraging is a vital behavioral task for living organisms. Behavioral strategies and abstract mathematical models thereof have been described in detail for various species. To explore the link between underlying neural circuits and computational principles, we present how a biologically detailed neural circuit model of the insect mushroom body implements sensory processing, learning, and motor control. We focus on cast and surge strategies employed by flying insects when foraging within turbulent odor plumes. Using a spike-based plasticity rule, the model rapidly learns to associate individual olfactory sensory cues paired with food in a classical conditioning paradigm. We show that, without retraining, the system dynamically recalls memories to detect relevant cues in complex sensory scenes. Accumulation of this sensory evidence on short time scales generates cast-and-surge motor commands. Our generic systems approach predicts that population sparseness facilitates learning, while temporal sparseness is required for dynamic memory recall and precise behavioral control. Our work successfully combines biological computational principles with spike-based machine learning. It shows how knowledge transfer from static to arbitrary complex dynamic conditions can be achieved by foraging insects and may serve as inspiration for agent-based machine learning.

mushroom body | sparse coding | navigation | artificial intelligence | spiking neural network

Navigating toward a food source during foraging requires dynamical sensory processing, accumulation of sensory evidence, and appropriate high-level motor control. Navigation based on an animal's olfactory sense is a challenging task due to the complex spatiotemporal landscape of odor molecules. A core aspect of foraging is the acquisition of sensory cue samples in the natural environment, where odor concentrations vary rapidly and steeply across space. Experimental access to the neural substrate is challenging in freely behaving insects. Biologically realistic models thus play a key role in investigating the relevant computational mechanisms. Consequently, recent efforts at understanding foraging behavior have focused on identifying viable computational strategies for making navigational decisions (1).

An odor plume is often considered a volume wherein odor concentration is generally above some behavioral threshold. At macroscopic scales and in a natural environment, however, plumes are turbulent (2, 3). In turbulent conditions, a plume breaks up into complex and intermittent filamentous structures that are interspersed with clean air pockets or below behavioral-threshold concentration patches (4, 5). The dispersing filaments form the cone-like shape of the macroscopic plume, where the origin of the cone yields the position of the odor source. When entering the cone, flying insects encounter odor filaments as discrete, short-lived sensory events in time.

Several features have been derived from the statistics of an odor plume that provide information regarding the location of the odor source (3, 4). The mean concentration varies smoothly in lateral and longitudinal directions of time-averaged (and laminar) plumes. However, for behavioral strategies, animals cannot afford the time it takes to obtain stable macroscopic estimates of mean concentrations (2). Hildebrand and colleagues

(6) proposed the time interval between odor encounters as an informative olfactory feature, while ref. 3 suggested intermittency, the probability of the odor concentration being above some behavioral threshold, as the relevant feature. However, similarly to estimating mean concentration, acquiring a sufficient number of samples for stable estimates of these quantities exceeds the time typically used to form behavioral decisions (2). Hence, obtaining time-averaged quantities is not an optimal strategy to guide navigational decisions, as concluded by ref. 7.

Most animals perform searches at large distances from the odor source, where the intermittency of plumes poses a more severe problem, as available sensory cues become more sparse in space and time. Thus, strategies that exploit brief, localized sensory cues for navigation have been studied by several groups. One strategy for medium- and long-range navigation that has consistently been observed across species of flying insects emerges from a sequence of chained sensorimotor reflexes: casting and surging (8). Encountering a whiff of odor triggers an upwind surge behavior, during which the insect travels parallel to the wind direction. After losing track of the plume, it evokes a cross-wind cast behavior, in which a flight path perpendicular to the direction of airflow is executed. Performing repeated casts by U-turning allows the insect to reenter and locate the plume in order to trigger the next upwind surge (8–10). As the subject approaches the source, it increasingly makes use of visual cues for navigation as the plume narrows down (8).

A number of studies have proposed abstract mathematical models for optimal search algorithms that assumed different

Significance

Living organisms demonstrate remarkable abilities in mastering problems imposed by complex and dynamic environments, and they can generalize their experience in order to rapidly adapt behavior. This paper demonstrates the benefits of using biological spiking neural networks, sparse computations, and local learning rules. It highlights the functional roles of temporal- and population-sparse coding for rapid associative learning, precise memory recall, and transformation into navigational output. We show how memory formation generalizes to perform precise memory recall under dynamic, nonstationary conditions, giving rise to nontrivial foraging behavior in a complex natural environment. Results suggest how principles of biological computation could benefit agent-based machine learning to deal with nonstationary scenarios.

Author contributions: H.R. and M.P.N. designed research; H.R. performed research; H.R. analyzed data; and H.R. and M.P.N. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: martin.nawrot@uni-koeln.de.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2009821117/-DCSupplemental>.

First published October 29, 2020.

types of relevant navigational cues. The infotaxis method proposed in ref. 11 depends on extensive memory and priors regarding a plume's structure. To the contrary, in ref. 8, only local cues are used. A standard algorithm for navigational problems in robotics is simultaneous localization and mapping (SLAM), which has been used in ref. 12 to study olfactory navigation in bumblebees. An algorithm that works without space perception has been proposed by ref. 13 using a standardized projection of the probability of source position and minimization of a free energy along the trajectory. Finally, the work of ref. 10 compares several models and shows that it is difficult to discriminate between different models based on behavioral responses. A recent work (7) using information-theoretic analysis shows that plumes contain both spatial and temporal information about the source's position.

While all of these previous mathematical methods for olfactory search algorithms have proven to successfully solve this task based on the respective assumptions, they share the same major drawback: None of them uses the computational substrate of the brain, spiking neurons and networks thereof. Instead, all methods make heavy use of symbolic math and advanced mathematical concepts that are not available to the biological brain. It is further unclear how and to what extent these methods could be implemented or learned by the nervous system. Additionally, the problem of navigation and foraging is often considered as an isolated task, independent from sensory processing.

Our approach distills recent experimental results to formulate a biologically plausible and detailed spiking neural-network model supporting adaptive foraging behavior. We thereby take advantage of the rapidly accumulating knowledge regarding the anatomy (e.g., refs. 14–16) and neurophysiology (e.g., refs. 17–19) of insect olfaction and basic computational features (20, 21). We follow the idea of compositionality, a widely used concept in mathematics, semantics, and linguistics. According to this principle, the meaning of a complex expression is a function of the meanings of its constituent expressions [Frege principle (22)]. In the present context of foraging and navigation, this means dynamically recombining memories of individual sensory cues present within a plume.

Results

We approached the problem of foraging by decomposition into four components: first, sensory processing with temporal sparse and population sparse coding in the mushroom body (MB); second, associative learning for assigning a valence to individual odor identities; third, the time-dependent detection of valenced cues resulting from encounters of discrete odor filaments to provide an ongoing and robust estimate of sensory cue evidence; and fourth, the translation into online motor command signals to drive appropriate behavior.

For sensory processing, we used a three-layer spiking neural network model of the insect olfactory pathway (Fig. 1). The generic blueprint of the insect olfactory system is homologous across species and comprises three successive processing stages (see *Materials and Methods* for details): the periphery with olfactory receptor neurons (ORNs), the antennal lobe (AL), and the MB. Excitatory feed-forward connections across layers from ORNs to projection neurons (PNs), from ORNs to local interneurons (LNs), and from PNs to the MB Kenyon cells (KCs) are fixed. Lateral inhibition within the AL uses fixed synaptic weights from LNs to PNs. For neuron numbers and their connectivity patterns, we here rely on the adult *Drosophila melanogaster*, where anatomical knowledge is most complete (14, 23, 24). A single MB output neuron (MBON) receives input from all KCs and plasticity at the synapses between KCs, and the MBON enables associative learning (25, 26).

Sparse Coding in Space and Time. The olfactory system transforms a dense olfactory code in the AL into a sparse stimulus code at

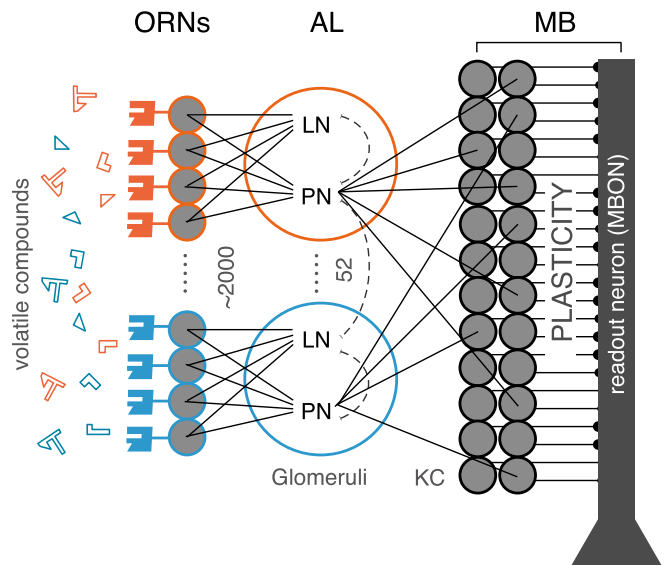


Fig. 1. Spiking network model of the insect olfactory system. Olfactory receptor neurons ($N = 2,080$) at the antenna bind and respond to volatile odorant compounds. ORNs expressing the same (1 of 52 different) genetic receptor type converge onto the same glomerulus in the AL. Each of the 52 glomeruli comprises one PN and one LN. Each LN forms lateral inhibitory connections with all PNs. PNs randomly connect to a large population of KCs ($N = 2,000$), where each KC receives input from, on average, 6 random PNs. All KCs project to a single MBON via plastic synapses.

the MB level. In the large population of KCs, a specific odor stimulus is represented by only a small fraction of all KCs (population sparseness), and each stimulus-activated KC responds with only a single or very few action potentials (temporal sparseness). In our model, temporal sparseness is achieved through the cellular mechanisms of spike-frequency adaptation (SFA; refs. 21, 27, and 28) implemented at two levels of the system.

ORNs show clear stimulus-response adaptation that has been attributed to the spike-generating mechanism (29). Based on this experimental evidence, we introduced a slow and weak SFA conductance in our model ORNs (*Materials and Methods*). At the level of the MB, KCs have been shown to express strong SFA-mediated channels (18). This is matched by the SFA parameters of our model KCs (*Materials and Methods*; ref. 21). As an effect of cellular adaptation in ORNs and KCs, odor stimulation (Fig. 2A) results in temporally precise and adaptive responses across all layers of the network (Fig. 2B). The effect of SFA implemented in ORNs is transitive and, thus, carries over to the postsynaptic PN and LN populations in agreement with experimental observations across species (30–33). SFA has the additional advantage of inherently reducing the variability of spiking output (28, 34).

In the KC population, the background firing rate is very low (~ 0.4 Hz). This is partially due to the outward SFA conductance and in agreement with experimental results (17). The KC population response is highly transitive where individual responding cells generate only a single or very few response spikes shortly after stimulus onset. This is in good qualitative and quantitative agreement with the temporal sparse KC spike responses measured in various species (17, 30, 35).

Population-sparse stimulus encoding at the level of KCs is supported by two major factors. First, the sparse divergent-convergent connectivity between the PNs and the 20-times-larger population of KCs is the anatomical basis for sparse odor representation (15, 20, 21, 36) (*SI Appendix, Table S1*). Second, lateral inhibition mediated by the LNs in the AL (37) facilitates decorrelation of odor representations (37) and contributes to population sparseness (21). The sparse code in the

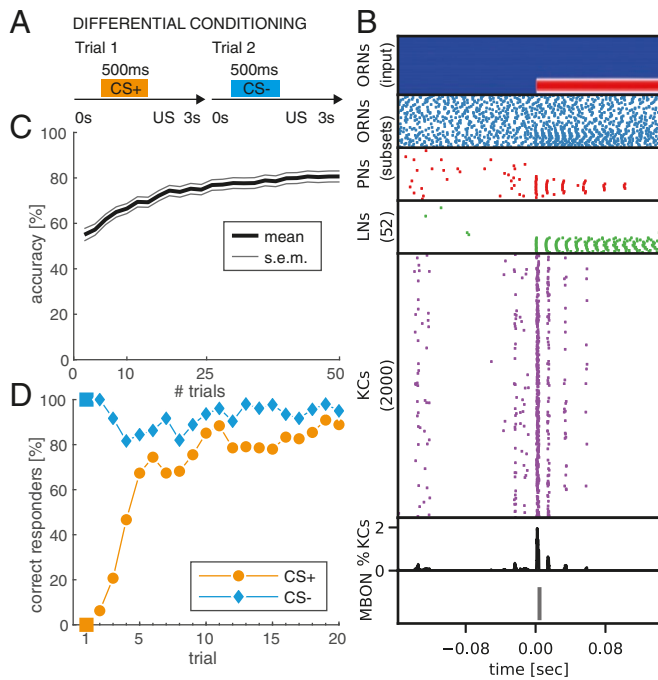


Fig. 2. Rapid associative learning expressed in neuronal plasticity and CR behavior. (A) Sketch of differential conditioning protocol. In appetitive trials, a first sensory cue (conditioned stimulus; CS+; orange) is paired with a reward (unconditioned stimulus; US). In aversive trials, a second sensory cue (CS-; blue) is paired with a punishment. Both trial types are presented randomized within blocks (*Materials and Methods*). (B) Sensory input and neuronal responses across all four circuit layers (ORN, AL, MB, and MBON) in response to a CS+ odor presentation during the 10th training trial. Stimulus onset is at $t = 0$ s. From top to bottom: Model input is provided through independent noise-current injection into the ORNs. The stimulus-induced input currents are clearly visible (hot colors) on top of the background noise for the subset of ORNs that are sensitive to the CS+ odor (stimulus profile). Stimulus response is clearly visible by an increase in the spiking activity across all neuron populations. For ORNs (blue) and PNs (red), relevant subsets of 60 and 35 neurons are shown. The population of 2,000 KCs (magenta) shows a temporal and spatial sparse odor response. Only 2% of all KCs are activated during a brief transient response following stimulus onset (black histogram). The MBON generates a single action potential in response to cue onset, which is the correct learned response to the CS+ odor. (C) Learning performance of the MBON across $N = 100$ independent models as a function of the number of training trials. In any given trial, the MBON response was correct if exactly one action potential was generated during CS+ presentation or if no action potential was generated during CS- presentation. (D) The behavioral learning curve expresses the percentage of individuals that showed a correct behavior in the respective CS+ or CS- trial. The behavioral output is binary with either response or no response. The model triggers a response if the MBON generates one or more spikes.

KC population has been shown to reduce the overlap between different odor representations (38, 39), and, consequently, population sparseness is an important property of olfactory learning and plasticity models in insects (40–44). The system response to a single odor presentation in Fig. 2B demonstrates the transformation of a dense olfactory code at the ORN and PN layers into a population-sparse representation at the KC layer, where less than $< 2\%$ of the total KC population is active at any time during stimulus presentation. This is in good agreement with quantitative estimates in the fruit fly (23, 39, 45).

Few-Shot Learning Rapidly Forms an Associative Memory of Single Cues with Rewards. Many insects exhibit a rapid learning dynamics when trained in classical olfactory conditioning tasks. They typically acquire high retention scores (test accuracy $> 60\%$) for

a binary conditioned response (CR) behavior within only very few trials (e.g., refs. 46–48).

We here mimic a standard experimental laboratory protocol for differential conditioning (or acuity learning) to form associative memories and to generate a binary CR behavior by training our network (Fig. 1). Across successive learning trials, we present two different odors in pseudo-random trial order (Fig. 2A). Each trial constitutes a single odor presentation for 500 ms, followed by a reinforcing stimulus (US) occurring shortly after the stimulus presentation. The CS+ odor is paired with a reward, the CS- odor with a punishment (*Materials and Methods*). In order to establish a neural representation of the odor valence at the MB output (49–52), the MBON is trained (25, 26) to elicit exactly one action potential in response to the CS+ stimulus that is paired with the reward and zero action potentials when the CS- stimulus is presented (*Materials and Methods*). The system response to a single CS+ stimulus after nine conditioning trials is shown in Fig. 2B.

In a first step, we quantified the learning performance by considering the accuracy of the MBON response. MBON output is counted as correct if exactly one spike is generated during a CS+ trial and zero spikes during a CS- trial. The average accuracy over $N = 100$ independently trained model instances across successive trials is shown in Fig. 2C. The learning dynamics shows a steep and steady increase, indicating that an accurate memory is formed rapidly reaching up to 80% accuracy after 50 ($25 \times$ CS+ and $25 \times$ CS-) training trials.

Next, we consider the behavioral learning curve, i.e., the acquisition of a binary CR behavior across successive learning trials. In each trial, the model generates a behavioral response if the MBON produces one or more action potentials in response to the stimulus. No response is generated if the MBON remains silent. A CR is counted as correct if the MBON generates a response to the CS+ cue or no response to the CS- cue. The learning curve in Fig. 2D represents the percentage of correctly responding individuals across $N = 100$ independently trained models. The untrained model, by default, does not generate any output spike; consequently, 100% of the independent models correctly respond to CS- trials from the beginning (Fig. 2D, blue). The orange curve shows a rapid learning success where up to 70% of individuals generated the correct, appetitive CR to the CS+ stimuli within only three to five trials. The learning curve saturates after ~ 10 trials with an asymptotic value of $\sim 80\%$ correct responders. This reproduces the rapid learning dynamics of insects in classical conditioning experiments and fits qualitatively and quantitatively the CR behavior in honeybees (for review, see ref. 47).

We conclude that our statically configured sensory network model with a single plastic readout neuron is capable to successfully form associative memories by few-shot learning, replicating the classical conditioning experiments in the typical laboratory situation. The computational mechanism of population sparseness implemented in our model increases discriminability of the two different stimuli supporting a rapid learning dynamics and a high accuracy of memory recall (*SI Appendix, Figs. S1 and S2*).

Robust Dynamic Memory Recall and Odor-Background Segregation in Complex Sensory Scenes. We now challenge our previously trained model (Fig. 2) in a novel task asking whether the already-learned odor associations can be reactivated in a complex and dynamic olfactory scene. To this end, we mimicked the encounter of odor filaments in a turbulent odor plume during a foraging flight (Fig. 3A). For this, we presented random sequences of nonoverlapping olfactory cues within $T = 10$ s (*Materials and Methods*). Each cue was of variable duration in the range between 1 and 200 ms. Odor identity of each cue was randomly assigned to either the CS+, CS- odor or one out of three additional background odors (Fig. 4A). The use of nonoverlapping

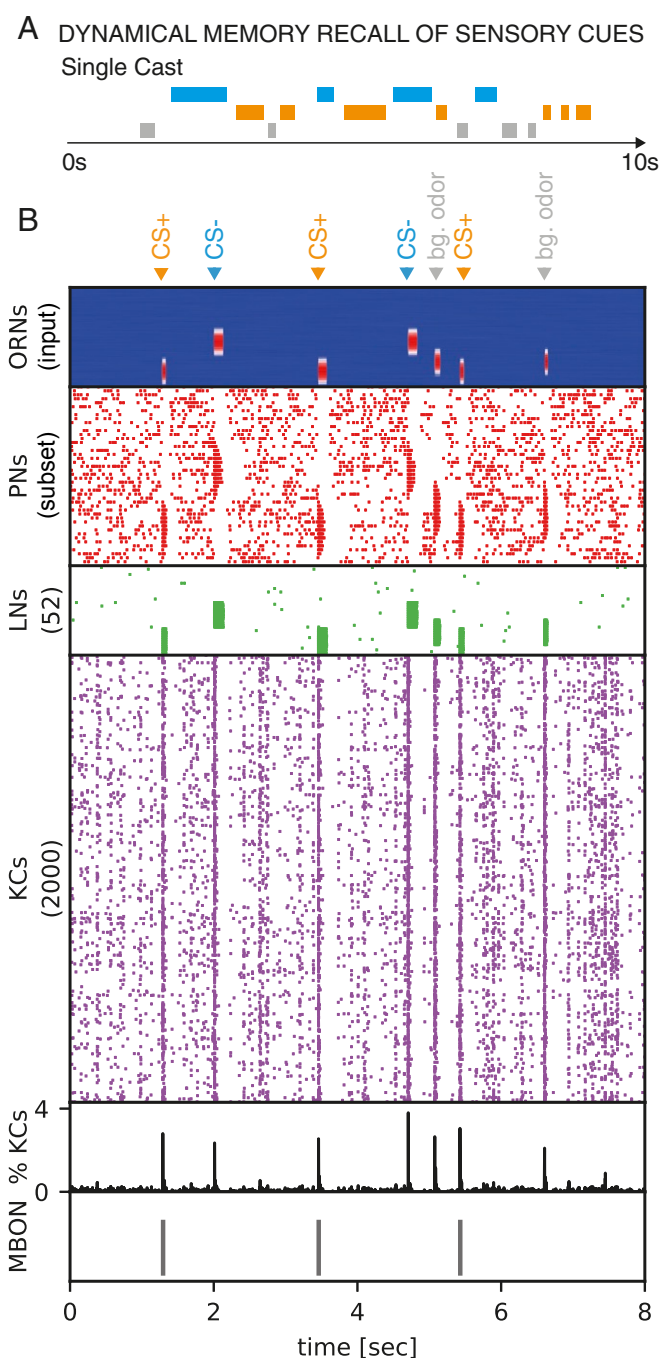


Fig. 3. Recognition of valenced odor cues in complex dynamic scenes. (A) Sketch of the dynamical memory-recall task mimicking the sensory experience during a natural foraging flight. In each single trial of 10s duration, the model encounters multiple (on average five) cues of different odor identities, including the CS+ odor (orange), the CS− odor (blue), and background odors (gray). (B) Network response to one example input sequence made up of three CS+ cues, two CS− cues, and two distractor cues (bg odor), as indicated at the top. The ORN input sequence indicates the fluctuating duration of odor cues. Transient PN and LN response profiles faithfully represent individual odor-cue onsets in time and odor identity across neuronal space. The KC population shows clear responses to all individual odor cues, albeit with < 4% of activated cells at any time. The MBON correctly produced a single action potential in response to each of the three CS+ cues, and zero output else.

cues follows the rationale that, in nature, filaments originating from different odors do not mix perfectly (53).

The objective in this memory-recall task is to correctly detect the occurrences of the positively valenced odor (CS+) by means of a single MBON action potential as model output, while no output should be generated for all other cues (CS− or distractor odors). Fig. 3B shows the system's response to a single random stimulus sequence where the MBON correctly generated a single action potential in response to each of three CS+ encounters. For quantification of task accuracy, we considered the overall response to a given sequence to be correct if the number of action potentials generated by the readout neuron is equal to the number of CS+ cues.

For assessing model performance, we systematically varied task difficulty by varying the number of possible background odors (between one and three) and their similarity with the CS+ odor (Fig. 4). In a first task variant, background-odor activation profiles are rather distinct from the CS+ odor and more similar to the CS− odor (Fig. 4A). Accuracy of the model response was computed across 200 test sequences as shown in Fig. 4B. We find that our previously trained model successfully generalized to this new task with ~80% accuracy for different sequence complexity in terms of identity and number of background odors. In a second task variant, we reversed the odor contingency of the CS+ and CS− odors during initial differential conditioning. Thus, the reward predicting odor CS+ is now more similar to two of the background odors, while similarity with the third background odor remains unchanged (Fig. 4C). In this more challenging case, accuracy reduces to ~50% of sequences for which the model produced the correct number of MBON output spikes. Note that the accuracy measure in Fig. 4 is based on the correct cumulative spike count during a complete trial of 10 s. The more similar a background-odor stimulus profile is to the CS+ odor, the more likely the model will produce false-positive (FP) action potentials in response to such a similar odor and, thus, a total spike count that is higher than the number of CS+ occurrences. This is reminiscent of the effect observed in insects and other animals in odor-discrimination tasks, where perceptually similar odors are more difficult to distinguish from previously learned CS+ odors than perceptually dissimilar odors during memory-retention tests. This might be overcome if similar odors are used during the initial differential conditioning.

We conclude that our network model is able to recall previously learned neural representation of odors and signal their valence in a temporally dynamic setting, where the rewarded and punished odors appear with up to 500-times shorter durations and within an unpredictable temporal cue sequence of previously unknown background odors. The model, thus, also solves the problem of odor vs. background segmentation under quasinatural conditions (54).

Accumulation of Sensory Evidences Informs Motor Control in Foraging. We now consider the situation of foraging within a natural environment (Fig. 5A). The objective is to locate the food source, which emits an attractive odor (CS+), by utilizing the sensory cues present in its turbulent odor plume. We show that cast-and-surge behavior can emerge by accumulation and exploitation of sensory evidence of sequentially experienced individual cues.

For this task, we assumed that thin odor filaments within a cross-wind plane of the concentric odor plume were approximately Gaussian-distributed. This is a reasonable assumption, particularly in a wind-tunnel setting with laminar flow, as typically used in experimental settings (8, 55). When the insect performs a cast through the plume, it encounters filaments as short-lived discrete, sequential events, where each encounter represents a single sensory cue (sketch in Fig. 5B). Therefore, in our simulation of casting flights, the agent encounters sequences of cues and distractors where cue onsets for the CS+ odor are drawn from a Gaussian distribution, while distractor cue onsets

appear uniformly distributed over time (*Materials and Methods*). We further assume that the subject has already formed an association of food with the attractive odor, either through learning or through some genetically predetermined innate valence. To this end, we again used the trained model from the classical conditioning task above (Fig. 2) without any further retraining.

We simulated four consecutive casting trajectories where the agent senses odor cues of sequentially experienced filament encounters. Ongoing accumulation of sensory evidence (Fig. 5C) by low-pass filtering of the readout neuron's output assumes positive values shortly after entering the plume cone and further increases while approaching the plume's center line. When traveling beyond the center line, sensory evidence slowly decreases until the agent leaves the plume cone boundary. When sensory evidence drops to zero and after a fix delay, the agent initiates a U-turn motor command to perform another cross-wind cast.

Responses from our model's readout neuron precisely followed the ground truth of CS+ odor cues, as shown by 10 random casting trajectories in Fig. 5D. Performing analysis by averaging of sensory evidence across these 10 casting trajectories yielded an average evidence (Fig. 5E) that faithfully resembled the underlying, true Gaussian profile of the simulated filaments.

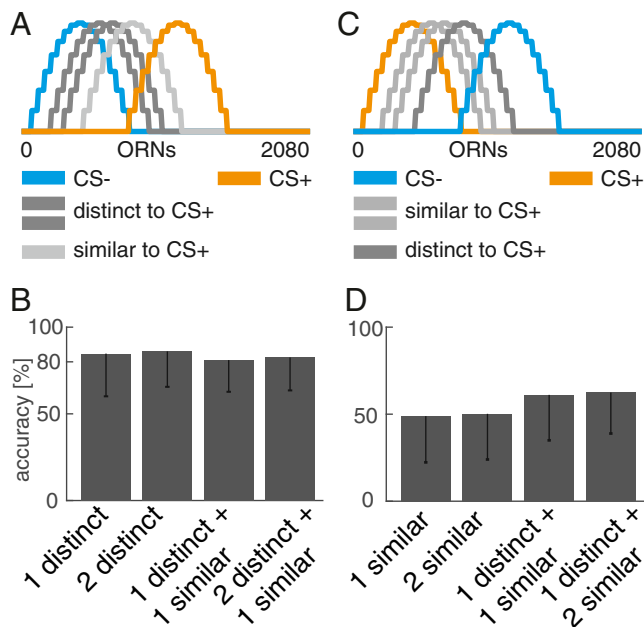


Fig. 4. Performance in the dynamic memory recall task. (A) Input activation profiles across ORNs for five different odors. The model had been previously trained in the differential conditioning protocol using the orange CS+ and the blue CS− odors. In the dynamic memory-recall task (see *Robust Dynamic Memory Recall and Odor-Background Segregation in Complex Sensory Scenes*), background odors (gray) were presented as distractor cues, along with CS+ and CS− cues. In the first task variant, two background odor profiles were rather distinct from CS+ (dark gray), and one was more similar (light gray). (B) Task accuracy across 200 trials in four different scenarios. A single trial consists of a random temporal sequence of sensory-cue presentations during 10 s. Each single cue is of random duration between 1 and 200 ms. The single trial model response was correct if the MBON generated exactly as many action potentials as CS+ cues had been presented. The four task scenarios varied number (between one and three) and type (distinct vs. similar) of distractor cues, as indicated. (C) In the second task variant, the same odors were used for cue presentation in a random temporal sequence. However, the model had been trained with reversed CS+/CS− odor contingency, such that two distractor odors (light gray) were now more similar to the CS+ odor (orange). (D) As in B, but for reversed CS+/CS− odor contingency, where distractor cues were overall more similar to the CS+ odor, increasing task difficulty.

We conclude that the model output provides an accurate and robust estimate of sensory evidence that can be used to reason about a plume's spatial extent and center line. Both pieces of information are crucial to generate appropriate motor commands for U-turn and upwind surge behavior, necessary to successfully execute the cast-and-surge strategy. Apart from the existence of filaments inside a plume and absence outside a plume's cone, our model does not make any specific assumption regarding the plume's structure and statistics. It, thus, provides a generic mechanism implemented in a neural system to perform cast-and-surge behavior during foraging flights.

Discussion

Distinct Functional Roles for Population and Temporal Sparse Stimulus Encoding. Population sparseness improves discriminability of different stimuli to facilitate associative learning. This has been demonstrated in theory and experiment (15, 20, 36, 39). We have shown that our neural network model implements this feature in a biologically realistic way, and our results confirm the functional role of population sparseness to support rapid and robust memory acquisition through associative learning.

Experimental (39) and theoretical (20, 56) studies in the fruit fly suggest that inhibitory feedback through the anterior paired lateral (APL) neuron improves population sparseness and learning in the KC population. In the adult fly, it likely receives input in both the calyx and the lobes of the MB, and it is thought to widely inhibit KCs and possibly PN synaptic boutons that are presynaptic to KCs in the calyx. Whether this neuron is generating sodium spikes or whether it is a nonspiking neuron is currently under debate (57). Here, we tested the performance of our model when implementing the APL neuron as spiking neuron that receives excitatory input from all KCs and inhibits all KCs via weak inhibitory synapses (*SI Appendix, Fig. S2*). Feedback inhibition increased the robustness of model function: In the case of more dense PN–KC connectivity and moderate to low sparseness (~5–10% response-activated KCs), inhibitory feedback could partly re-establish rapid learning during training as well as high performance in the transfer test (*SI Appendix, Fig. S2*) and the foraging task (*SI Appendix, Fig. S3*). GABAergic feedback from MB output onto MB input is a recurring motif across species. In the locust, this is provided by a single nonspiking neuron (58). In the cockroach, there are exactly four spiking inhibitory feedback neurons per hemisphere that have been shown to receive input from the calyx and the lobes (59), while in the honey bee, a population of ~50 spiking and plastic neurons provides GABAergic feedback onto the synaptic boutons of PNs in the calyx (60, 61).

Our model demonstrates how temporal sparseness can be exploited to generate short-patterned signaling of cue identity. This enables perception of high temporal stimulus dynamics. In our model, this is achieved independently of the duration of individual stimulus incidents and their distribution in time and makes temporally precise and robust sensory evidence available. It allows for the ongoing computation of derived estimates such as cue distributions or changes in cue density. Maintaining temporally sparse representations mechanistically supports the principle of compositionality [or Frege principle (22)], where an atomic stimulus entity is represented and can be learned by the readout neuron before processing this output—for example, by estimation of densities or recombination with other entities to form composite perception or memory readout. The temporal stimulus dynamics remains intact throughout the system even after learning of stimulus relevance. Thus, valence is encoded with the same dynamics and faithfully captures occurrences of relevant cues. This allows compression of code to relevant stimuli, while retaining full stimulus dynamics of the external world. Compression of code along the sensory-processing pipeline is particularly relevant for small-brained animals like insects, which need to economize on their neuronal resources.

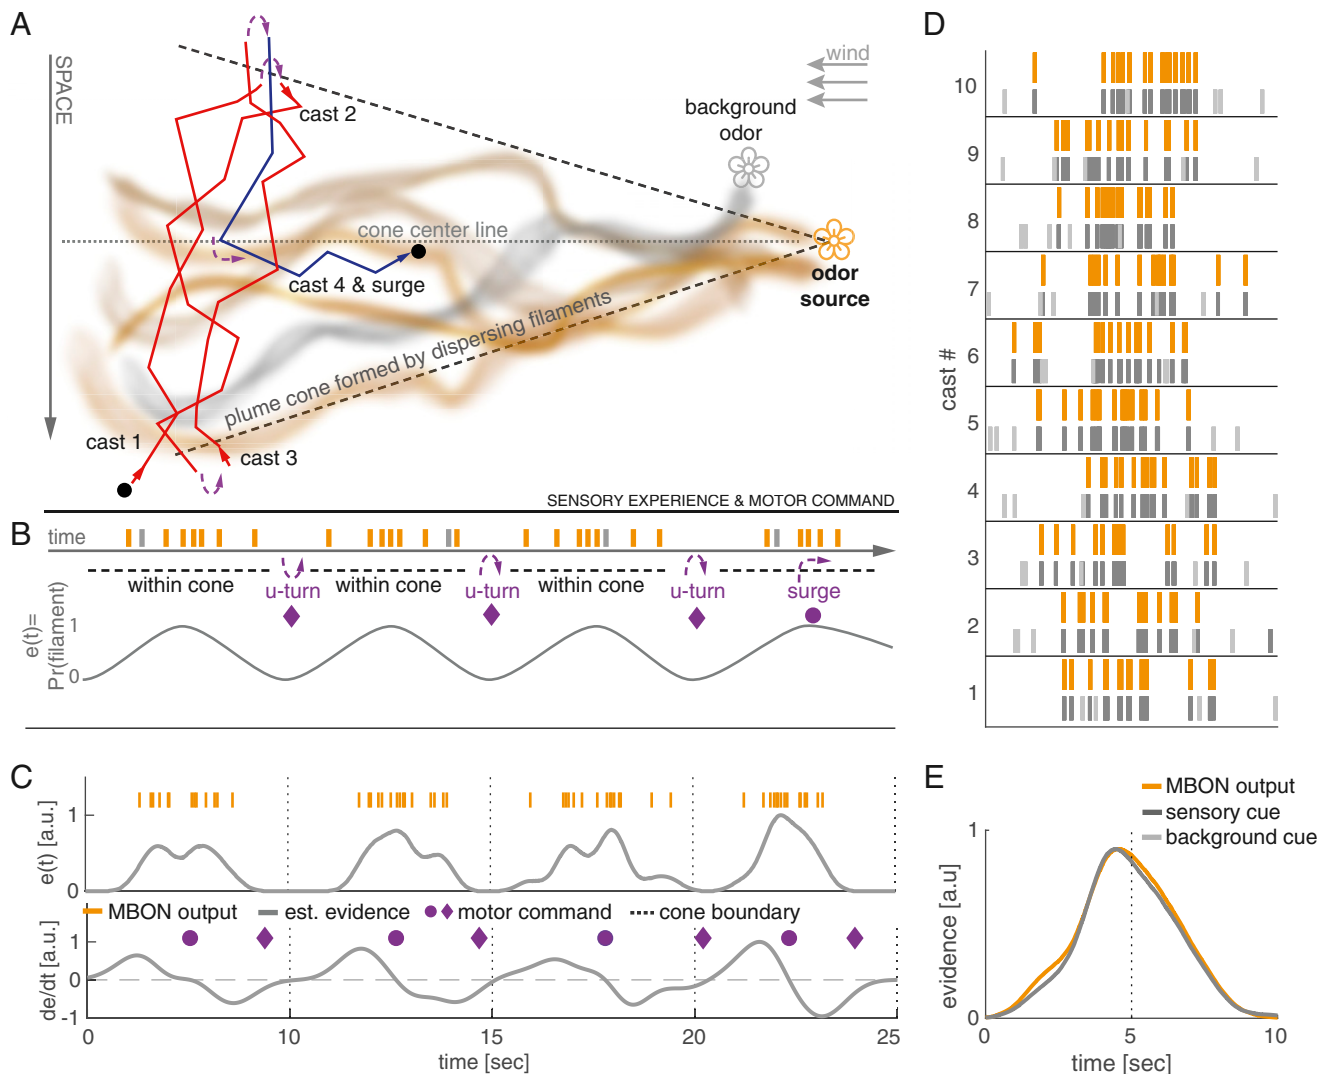


Fig. 5. Dynamical sensory processing and motor control serving chemotaxis. (A) Sketch of a typical olfactory experimental setup in a wind tunnel with a pleasant odor source (orange flower) and a second distractor source (gray flower). Due to turbulence, the odor molecules emitted by a single source form dispersing, intermittent filaments within a cone-like boundary that constitutes the odor plume. The plume is modeled as Gaussian-distributed filaments. A behaving model insect (here, *D. melanogaster*) performs stereotypic cast-and-surge behavior to locate the food source. This constitutes alternating between scanning cross-wind and U-turning after running past the plume-cone boundary where no filaments are present. Eventually, after several casts (here, three), it surges upwind until it loses track of the plume cone and starts over. (B) Filament encounters during this behavior result in sequential brief on/off stimulations of the olfactory system. The probability of encountering filaments is > 0 within the plume and zero outside of the plume. Sensory evidence $e(t)$ can be viewed as a likelihood function of filament encounters that increases toward the plume's center line and is zero outside of the plume. The properties of this function can be used to generate optimal motor commands for chemotaxis. (C) Evidence $e(t)$ and derivative $\frac{de}{dt}$ over four simulated successive casting trajectories estimated from the MBON spiking activity. U-turn motor commands (purple diamonds) are generated when $e(t)$ runs below a fixed threshold (0.01), and surge motor commands (purple circles) are generated when $\frac{de}{dt}$ turns negative. The motor commands generated by the model match well with the theoretically optimal commands, as sketched in B. (D) Spiking activity of the MBON (orange) in response to 10 casting trajectories. The MBON reliably predicts the true sensory cues of positively valenced filaments (dark gray) and ignores background cues (light gray). (E) Smooth peristimulus time histogram computed over 10 casting trials recovers an accurate estimate of the true underlying sensory-cue distribution simulated as Gaussian distribution.

Odor-Background Segregation: A Joint Effect of Temporal and Population Sparse Cue Representation. The task presented in Fig. 3 implicitly addresses the issue of odor-background segregation. This refers to the problem that, in nature, cues of multiple odors of different sources are present, either in terms of mixtures or stimulus-onset asynchrony due to turbulent conditions (53, 54). For behavior, it is relevant to reliably isolate and detect the relevant cues from any background or distractor cues. The results presented in Fig. 4 show that this works nicely in our system. This is achieved by exploiting the joint effect of temporal and population sparseness. Optimal discrimination of cue representation is guaranteed by population sparseness and tem-

poral precision by means of temporal sparseness. Our plastic output neuron requires population sparseness for learning, and the plasticity rule (25, 26) allows for temporally precise memory recall. We predict that our model can solve the challenge of odor-background segregation.

Rapid Learning within Few Trials. The ability of insects to quickly form associative memories after three to five trials has been demonstrated experimentally (47). However, in general, few-shot learning remains a difficult task for computational models, including insect-inspired models (62). We find that, when compared with learning-dynamics data of real insects (47), our model

is able to show realistic learning dynamics that matches with the experimental observations. Due to frequent changes in the environment, it might be a better strategy to trade off fast and reasonable accurate learning against slow and high-precision learning. Additionally, acquisition of training samples might be costly, or they generally occur very sparsely.

In fact, few-shot learning likely is a fundamental skill for survival. We have demonstrated that our neurobiologically motivated approach using spike-based computations is capable to perform few-shot learning with similar speed as insects. We further showed that our model can transfer learned associations to novel, complex combinations that have not been part of the training data.

Innate vs. Learned Behavior. Cast-and-surge behavior belongs to the innate behavioral repertoire of airborne insects and emerges from a set of sensorimotor reflexes (8). It can be considered as a base strategy that guarantees survival. The base system can be modulated and improved throughout an animal's lifespan by experience-based learning. This is superior to alternative strategies that would solely rely on learning appropriate behaviors and, thus, require constant retraining, as is the case in machine-learning (ML) approaches. Here, we assumed that our readout neuron was tuned to a pleasant odor. In the present work, this tuning was learned (adaptive process) in a classical conditioning task. However, a tuning can generally be learned by other mechanisms, e.g., reinforcement or unsupervised learning. We demonstrated that the existence of such a tuned neuron allows cast-and-surge foraging behavior to emerge.

There are other ways how such a tuned neuron can come about, for example, due to genetically predetermined wiring or during development from larval to adult stage. The cast-and-surge behavior can be executed on innately valenced olfactory cues, and our suggested model for motor control during cast and surge (Fig. 5 *A* and *B*) also works for innate valenced stimuli. Learning is important to adapt behavior to changing environmental circumstances, and associative learning provides a means to learn new valences on demand in such situations. Our model learns odor valences at the MB output, and it has been shown that MBONs signal odor valence (49–52). We suggest that this valence is then used downstream to execute higher-level functions of motor control. At this processing stage, it might be integrated with innate valences and other necessary sensory modalities to form behavioral decisions.

Implications for Other Sensory Systems. Sparse stimulus encoding has been identified as a powerful principle used by higher-order brain areas to encode and represent features of the sensory environment in invertebrate (17, 23, 45) and vertebrate (63–66) systems. Sensory systems with similar coding principles may share similar mechanisms when it comes to learning and multimodal sensory integration. The MB is a center for integration of multimodal sensory information. Thus, our model can be extended to incorporate input from different sensory modalities. It is known that olfactory search and foraging strategies do not solely rely on olfactory cues, but require additional sensory information from at least visual cues and wind direction. Extending our model to include additional sensory-processing systems for vision and wind direction can provide a comprehensive functional model to study foraging and navigation.

Potential Improvement through Multiple Readout Neurons. Our current approach only comprises the simplest case of a single readout neuron. This model can be extended to multiple readout neurons. Different readout neurons can be tuned to different odors or groups of odorants. This would allow foraging for different types of food sources and, further, be useful for multi-

modal sensory integration and learning of valences of multiple odors. Another way to use multiple readout neurons is to create an ensemble learning model. Particularly, one can perform bootstrap aggregation (bagging) to decrease variance of predictions. With this technique, multiple, independent readout neurons can be trained for the same target, and their outputs are averaged to produce a single output. This approach can be useful when the level of noise increases due to different input models used to drive the network. Another possible extension is to use a single readout neuron to code for multiple odors by associating different numbers of action potentials to different odors (e.g., two or three). The choice of model for the readout neuron and the plasticity rule allows us to do this (25).

Top-Down Motor Control and Lateral Horn. The model currently lacks a neuronal implementation of sensory-evidence integration and generation of motor commands. Integration of sensory evidence is modeled by low-pass filtering of the readout neuron's spike train, and its derivative is numerically estimated. Drift-diffusion models (67) provide a straightforward alternative to implement integration of evidence to generate a stochastic decision variable that can be thresholded in order to take a behavioral decision for turn or surge (68, 69). Our ultimate goal is a model extension by a neuronal circuit for behavioral decision and top-down motor control. This circuit may involve basic neuronal mechanisms. For example, ref. 70 has shown that a single-compartment Hodgkin–Huxley neuron can operate in two computational regimes. One is more sensitive to input variance, and the model then acts as differentiator, while in the other regime it acts as integrator. Similarly, ref. 71 has shown that the subthreshold current of neurons can encode the integral or derivative of their inputs based on their tuning properties. This and other suggested mechanistic implementations (e.g., ref. 72) could serve as a basis for estimating the low-pass-filtered sensory evidence and its derivative solely by using neural computations. The initiation of a turning behavior based on a time-dependent evidence signal could be implemented, e.g., through disinhibition of motor-command neurons when reaching a decision threshold. The mechanism for U-turning could rely on either cell-intrinsic properties such as SFA, where a neuron initiates a fast turning movement that decays with a fixed time constant, or through state-switching dynamics in neuronal populations.

Relevance for ML and Artificial Intelligence. Learning and building artificial intelligence (AI) agents capable of interacting with their environment are major objectives in the fields of ML and AI. Deep artificial neural networks (73) have demonstrated great success over the recent years, particularly in the domains of image recognition, natural language processing, and deep reinforcement learning (74, 75). Three challenges currently hamper further progress in the theoretical understanding of deep neural networks. These are expressivity, optimization, and generalizability (76). The latter is of interest in the context of the present results. It refers to a model's ability to correctly classify previously unseen data samples during testing that stems from the same a priori defined sample distribution. A related, but much harder, problem is classification of unseen out-of-distribution (OOD) samples that are drawn from a different distribution. Several methods have been proposed for mastering generalization or OOD in deep-learning models, including retraining on new tasks, continual learning (77), meta-learning (78), and transfer-learning (79). However, these methods bear new challenges, such as catastrophic forgetting (80), and some progress has been made toward novel solutions (81). We trained our model on exactly two odor stimuli (CS+ and CS–; Fig. 24). These were drawn from a distribution that is predefined in time (fix stimulus duration) and in neuronal space (receptor activation model; Fig. 44). After training, the model performed well in the dynamical memory-recall task (Fig. 4 *B* and *D*) and in the cast-and-surge task

(Fig. 5 D and E). In both test situations, the previously trained model faced novel stimuli of varying receptor-activation profiles and random stimulus durations, both drawn from different distributions, as used during training. Adding to previous studies on generalizability and learning convergence of the multispike tempotron (MST) model in ML-related problems (25, 26), our results support the idea that spiking neural networks can generalize and solve OOD problems in a biological context, which can possibly be carried over to ML-related tasks. We predict that knowledge transfer from research in neural computation to research in ML and AI will become increasingly important.

Materials and Methods

Code and data sets are available through our GitHub profile at: <https://github.com/nawrotlab> (82).

Spiking Network Model. All neurons of the olfactory network are modeled as conductance-based leaky integrate-and-fire neurons with SFA. Specifically, the membrane potential follows the dynamical current balance (Eq. 1). On threshold crossing, a hard reset of the membrane potential is performed by Eq. 2. SFA is modeled as outward current by term 4 of Eq. 1. Strength of the adaptation current is modeled by a constant (b) decrease on each threshold crossing. Input to the model is modeled as direct, time-dependent current injection of shot noise to all ORNs by the term $I_{stim}(t)$. All simulations of the network are carried out by using the BRIAN2 (83) simulator. The membrane potential of each neuron within a population is initialized randomly $\in [V_{rest}, V_{threshold}]$. To avoid any artifacts, the network is brought to equilibrium by driving the network for 2 s with background activity only before starting the actual simulation.

$$C_m \frac{dv}{dt} = g_i(E_i - v) \quad [1]$$

$$+ g_e(E_e - v) - g_i(E_i - v) - g_{la}(E_{la} - v) + \underbrace{I_{stim}(t)}_{\text{only for ORNs}}$$

$$v = V_{rest} \quad \text{on threshold crossing} \quad [2]$$

$$\tau_{la} \frac{dg_{la}}{dt} = -g_{la} \quad [3]$$

$$g_{la} = g_{la} - b \quad \text{on threshold crossing} \quad [4]$$

For this work, the number of neurons within each layer and connectivity schemes are chosen to match the numbers found in the adult *D. melanogaster* (14, 24). Our model comprises 2,080 explicitly modeled ORNs organized in 52 different receptor types. ORNs of the same receptor type converge onto the same glomerulus (52) by feed-forward excitatory synapses. Each glomerulus is formed by a PN and a LN. LNs provide lateral inhibition to all other PNs and LNs. PNs randomly project to a large population (2,000) of KCs with excitatory synapses, such that each KC on average receives input from six random PNs (connectivity degree parameter K). This sparse random convergence implements population-sparse responses. The single, plastic MBON is fully connected to all KCs.

We used the cellular mechanism of SFA to achieve temporal sparseness. ORNs are configured to have slow and weak SFA in accordance with experimental findings (29, 32). For PNs and LNs SFA has been turned off and KCs are set to produce fast and strong adaptation currents (18, 84). The property of temporal sparseness can also be achieved by an alternative implementation through feedback inhibition as proposed by (56, 85).

The synaptic weights of all connections within the network have been manually determined such that an average background firing rate of 8 to 10 Hz is achieved in the LN population.

Stimulus-Response Profile of ORNs. The stimulus-response profile of ORNs is determined by the ORN tuning curves. We follow a similar method as used in ref. 21, where cyclical tuning over receptor types is modeled as half-period sine waveforms. Our model comprises $N_{type} = 52$ receptor types and supports 52 different stimuli (e.g., different odors), where k_{type} refers to the receptor-type index ($\in [0, 51]$) and k_{odor} to the stimulus index ($\in [0, 51]$). $N_{orn} = 15$ determines the number of receptor types activated by a stimulus. The tuning strength r of the ORNs can be computed as 0.5 cycle of a sine wave with peak amplitude $r_{max} = 1$. In the present work, all tuning profiles are normalized to have a peak amplitude of one.

$$x = \frac{k_{type} - k_{odor} \bmod N_{type}}{N_{orn} + 1} \quad [5]$$

$$r = r_{max} \begin{cases} \sin(x\pi) & \text{for } 0 < x < 1 \\ 0 & \text{else} \end{cases} \quad [6]$$

Model Input. Input to the MB model is modeled as time-dependent, direct-current injection into all ORNs. In the absence of any stimuli, ORNs exhibit spontaneous activity (29). The model input thus consists of spontaneous background activity and stimulus-related activity. To generate the background activity, a current time-series is generated for each ORN by simulating shot noise. For each ORN, background activity events are generated from a Poisson process with high rate ($\lambda = 300$) (independent Poisson processes are drawn for each individual neuron). Events of the Poisson process are filtered by a low-pass filter with $\tau = 0.6$ s. Using this shot-noise model is consistent with experimental findings of odor transduction at the ORNs (29). To induce stimulus-related activity to this time-series of ORN j , it is multiplied point-wise with a stimulation protocol time-series $s_j(t)$, which is rescaled by a constant determined by the tuning strength ($r_j \in [0, 1]$) to the specific odor of the ORN. This results in a current time-series, where during stimulus, the current magnitude is increased proportional to the ORN's tuning strength and otherwise remains at the magnitude of the background activity.

We define a stimulation-protocol function $s(t)$, which is a step function taking on the value one at all time points t where a stimulus or sensory cue is active. For each ORN, a rescaled instance of the stimulation protocol is defined as $s_j(t) = r_j s(t)$, where the scaling parameter $r_j \in [0, 1]$ is given by the stimulus response profile (Eq. 6) of the ORN to the specific stimulus.

$$s(t) = \begin{cases} 1 & \text{if some stimulus is present} \\ 0 & \text{else.} \end{cases}$$

Sequences of Sensory Cues. Each sequence has a duration of 10 s. Sequences of sensory cues are generated by drawing the total number of cues within a single sequence from a Poisson distribution with mean $\lambda = 8$. Onset times of the cues between 0 and 10 s are drawn from a random uniform distribution, and it is assured that there is no temporal overlap between cues. A stimulus relates to a single sensory cue, and its duration is drawn uniformly between [1, 200] ms. Finally, each sensory cue is associated with a random odor drawn from a fixed set of possible odors (random sampling with replacement and equal probability). This results in sequences with a random number of sensory cues, random onset, random duration, and randomized odor and distractor combinations.

Model of Sensory Cues within (Gaussian) Plume. The same procedure is used as above to simulate the experience of sensory cues during a single casting trajectory within a turbulent odor plume. The number of pleasant cues experienced in a casting trajectory is drawn from a Poisson distribution with mean $\lambda = 14$. The cue onset times are drawn from a Gaussian distribution with $\mu = 5, \sigma = 1.5$. The number of distractor cues is drawn from a Poisson distribution with mean $\lambda = 5$ and are distributed uniformly in time. Duration of both, pleasant and distractor cues, is drawn uniformly between [100, 500] ms. In total, 200 different casting trajectories have been generated by using this procedure.

Readout Neuron and Learning Rule. To fit the readout neuron to the stimuli such that it generates one spike for pleasant odor stimuli (CS+) and zero spikes for any other stimuli (CS-), we used a modified implementation of the MST (25, 26). Thus, the readout neuron is modeled as voltage-based leaky integrate-and-fire neuron with soft reset following the dynamical Eq. 7. Incoming spikes evoke exponentially decaying postsynaptic potentials. When the membrane potential reaches the spiking threshold at some time t_0 , an output spike is generated, and the membrane potential is reset by the last term of Eq. 7.

$$V(t) = \underbrace{V_{rest}}_{:=0} + \sum_{i=1}^N \omega_i \sum_{t'_i < t} \underbrace{K(t - t'_i)}_{\text{exp. PSP kernel}} - \underbrace{(\vartheta - V_{rest})}_{:=1} \sum_{t'_sp} e^{-\frac{t - t'_sp}{\tau_m}} \quad [7]$$

The dynamical equation can be decomposed into two parts, the unreset subthreshold potential $V_0(t)$ (Eq. 8) minus the remaining terms for the soft

reset. The neuron was trained to generate one spike for pleasant odor stimuli (CS+) and zero spikes for any other stimuli (CS−). To fit the desired neural code, a training step was performed after each stimulus presentation. A training step was performed only if the number of spikes generated in response to a stimulus was not correct. The training target was given by the difference between number of output spikes the model generated and the number of output spikes associated with the stimulus. We denote the desired critical threshold value, the voltage value that generates $d = 1$ spike, as ϑ^* and the time point where this voltage value is reached by t^* (more generally the critical threshold value to generate d spikes). We briefly sketch the idea and intuition of the MST learning rule. For detailed derivation of the rule, we refer to ref. 25 (section *the ϑ^* gradient*). The MST training algorithm works by differentiating the membrane potential of the critical threshold with respect to the synaptic weights ($\bar{\omega}$). This can be done since ϑ^* is a regular voltage value, that can be expressed by the neuron's dynamical equation (Eq. 7), with the special identities shown in Eq. 10. This allows us to take the full derivative, as shown in Eq. 11.

$$V_0(t) = \sum_{i=1}^N \omega_i \sum_{t_j^i < t} K(t - t_j^i) \text{ unresetsub} - \text{thresh.potential}, \quad [8]$$

$$V(t) = V_0(t) - \vartheta \sum_{t_j^i} e^{-\frac{t - t_j^i}{\tau_m}}, \quad [9]$$

$$\vartheta^* = V(t^*) = V(t_{sp}^i) \text{ critical thresh. that makes } d \text{ spikes}, \quad [10]$$

$$\nabla_{\bar{\omega}} \vartheta^* = \frac{\partial}{\partial \omega} V(t^*) + \sum_{j=1}^m \frac{\partial}{\partial t_{sp}^j} V(t^*) \frac{d}{d\omega} t_{sp}^j. \quad [11]$$

The gradient of the critical threshold with respect to a single synapse i is given by Eq. 12:

$$(\vartheta_i^*)' = \frac{d}{d\omega_i} \vartheta^* = \frac{d}{d\omega_i} V(t^*) = \frac{d}{d\omega_i} V(t_{sp}^i), \quad [12]$$

$$(\vartheta_i^*)' = \frac{\partial}{\partial \omega_i} V(t^*) + \sum_{j=1}^m \frac{\partial}{\partial t_{sp}^j} V(t^*) \frac{d}{d\omega_i} t_{sp}^j \text{ recursive expr. exists.} \quad [13]$$

Data Availability. Code and datasets are available through our GitHub profile at <https://github.com/nawrotlab>. Datasets and source code data have been deposited on GitHub at <https://github.com/nawrotlab/SpikingNeuralProgramForagingInsect-PNAS>.

ACKNOWLEDGMENTS. This research is supported by German Research Foundation Grant 403329959 (to M.P.N.) within the Research Unit "Structure, Plasticity and Behavioral Function of the *Drosophila* Mushroom Body" (DFG-FOR 2705, www.uni-goettingen.de/en/601524.html).

1. K. L. Baker *et al.*, Algorithms for olfactory search across species. *J. Neurosci.* **38**, 9383–9389 (2018).
2. J. Murlis, J. S. Elkinton, R. T. Cardé, Odor plumes and how insects use them. *Annu. Rev. Entomol.* **37**, 505–532 (1992).
3. J. P. Crimaldi, M. B. Wiley, J. R. Koseff, The relationship between mean and instantaneous structure in turbulent passive scalar plumes. *J. Turbul.* **3**, 1–24 (2002).
4. A. Celani, Odor landscapes in turbulent environments. *Phys. Rev. X* **4**, 041015 (2014).
5. E. G. Connor, M. K. McHugh, J. P. Crimaldi, Quantification of airborne odor plumes using planar laser-induced fluorescence. *Exp. Fluid* **59**, 137 (2018).
6. N. Vickers, T. Christensen, T. Baker, J. Hildebrand, Odour-plume dynamics influence fly brain's olfactory code. *Nature* **410**, 466–470 (2001).
7. S. D. Boie *et al.*, Information-theoretic analysis of realistic odor plumes: What cues are useful for determining location? *PLoS Comput. Biol.* **14**, e1006275 (2018).
8. F. van Breugel, M. H. Dickinson, Plume-tracking behavior of flying *Drosophila* emerges from a set of distinct sensory-motor reflexes. *Curr. Biol.* **24**, 274–286 (2014).
9. J. A. Riffell *et al.*, Flower discrimination by pollinators in a dynamic chemical environment. *Science* **344**, 1515–1518 (2014).
10. R. Pang, F. van Breugel, M. Dickinson, J. A. Riffell, A. Fairhall, History dependence in insect flight decisions during odor tracking. *PLoS Comput. Biol.* **14**, e1005969 (2018).
11. M. Vergassola, E. Villermaux, B. I. Shraiman, "Infotaxis" as a strategy for searching without gradients. *Nature* **445**, 406–409 (2007).
12. B. Baddeley *et al.*, What can be learnt from analysing insect orientation flights using probabilistic SLAM? *Biol. Cybern.* **101**, 169–182 (2009).
13. J. B. Masson, Olfactory searches with limited space perception. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 11261–11266 (2013).
14. Y. Aso *et al.*, The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife* **3**, e04577 (2014).
15. S. J. Caron, V. Ruta, L. Abbott, R. Axel, Random convergence of olfactory inputs in the *Drosophila* mushroom body. *Nature* **497**, 113–117 (2013).
16. C. S. Xu *et al.*, A connectome of the adult *Drosophila* central brain. <https://www.biorxiv.org/content/10.1101/2020.01.21.911859v1> (21 January 2020).
17. I. Ito, R. Cy. Ong, B. Raman, M. Stopfer, Sparse odor representation and olfactory learning. *Nat. Neurosci.* **11**, 1177–1184 (2008).
18. H. Demmer, P. Kloppenburg, Intrinsic membrane properties and inhibitory synaptic input of Kenyon cells as mechanisms for sparse coding? *J. Neurophysiol.* **102**, 1538–1550 (2009).
19. P. Szyszka, R. C. Gerkin, C. G. Galizia, B. H. Smith, High-speed odor transduction and pulse tracking by insect olfactory receptor neurons. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 16925–16930 (2014).
20. A. Litwin-Kumar, K. D. Harris, R. Axel, H. Sompolinsky, L. Abbott, Optimal degrees of synaptic connectivity. *Neuron* **93**, 1153–1164.e7 (2017).
21. R. Betkiewicz, B. Lindner, M. P. Nawrot, Circuit and cellular mechanisms facilitate the transformation from dense to sparse coding in the insect olfactory system. *eNeuro* **7**, ENEURO.0305-18.2020 (2020).
22. J. Hintikka, A hundred years later: The rise and fall of Frege's influence in language theory. *Synthese* **59**, 27–49 (1984).
23. G. C. Turner, M. Bazhenov, G. Laurent, Olfactory representations by *Drosophila* mushroom body neurons. *J. Neurophysiol.* **99**, 734–746 (2008).
24. S. Takemura *et al.*, A connectome of a learning and memory center in the adult *Drosophila* brain. *eLife* **6**, e26975 (2017).
25. R. Gütig, Spiking neurons can discover predictive features by aggregate-label learning. *Science* **351**, aab4113 (2016).
26. H. Rapp, M. P. Nawrot, M. Stern, Numerical cognition based on precise counting with a single spiking neuron. *iScience* **23**, 100852 (2020).
27. J. Benda, A. V. Herz, A universal model for spike-frequency adaptation. *Neural Comput.* **15**, 2523–2564 (2003).
28. F. Farkhooi, A. Froese, E. Muller, R. Menzel, M. P. Nawrot, Cellular adaptation facilitates sparse and reliable coding in sensory pathways. *PLoS Comput. Biol.* **9**, e1003251 (2013).
29. K. I. Nagel, R. I. Wilson, Biophysical mechanisms underlying olfactory receptor neuron dynamics. *Nat. Neurosci.* **14**, 208–216 (2011).
30. M. Stopfer, V. Jayaraman, G. Laurent, Intensity versus identity coding in an olfactory system. *Neuron* **39**, 991–1004 (2003).
31. R. I. Wilson, G. C. Turner, G. Laurent, Transformation of olfactory representations in the *Drosophila* antennal lobe. *Science* **303**, 366–370 (2004).
32. S. Krcfczik, R. Menzel, M. P. Nawrot, Rapid odor processing in the honeybee antennal lobe network. *Front. Comput. Neurosci.* **2**, 9 (2009).
33. H. Watanabe, H. Ai, F. Yokohari, Spatio-temporal activity patterns of odor-induced synchronized potentials revealed by voltage-sensitive dye imaging and intracellular recording in the antennal lobe of the cockroach. *Front. Syst. Neurosci.* **6**, 55 (2012).
34. F. Farkhooi, E. Muller, M. P. Nawrot, Adaptation reduces variability of the neuronal population code. *Phys. Rev. E* **83**, 050905 (2011).
35. E. Gruntman, G. C. Turner, Integration of the olfactory code across dendritic claws of single mushroom body neurons. *Nat. Neurosci.* **16**, 1821–1829 (2013).
36. R. A. Jortner, S. S. Farivar, G. Laurent, A simple connectivity scheme for sparse coding in an olfactory system. *J. Neurosci.* **27**, 1659–1669 (2007).
37. R. I. Wilson, Early olfactory processing in *Drosophila*: Mechanisms and principles. *Annu. Rev. Neurosci.* **36**, 217–241 (2013).
38. S. X. Luo, R. Axel, L. F. Abbott, Generating sparse and selective third-order responses in the olfactory system of the fly. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 10713–10718 (2010).
39. A. C. Lin, A. M. Bygrave, A. De Calignon, T. Lee, G. Miesenböck, Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination. *Nat. Neurosci.* **17**, 559–568 (2014).
40. R. Huerta, T. Nowotny, Fast and robust learning by reinforcement signals: Explorations in the insect brain. *Neural Comput.* **21**, 2123–2151 (2009).
41. J. Wessnitzer, J. M. Young, J. D. Armstrong, B. Webb, A model of non-elemental olfactory learning in *Drosophila*. *J. Comput. Neurosci.* **32**, 197–212 (2012).
42. P. Ardin, F. Peng, M. Mangan, K. Lagogiannis, B. Webb, Using an insect mushroom body circuit to encode route memory in complex natural environments. *PLoS Comput. Biol.* **12**, e1004683 (2016).
43. F. Peng, L. Chittka, A simple computational model of the bee mushroom body can explain seemingly complex forms of olfactory learning and memory. *Curr. Biol.* **27**, 224–230 (2017).
44. J. Müller, M. Nawrot, R. Menzel, T. Landgraf, A neural network model for familiarity and context learning during honeybee foraging flights. *Biol. Cybern.* **112**, 113–126 (2018).
45. K. S. Honninger, R. A. Campbell, G. C. Turner, Cellular-resolution population imaging reveals robust sparse coding in the *Drosophila* mushroom body. *J. Neurosci.* **31**, 11772–11785 (2011).
46. M. Bitterman, R. Menzel, A. Fietz, S. Schäfer, Classical conditioning of proboscis extension in honeybees (*Apis mellifera*). *J. Comp. Psychol.* **97**, 107–119 (1983).
47. E. Pamiir, P. Szyszka, R. Scheiner, M. P. Nawrot, Rapid learning dynamics in individual honeybees during classical conditioning. *Front. Behav. Neurosci.* **8**, 313 (2014).
48. L. Scheunemann *et al.*, AKAPS act in a two-step mechanism of memory acquisition. *J. Neurosci.* **33**, 17422–17428 (2013).
49. D. Oswald *et al.*, Activity of defined mushroom body output neurons underlies learned olfactory behavior in *Drosophila*. *Neuron* **86**, 417–427 (2015).
50. M. F. Strube-Bloss, M. P. Nawrot, R. Menzel, Mushroom body output neurons encode odor—Reward associations. *J. Neurosci.* **31**, 3129–3140 (2011).

51. Y. Aso *et al.*, Mushroom body output neurons encode valence and guide memory-based action selection in *Drosophila*. *Elife* **3**, e04580 (2014).
52. M. F. Strube-Bloss, M. P. Nawrot, R. Menzel, Neural correlates of side-specific odour memory in mushroom body output neurons. *Proc. Royal Soc B Biol. Sci.* **283**, 20161270 (2016).
53. A. Sehdev, P. Szyszka, Segregation of unknown odors from mixtures based on stimulus onset asynchrony in honey bees. *Front. Behav. Neurosci.* **13**, 155 (2019).
54. A. Grabska-Barwińska *et al.*, A probabilistic approach to demixing odors. *Nat. Neurosci.* **20**, 98–106 (2016).
55. A. Sehdev, Y. G. Mohammed, T. Triphan, P. Szyszka, Olfactory object recognition based on fine-scale stimulus timing in *Drosophila*. *iScience* **13**, 113–124 (2019).
56. C. Assisi, M. Stopfer, M. Bazhenov, Optimality of sparse olfactory representations is not affected by network plasticity. *PLoS Comput. Biol.* **16**, e1007461 (2020).
57. H. Amin, R. Suárez-Grimalt, E. Vrontou, A. C. Lin, Localized inhibition in the *Drosophila* mushroom body. <https://www.biorxiv.org/content/10.1101/2020.03.26.008300v1> (26 March 2020).
58. M. Papadopoulou, S. Cassenaer, T. Nowotny, G. Laurent, Normalization for sparse encoding of odors by a wide-field interneuron. *Science* **332**, 721–725 (2011).
59. N. Takahashi, H. Nishino, M. Domae, M. Mizunami, Separate but interactive parallel olfactory processing streams governed by different types of GABAergic feedback neurons in the mushroom body of a basal insect. *J. Neurosci.* **39**, 8690–8704 (2019).
60. B. Grünewald, Morphology of feedback neurons in the mushroom body of the honeybee, *Apis mellifera*. *J. Comp. Neurol.* **404**, 114–126 (1999).
61. J. Haenicke, N. Yamagata, H. Zwaka, M. Nawrot, R. Menzel, Neural correlates of odor learning in the presynaptic microglomerular circuitry in the honeybee mushroom body calyx. *Eneuro* **5**, ENEURO.0128-18.2018 (2018).
62. C. B. Delahunt, J. N. Kutz, Putting a bug in ML: The moth olfactory network learns to read MNIST. *Neural Network.* **118**, 54–64 (2019).
63. T. Hromádka, M. R. DeWeese, A. M. Zador, Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biol.* **6**, e16 (2008).
64. W. E. Vinje, Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273–1276 (2000).
65. J. Wolfe, A. R. Houweling, M. Brecht, Sparse and powerful cortical spikes. *Curr. Opin. Neurobiol.* **20**, 306–312 (2010).
66. J. S. Isaacson, Odor representations in mammalian cortical circuits. *Curr. Opin. Neurobiol.* **20**, 328–331 (2010).
67. J. I. Gold, M. N. Shadlen, The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
68. J. D. Davidson, A. El Hady, Foraging as an evidence accumulation process. *PLoS Comput. Biol.* **15**, e1007060 (2019).
69. G. Meckenhäuser, S. Krämer, F. Farkhooi, B. Ronacher, M. P. Nawrot, Neural representation of calling songs and their behavioral relevance in the grasshopper auditory system. *Front. Syst. Neurosci.* **8**, 183 (2014).
70. B. N. Lundstrom, S. Hong, M. H. Higgs, A. L. Fairhall, Two computational regimes of a single-compartment neuron separated by a planar boundary in conductance space. *Neural Comput.* **20**, 1239–1260 (2008).
71. S. Ratté, M. Lankarany, Y. A. Rho, A. Patterson, S. A. Prescott, Subthreshold membrane currents confer distinct tuning properties that enable neurons to encode the integral or derivative of their input. *Front. Cell. Neurosci.* **8**, 452 (2015).
72. B. P. Tripp, C. Eliasmith, Population models of temporal differentiation. *Neural Comput.* **22**, 621–659 (2010).
73. J. Schmidhuber, Deep learning in neural networks: An overview. *Neural Netw.* **61**, 85–117 (2015).
74. T. J. Sejnowski, The unreasonable effectiveness of deep learning in artificial intelligence. *Proc. Natl. Acad. Sci. U.S.A.*, 10.1073/pnas.1907373117 (2020).
75. M. Botvinick, J. X. Wang, W. Dabney, K. J. Miller, Z. Kurth-Nelson, Deep reinforcement learning and its neuroscientific implications. *Neuron* **107**, 603–616 (2020).
76. Y. B. K. Kawaguchi, L. P. Kaelbling, *Generalization in Deep Learning* (Cambridge University Press, Cambridge, UK, 2018).
77. G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, S. Wermter, Continual lifelong learning with neural networks: A review. *Neural Netw.* **113**, 54–71 (2019).
78. C. Lemke, M. Budka, B. Gabrys, Metalearning: A survey of trends and technologies. *Artif. Intell. Rev.* **44**, 117–130 (2013).
79. S. J. Pan, Q. Yang, A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **22**, 1345–1359 (2010).
80. J. Kirkpatrick *et al.*, Overcoming catastrophic forgetting in neural networks. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 3521–3526 (2017).
81. F. Zenke, B. Poole, S. Ganguli, Continual learning through synaptic intelligence. *Proc. Mach. Learn. Res.*, **70**, 3987–3995.
82. H. Rapp, A spiking neural program for sensorimotor control during foraging in flying insects. GitHub. <https://github.com/nawrotlab/SpikingNeuralProgramForagingInsect-PNAS>. Deposited 6 October 2020.
83. M. Stimberg, R. Brette, D. F. Goodman, Brian 2, an intuitive and efficient neural simulator. *eLife* **8**, e47314 (2019).
84. D. G. Wustenberg *et al.*, Current- and voltage-clamp recordings and computer simulations of Kenyon cells in the honeybee. *J. Neurophysiol.* **92**, 2589–2603 (2004).
85. J. E. M. Bennett, A. Philippides, T. Nowotny, Learning with reward prediction errors in a model of the *Drosophila* mushroom body. <https://www.biorxiv.org/content/10.1101/776401v1> (19 September 2019).