

Habituation as a neural algorithm for online odor discrimination

Yang Shen^a, Sanjoy Dasgupta^b, and Saket Navlakha^{a,1}

^aSimons Center for Quantitative Biology, Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724; and ^bComputer Science and Engineering Department, University of California San Diego, La Jolla, CA 92093

Edited by Timothy O'Leary, University of Cambridge, Cambridge, United Kingdom, and accepted by Editorial Board Member Gina G. Turrigiano April 13, 2020 (received for review September 5, 2019)

Habituation is a form of simple memory that suppresses neural activity in response to repeated, neutral stimuli. This process is critical in helping organisms guide attention toward the most salient and novel features in the environment. Here, we follow known circuit mechanisms in the fruit fly olfactory system to derive a simple algorithm for habituation. We show, both empirically and analytically, that this algorithm is able to filter out redundant information, enhance discrimination between odors that share a similar background, and improve detection of novel components in odor mixtures. Overall, we propose an algorithmic perspective on the biological mechanism of habituation and use this perspective to understand how sensory physiology can affect odor perception. Our framework may also help toward understanding the effects of habituation in other more sophisticated neural systems.

habituation | odor discrimination | background subtraction | systems neuroscience | unsupervised algorithm

Habituation is a type of nonassociative plasticity in which neural responses to repeated, neutral stimuli are suppressed over time (1). This computation enables organisms to focus their attention on the most salient signals in the environment, even if these signals are embedded within high background noise. For example, if a dog sitting in a garden habituates to the smell of the flowers, then any change to the environment—for example, a coyote appearing in the distance—would more likely be detected by the dog, despite the fact that the coyote's smell represents only a small component of the raw odor entering the dog's nose (Fig. 14). Habituation is a ubiquitous aspect of sensory processing in many neural systems, and its impairment has been linked to various neural diseases, including autism and Parkinson's (2, 3).

Habituation has also attracted the attention of computer scientists (4). Computational methods that demonstrate the primary effects of habituation (i.e., background subtraction) have been used in robotics applications (5–8) and in deep learning networks to enhance object recognition (9). These studies have begun to reveal the benefits of using habituation in machine learning applications and suggest that models that incorporate additional features of habituation could lead to even more powerful algorithms. Additional models have used habituation for information filtering, pattern classification, and novelty detection (10). Some of these models take inspiration from biology [e.g., habituation in the toad visual system (11, 12)]; however, most do not model any specific biological circuit (13, 14). Our goal here is not to benchmark our approach against other background subtraction algorithms proposed in the literature (biological or otherwise); rather, our goal is to interpret a specific biological mechanism of olfactory habituation from an algorithmic perspective to highlight a connection between sensory physiology and cognition (15). This perspective helps bridge classical neurophysiology and systems neuroscience by describing how an active and dynamical sensory process affects abstract, circuit-level computation.

Specifically, here we focus on short-term habituation in *Drosophila melanogaster* (16), which occurs after about 30 min of exposure to an odor stimulus and lasts for another 30 min after

the extinction of the stimulus. This is in contrast with faster forms of habituation, such as odorant receptor neuron (ORN) adaptation (17, 18) occurring on the order of hundreds of milliseconds, and longer forms of habituation, such as long-term habituation (16), which lasts for days. Different time scales of habituation have also been described in the vertebrate olfactory system (1). Short-term habituation is of interest because it requires a relatively fast online mechanism for its initiation and decay, yet it is also stable enough to have a lasting effect.

Prior work has studied the mechanisms of short-term habituation in the mouse and honeybee olfactory circuits (19–22), but less understood is how habituation in early sensory processing affects odor coding and odor identification that occurs downstream. For example, in honeybee models (21), habituation is achieved through the potentiation of synaptic strength between local inhibitory neurons (LNs) and projection neurons (PNs), as in fruit flies. However, this model focuses on odor representations in the antennae lobe (22) and not in the mushroom body, where odor learning occurs. Similarly, in mice models, habituation has been implemented via short-term depression of synaptic strength between mitral cells in the olfactory bulb (analogous to PNs in the fly antennae lobe) and principal cells in the cortex (analogous to Kenyon cells [KCs] in the fly mushroom body) (19). However, experimental evidence suggests that the effects of habituation are also observed in the olfactory bulb itself (23, 24), and it remains unclear how these two mechanisms work in tandem to facilitate both habituation and dishabituation across multiple time scales (25). The olfactory system of *D. melanogaster*

Significance

Identifying relevant information embedded within a constant and noisy background is a challenge faced by many sensory systems in the brain. Habituation—the reduction in responsiveness to repeated, neutral stimuli—is an unsupervised learning method the brain uses to overcome this challenge. Here, we propose a neural algorithm for odor habituation based on the olfactory circuit in fruit flies. We demonstrate how habituation enhances discrimination between similar odors and improves foreground detection by subtracting background noise. More broadly, this perspective helps bridge the gap between sensory neurobiology and cognition.

Author contributions: Y.S., S.D., and S.N. designed research; Y.S., S.D., and S.N. performed research; Y.S. analyzed data; and Y.S., S.D., and S.N. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. T.O. is a guest editor invited by the Editorial Board.

Published under the [PNAS license](#).

Data deposition: Code for the habituation algorithms have been deposited on GitHub, <https://github.com/aspenshen/Habituation-as-a-neural-algorithm-for-online-odor-discrimination>.

¹To whom correspondence may be addressed. Email: navlakha@cshl.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1915252117/-DCSupplemental>.




Fig. 1. Illustration of the fruit fly olfactory circuit and two habituation-related problems. (A) Cartoon illustration of the effects of short-term habituation. (B) Overview of the fruit fly olfaction circuit. For each input odor, ORNs each fire at some rate, denoted by gray shade. ORNs send feed-forward excitation to PNs of the same type. A single lateral inhibitory neuron (LN1) also receives feed-forward inputs from ORNs and then sends inhibition to PNs. The net firing rate of each PN is the summation of the excitation it received from its ORN minus the inhibition received from the LN1. The locus of habituation lies at the LN1 $\rightarrow$ PN synapse, and the weights of these synapses are shown via red line thickness. PNs then project to KCs, and each KC samples sparsely from a few PNs. KCs that receive a summed input above a threshold will output to a single inhibitory neuron, called APL. APL sends strong feed-back inhibition to the KCs, which silences about 95% of the lowest-firing KCs; the remaining 5% of highest-firing KCs form a “neural tag” for the odor. (C) Illustration of fine discrimination. Prior to habituation, three similar odors (s_A , s_B , and s_C) share overlapping neural tags. After habituation to odor s_A , many of the KCs shared by all three odors stop firing. As a result, the differences between s_B and s_C are heightened leading to improved discrimination. (D) Illustration of foreground (mixture) discrimination. Before habituation, the neural tag of the odor mixture with 80% of odor s and 20% of odor s' is similar to the neural tag of pure odor s . After habituation to s , the neural tag of the mixture resembles the tag for pure odor s' , since the s component of the mixture has been subtracted. KCs activated by odor s (blue); KCs activated by odor s' (green).

is attractive to study because its anatomy and physiology, from the receptor level to the mushroom body, have been relatively well mapped (26–34), allowing us to interpret the effects of habituation on downstream odor coding, learning, and behavior. Short-term habituation has also been described in other systems, such as in mice, cats, and worms (35–38), and may follow similar mechanisms.

Here, we follow experimentally derived mechanisms of short-term habituation in the fruit fly olfactory circuit to develop an unsupervised algorithm (39) for enhancing odor discrimination. Our algorithm is based on the “negative image” model of habituation (16, 40), in which a simple plasticity rule is used to store an inhibitory image of a background odor signal, which is then subtracted from future inputs. Our model can replicate three general features of habituation: *stimulus adaptation* (reduced responsiveness to neutral stimuli that do not have a learned or innate valence associated with them), *stimulus specificity* (habituation to one stimulus does not reduce responsiveness to other, distinct stimuli), and *reversibility* (dishabituation to the background in case it becomes relevant). We also demonstrate the relevance of these features for computational problems, such as online similarity search (41) and background subtraction (42, 43), using both neural and machine learning datasets.

Methods

Basic Anatomy of the Fruit Fly Olfactory Circuit. Odors in the fruit fly olfactory circuit are encoded using five main steps (Fig. 1B).

- 1) In the first step, ORNs are activated when odor molecules bind to specific receptors. There are approximately 50 types of odorant receptors in the fruit fly antennae, and a single ORN contains only one type of receptor. Thus, an odor is initially encoded as a point in $\mathbb{R}_+^d$, where $d = 50$. The coding of odors across the ORNs is combinatorial (44–46), such that every

ORN responds to almost every odor with a different activation level, and only a few ORNs respond highly to a given odor.

- 2) In the second step, ORNs extend axons into structures called glomeruli in the antennae lobe, where they synapse onto dendrites of PNs. There are approximately 50 PN types, and ORNs of the same type send odor information to corresponding PNs of the same type. This step normalizes PN responses such that the mean firing rate of PNs is nearly the same for all odors (47–50). In addition, PNs receive inhibition from lateral inhibitory neurons in the antennae lobe, including LN1 interneurons, which are involved in habituation, as described below.
- 3) In the third step, the dimensionality of the representation expands: The 50 PNs project to about 2,000 KCs in a structure called the mushroom body. Each KC receives input from approximately six random PNs (30), and KCs respond by thresholding the sum of their PN inputs (26, 51, 52). Thus, KCs have a rectilinear threshold activation function, which helps suppress PN noise (53, 54).
- 4) In the fourth step, each KC sends feed-forward excitation to a single anterior paired lateral inhibitory neuron (APL) (55), which, in turn, sends inhibitory feedback to each KC. As a result of this feedback loop, only $\sim 5\%$ of the most highly activated KCs fire for each odor in what is called a “winner take all” computation (28, 29). Thus, an odor is finally encoded as a sparse point in a 2,000-dimensional space. We refer to the 5% of KCs that fire for a given odor as the neural “tag” (the term used in neuroscience) or “hash” (the term used in computer science) for the odor. Because this circuit generates tags that are similarity preserving (41), tags for similar odors share more overlapping KCs than tags for different odors.
- 5) In the fifth step, KCs send feed-forward excitation to mushroom body output neurons (MBONs) (31) that “decode” the input and drive behavior. For example, there are MBONs involved in learning approach and avoidance behavior. When an odor is paired with some reward or punishment, the strength of synapses between KCs activated for the odor and downstream approach and avoidance MBONs are modified (32). While there are many complexities in the decoding process still to be determined, it is generally accepted that it is easier to learn to associate different behaviors to two different odors when the KCs activated for

the two odors are nonoverlapping, so that the modified $KC \rightarrow MBON$ synapses do not “interfere” with one another (29).

More detailed circuit anatomy and physiology in the fruit fly olfactory system can be found in previous reviews (56–58).

The “Negative Image” Model of Habituation. We now introduce a circuit mechanism for habituation (16, 40) in the fruit fly olfactory system, and how it may be used to enhance discrimination between similar odors.

Expanding on step 2 of the circuit above, the primary locus for habituation occurs at the synapses between LN1 inhibitory interneurons and PNs in the antennae lobe (16, 40, 59) (Fig. 1B). ORNs, in addition to projecting to PNs, also send feed-forward excitation to LN1 inhibitory neurons, which, in turn, synapse onto PNs to inhibit their firing. Habituation is driven by the increased synaptic strength between LN1s and PNs. In our model, this synapse is subject to activity-dependent Hebbian plasticity, such that coincident firing of an LN1 and PN increases the weight of the $LN1 \rightarrow PN$ inhibitory synapse. If the ORN is persistently active, the corresponding $LN1 \rightarrow PN$ synapse weight will increase, causing the firing of the PN to decrease over time. To achieve odor-specific habituation, it is important that only the PNs active for an odor be suppressed, and that the amount of suppression is related to the activity of the PN.

Although, to our knowledge, there’s no direct experimental evidence of potentiation in synaptic strength between LN1 and PNs, studies have shown that inhibitory LN1 is required for short-term habituation and that γ -aminobutyric acid (GABA)ergic transmission from LN1 is necessary and sufficient for habituation (16). In addition, computational modeling and experimental work in the honeybee olfactory system have shown that Hebbian modification of synapses from inhibitory LN1s to PNs, rather than ORN adaptation, most reliably explains the shift of an odor mixture’s representation in the antennae lobe due to habituation (22). Finally, in honeybees, inhibitory synapses between LN1s and PNs have been shown to modify the PN response to an odor mixture after exposure to the pure odor component in the mixture (21).

Formally, we are given an ORN (s_i), its corresponding PN (x_i) of the same type, a single LN1 inhibitory neuron (ℓ) that receives input from all ORNs, and the weight (w_i) of the synapse between the LN1 and the PN (Fig. 1B). We assume each input is normalized such that the total activation of the LN1 neuron is a constant, and, without loss of generality, we assume the constant to be 1 (i.e., $\ell = 1$). Then the PN activity equals

$$\begin{aligned} x_i &= \max(s_i - \ell w_i, 0) \\ &= \max(s_i - w_i, 0). \end{aligned} \quad [1]$$

If s_i is a consistently similar background odor, then we would like w_i to converge over time to s_i to form a “negative image” of the input. After subtracting this negative image from the next input, what remains is the foreground component of the odor. To achieve this, the weight update rule for w_i at time t is

$$\begin{aligned} w_i^t &= w_i^{t-1} + \alpha x_i^{t-1} - \beta w_i^{t-1} \\ &= w_i^{t-1} + \alpha x_i^{t-1} - \beta w_i^{t-1}. \end{aligned} \quad [2]$$

The second term is the activity-dependent Hebbian update rule where coincident activity of the LN1 and PN increases w_i , weighted by the habituation rate, $\alpha \in [0, 1]$. Thus, the magnitude of the increase in synaptic weight depends on the magnitude of the PN firing rate. Each PN (x_i) is associated with its own weight w_i , and thus the amount of habituation can be different for each PN. The third term allows for gradual dishabituation, controlled by a weight recovery rate $\beta \in [0, 1]$. We initialize $w_i^0 = 0$ for all i .

There are two simplifications made in the above circuit. First, anatomy shows that there is not a single LN1 inhibitory neuron, but rather about 18 to 20 LN1 neurons, per antennae lobe (60, 61). However, each LN1 neuron innervates nearly every glomerulus (60–63), and thus we simplify the architecture to contain a single panglomerular LN1 neuron. Second, anatomy shows that habituation is actually a three-step process, in which ORNs excite PNs, which excite LN1s, which inhibit PNs. We simplify this to a two-step procedure, in which ORNs excite both PNs and LN1s, and then LN1s inhibit PNs. While there may be subtle differences between the two, computational modeling (59, 64) suggests that, for creating inhibitory images, inhibitory synapses can be either feed-forward (as in our model) or feed-back, as long as feed-back inhibition from LN1 to PNs is fast. This would make the time lag between PNs firing (due to ORN excitation) and PNs being inhibited (via feedback from LN1) to be small, ensuring that the effects of habituation are applied before PNs pass information downstream to KCs.

Problem 1: Habituation and Fine Discrimination. The first problem we consider is motivated by observations of fine odor discrimination in humans. Imagine three odors that are all very similar to each other, such as three types of rose flowers. An experimental study showed that, after habituating to one flower’s scent, even untrained noses could discriminate between the other two flower scents with higher accuracy than a control group that did not habituate (65). Intuitively, habituation to one odor causes some common representation shared by all three odors to be subtracted; this expands the differences between the two remaining odors, making them easier to discriminate (Fig. 1C).

Formally, as input, we are provided a stream of odors, $S = [s_A^{(1)}, s_A^{(2)}, \dots, s_A^{(k)}, s_B, s_C]$, where one odor is observed in each time step. Each odor is a point in $\mathbb{R}_+^d$, where $d = 50$ corresponds to the 50 ORN firing rates for the odor. The $s_A^{(i)}$ represent noisy versions of odor s_A . In the simplest case, all $s_A^{(i)}$ are the same odor, that is, $s_A^{(i)} = s_A$ for all $i \in [1, k]$. After experiencing the odor k times, two test odors follow. All three odors (s_A, s_B, s_C) are highly correlated with each other, and thus, without any habituation, s_B and s_C would smell very similar. The goal is to apply habituation to s_A to better discriminate between s_B and s_C .

What does it mean to “discriminate” between two odors? As described in the steps above, each odor is assigned a sparse point in a high-dimensional space, consisting of $m = 2,000$ KCs, of which only roughly 5% are active. This transformation can be viewed as a hash function $I: \mathbb{R}_+^d \rightarrow \mathcal{K}$, where $\mathcal{K}$ is a set of KCs that are active for the odor, and $|\mathcal{K}| = 0.05m$. The function I is locality sensitive (41, 66), meaning that, if two inputs (such as s_B and s_C) are highly correlated, they will share many active KCs. We denote $I(s_B)$ to be the tag assigned to odor s_B , and $I(s_B; s_A)$ to be the tag assigned to s_B after habituation to odor s_A .

The ability to discriminate, then, is related to the amount of overlap between $I(s_B; s_A)$ and $I(s_C; s_A)$, where fewer shared KCs leads to better discrimination (57). Biologically, KCs synapse onto approach and avoidance circuits, and the relative strengths of these synapses for odor-activated KCs determine which behavior will be applied when the odor is smelled (31, 32, 67, 68). For the fly to learn to, say, approach s_B and avoid s_C , it is important that synaptic interference between the two odors is low (27); that is, there are KCs active for s_B that are not active for s_C , and vice-versa. Thus, the smaller the overlap between the KCs active for s_B and s_C (both after habituation to s_A), the easier it is to discriminate them.

Problem 2: Habituation and Foreground Discrimination. The second problem we consider highlights the role of habituation for detecting novel components of odor mixtures.

As input, we are given a stream of odors, $S = [s^{(1)}, s^{(2)}, \dots, s^{(k)}, s'']$. The odors $s^{(i)}$ again represent a persistent background odor (e.g., a garden with flowers). Then, a new odor s'' is presented (e.g., a coyote appearing in the garden), which is a mixture of s and a new odor s' (Fig. 1A). We are particularly interested in regimes with low foreground to background ratios; for example, s'' could be a mixture consisting of 80% background (s) and 20% foreground (s'). The goal is to robustly detect the foreground component of the mixture (i.e., the coyote) by eliminating or ignoring the background (Fig. 1D).

What does it mean to “detect” the foreground? As before, the more overlap between the KCs responding to the mixture (s'') and the KCs responding to pure odor s' , the easier it is to detect the coyote (s') and recall the appropriate learned response.

Experimental Dataset and Setup. We tested our habituation algorithm for these two problems using published experimental recordings in the fruit fly for $d = 24$ ORNs responding to 110 odorants (46). We normalized each input vector such that all values are nonnegative and the mean of the vector is a constant (set to 10), mimicking divisive normalization that occurs in the antennae lobe (47).

Applying habituation. Each time an input vector s is presented, we apply Eq. 1 to compute the activity of each PN, and then we apply Eq. 2 to update each habituation weight. Each weight w_i is initialized to 0. We set $\alpha = 0.05$ and $\beta = 0.01$. The effect of different α, β values on habituation is shown in [SI Appendix, Fig. S1](#).

Computing the neural tag of an input. Given a PN vector x , the KC responses y are computed as $y = \Theta x$. Here, Θ is a sparse random matrix, containing $m = 1,000$ rows and $d = 24$ columns. This value for m was selected to preserve a dimensionality expansion ratio of about 40 to 50, as in the fly. In each row, $c = 3$ random positions are set to 1, and the rest of the positions in the row are set to 0. This value for c was selected to approximately match the sampling ratio of each KC (6 glomeruli out of $d = 50$) (30). After the

projection step, a rectilinear function is used to filter out noise in the KC responses. For each $1 \leq i \leq m$, we set

$$y'_i = \begin{cases} y_i & \text{if } y_i \geq \tau_o \\ 0 & \text{otherwise,} \end{cases} \quad [3]$$

where τ_o is a small constant equal to the mean of the input vectors. Finally, the winner-take-all sparsification step sets the top 5% of the highest firing KCs to 1, and the remaining KCs are set to 0. This gives us a vector $z \in \{0, 1\}^m$, where

$$z_j = \begin{cases} 1 & \text{if } y'_j \in \{[0.05m] \text{ largest nonzero entries of } y'\} \\ 0 & \text{otherwise} \end{cases} \quad [4]$$

Why is the first KC thresholding step needed, in addition to the second? The winner-take-all alone always allows about 5% ($0.05m$) of KCs to fire, even if the KCs are firing at very low rates (which, e.g., could occur when an odor is close to being habituated). In such a case, the activity of these KCs could drive approach or avoidance behavior, despite mostly encoding noise. As a result, a minimum threshold (τ_o) for firing is required to suppress KC noise; this is similar to the notion of a “firing threshold” for a neuron or the “rectilinear threshold” of rectified linear units in artificial neural networks.

Thus, the tag of an input is a binary vector of length m , with up to $0.05m$ values as ones, and the rest as zeros. Equivalently, the tag can be represented as a set of KC indices that are active for the odor: $l(x) = \{j: z_j = 1\}$. We do not arbitrarily break ties, and thus, if there are ties in the KC firing rates, there may be slightly more ones in the vector.

Neural tag similarity. The KC tag $l(s)$ for an odor s is a set of indices (each between 1 and m), corresponding to the KCs active for odor s . The similarity between the tags of two odors s and s' is defined as the Jaccard coefficient between $l(s)$ and $l(s')$,

$$\frac{|l(s) \cap l(s')|}{|l(s) \cup l(s')|}.$$

The Jaccard coefficient is a widely used measure to gauge the similarity between two sets. Its value lies between 0 and 1, with 0 meaning that the two sets have nothing in common, and 1 meaning the two sets have identical items.

Forming binary odor mixtures. Predicting ORN responses to an odor mixture remains a challenging problem, with evidence that ORNs for mixtures can respond both linearly (69, 70) and nonlinearly (71–73) with respect to the mixture’s individual components. We explore two models here to form mixtures. For the linear model, given two odors s and s' , the binary mixture is computed simply as $s'' = \xi s + (1 - \xi)s'$, where ξ determines the concentration of each odor component. For experiments in this paper, we set $\xi = 0.8$ to explore regimes with a dominant background odor (s) in the mixture, and a low signal foreground (s'). For the nonlinear model, the binary mixture is computed as $s'' = s^\xi + s'^{(1-\xi)}$. For both models, s'' is normalized to have the same average firing rate as the single odors (45).

Nearest neighbors analysis for online similarity search. For each of the 110 odors, we first compute a ground-truth list of the top N nearest neighbors for the odor. For odor s , this list is computed based on the Jaccard similarity between $l(s)$ and the tag for each other odor. For a mixture, $s'' = \xi s + (1 - \xi)s'$, we seek to determine how well the predicted nearest neighbors for s'' after habituation to s overlap with the ground-truth nearest neighbors of pure s' .

The overlap between two lists of N nearest neighbors is computed using the mean average precision (mAP) (74). The mAP measures the similarity between two ranked lists. Common items between the two lists are weighted by their ranks in the lists. Formally, assume there are N_{tot} items (odors), and let l_1 and l_2 be two ranked lists of $N < N_{\text{tot}}$ of the items. Each list can contain items not present in the other list. The list l_1 is the ground-truth list, and l_2 is the predicted list whose similarity to l_1 we wish to compute. Let l_{12} be the list of the N_k common items in l_1 and l_2 , in the same order as they appear in l_1 . Finally, let $r_{12}(i)$ be the rank in l_1 of the i th item in l_{12} . Then,

$$\text{mAP}(l_1, l_2) = (1/N_k) \times \sum_{i=1}^{N_k} i/r_{12}(i).$$

We compute the mAP over three values of $N = \{10, 20, 30\}$ and report the average. The mAP ranges from 0 (least similar) to 1 (most similar). We compare the similarity of two lists using the mAP, as opposed to the Jaccard coefficient, because the former takes the rank of each nearest neighbor into account, as well.

Results

First, we provide an example of habituation and prove that, under our model, habituation weights converge to form a “negative image” of a constant input. Second, we use habituation to address problem 1 (fine discrimination) and problem 2 (foreground discrimination) both empirically and analytically. Third, we demonstrate applications of habituation for improving online similarity search.

An Example of Habituation and Stimulus Specificity. Fig. 2 shows an example of how PN responses to repeated presentations of the same odor (odor A: acetic acid) change over time. Initially, PN firing rates for the odor are not suppressed, and the inhibitory synaptic weights between the LN1 and PNs are low. With time, these synaptic weights become stronger at different rates according to the firing rate for each PN, and this causes PN firing to decrease (Fig. 2A).

Habituation is also stimulus specific (Fig. 2B). Before habituation, about 50 (5% of the 1,000) KCs respond to odor A, but, after habituation, these KCs are almost entirely silenced. On the other hand, KCs remain responsive to a different odor (odor B: ethyl hexanoate) after habituation to odor A (Fig. 2B), suggesting that habituation to one odor does not hamper responses to different odors. In addition, habituation is reversible, where the circuit can dishabituate to one odor when it is not observed, and then habituate to a different odor (SI Appendix, Fig. S2).




Fig. 2. Changes in circuit activity due to habituation. (A) Over time (x axis), the synaptic weights between the LN1 and activated PNs increase (y axis, left), and the corresponding PN firing rates decrease (y axis, right). The weights between the LN1 and all nonactivated PNs remain unchanged (green line at $y = 0$). There are 24 total PNs, and each curve corresponds to a different PN. (B) Before habituation, about 5% of the 1,000 KCs respond (vertical bar) to odor A. After habituation to A, only a few KCs respond for A. However, after habituation to odor A, KCs still respond to a different odor B, suggesting that habituation is odor specific. (C) *Left* shows that odor A and odor B are relatively distinct odors with little overlap in their tags. *Right* shows that, after habituation to odor A, a substantial number ($44/50 = 88\%$) of KCs in original tag of odor B still respond to odor B after habituation to odor A. Odor A: acetic acid. Odor B: ethyl hexanoate.

Convergence of the habituation weight vector to form a “negative image.” Suppose w is initialized to zero and odor $s \in \mathbb{R}^d$ is repeatedly presented. A quick calculation shows that the limiting value of w is then

$$w^* = \frac{\alpha}{\alpha + \beta} s,$$

indicating that w indeed forms a “negative image” of the input (modulo the dishabituation rate, β). The following lemma reaffirms this and shows a geometric rate of convergence.

Lemma 1: Suppose the current weight vector is some $w \preceq w^*$; that is, w is coordinate-wise less than or equal to w^* . Assume $0 < \alpha, \beta < 1$ and $\alpha + \beta < 1$. Let $w' = (1 - \beta)w + \alpha(s - w)_+$ denote the updated vector upon presentation of stimulus s . Then $w' \preceq w^*$ and

$$w^* - w' = (1 - (\alpha + \beta))(w^* - w).$$

Proof: Since $w \preceq w^* \preceq s$, we have $(s - w)_+ = s - w$, and thus

$$\begin{aligned} w^* - w' &= w^* - (1 - \beta)w - \alpha(s - w) \\ &= w^* - \alpha s - (1 - (\alpha + \beta))w \\ &= w^* - (\alpha + \beta)w^* - (1 - (\alpha + \beta))w \\ &= (1 - (\alpha + \beta))(w^* - w). \end{aligned}$$

□

Neural Tags Remain Robust after Habituation to Another Odor. For habituation to be stimulus specific, the KC tag for an odor s' should remain relatively unchanged before and after habituation to a very different odor, s . Indeed, long-term learning would be difficult if odor tags constantly changed due to habituation.

To test this, we first generated a synthetic dataset of $d = 50$ ORNs responding to 110 odors. Each odor is a list of 50 values, sampled independently from an exponential distribution with $\lambda = 10$. For all odor pairs s and s' , we calculated the neural tag similarity between $I(s')$ and $I(s'; s)$ —the tags of s' before and after habituation to s , respectively. We found that, on average, 67% (median of 68%) of KCs in the original tag of odor s' remain active in the tag after habituation. In addition, the less overlap between $I(s)$ and $I(s')$, the more KCs remain unchanged after habituation (correlation coefficient = -0.47 ; *SI Appendix*, Fig. S3A). As an example, we show that, after habituation to odor A, the tag of odor B remains largely unchanged (Fig. 2C).

Second, we repeated this analysis for real odorant data (46), which also contains 110 odors, but with responses of $d = 24$ ORNs. We observed similar results—mean and median Jaccard similarity of 61% and 63%, respectively (*SI Appendix*, Fig. S3B)—but with larger spread of the data compared to the synthetic data. The larger spread is likely due to the biased sampling (24 out of 50 ORNs) and because the real odors are more correlated with each other, whereas synthetic odors were sampled independently (*SI Appendix*, Fig. S3C).

Overall, after habituation to an odor, the tag of another odor remains relatively stable; the less similar the two odors, the more stability. These results are also supported by theoretical analysis of tag robustness due to habituation, described in *SI Appendix*, Lemma S2 and Corollary S2.

Exactly how many KCs need to be shared between two odor tags such that the learned behavior for one odor can be generalized to the other odor remains unclear, biologically. However, large variations are observed experimentally in KC responses to the same odor across multiple trials, with only about one-third of KCs responding reliably across trials (75). Thus, the overlap observed here ($\sim 65\%$) may be sufficient to recognize odors after habituation, and hence, learned behaviors for an odor can still be recalled and correctly applied after habituating to another odor. Other processes that occur upstream, such as different LN1–PN architectures, or that may occur downstream, such as pattern completion, could further facilitate odor recognition.

Habituation Enhances the Separation of Similar Odors (Problem 1). Here, we show that triplets of similar odors can be better discriminated after habituating to one of the odors.

As an example, we selected three odors (s_A : dimethyl sulfide; s_B : acetone; s_C : 2-butanone) whose ORN representations are highly correlated pairwise (Fig. 3A). As a result, the tags $I(s_B)$ and $I(s_C)$ for odors s_B and s_C are highly overlapping, sharing 38 out of the ~ 50 (5% of 1,000) KCs (Fig. 3B, Left). We then habituated to odor s_A , calculated the tags $I(s_B; s_A)$ and $I(s_C; s_A)$, and found that the number of overlapping KCs shared by these two tags reduced to 25 (Fig. 3B, Right), making the two odors easier to distinguish. This reduction is caused by suppression of PNs responsive to s_A (Fig. 3C), which also fire for s_B and s_C , since all three odors are highly correlated. As a result, some KCs common to all three odors are suppressed, allowing other KCs previously below the 5% threshold to become activated and part of the posthabituation tags for odors s_B and s_C .

To show that this result generalizes, we repeated the analysis above using all 99 odor triplets in the dataset with pairwise correlation of > 0.80 . We computed the tag similarity between two of the odors before habituation and then after habituation to the remaining odor. We observed an average reduction in KC overlap of $26.7 \pm 15.4\%$ ($p < 0.001$, paired t test) before versus after habituation (Fig. 3D). Thus, our model demonstrates how habituation can help better discriminate between similar odors and how prior experience may shape the perception of odors (65).

Critically, even though habituation to odor s_A changed the tags of odors s_B and s_C , on average, about 60% of KCs that fire for s_B remain firing for s_B after habituation to s_A (similar statistics for odor s_C). Thus, habituation helps discriminate two similar odors, while also not hindering the recognition of these odors.

Habituation Enables Robust Foreground Discrimination in Odor Mixtures (Problem 2). Experimental studies in the honeybee antennae lobe show that PN responses to a binary odor mixture become more similar to one of its components after habituating to the other component (21, 22). Similar results have been observed in the locust olfactory system (69), where KCs that respond to one component also often responded to mixtures containing that component. Below, we show that our model can replicate this result and also provide insight into how habituation in the antennae lobe may affect odor coding downstream in the mushroom body.

The following exemplifies this idea (Fig. 4A). Computationally, we first calculated the tags of pure odor s (acetic acid), pure odor s' (ethyl hexanoate), and their mixture, $s'' = 0.80s + 0.20s'$ prior to any habituation. The three tags are denoted $I(s)$, $I(s')$, and $I(s'')$, respectively. Because s'' is composed of 80% s and 20% s' , there are more KCs shared between s'' and s (59% Jaccard similarity between $I(s'')$ and $I(s)$) than s'' and s' (32% Jaccard similarity). We then habituated to s and recalculated the tag of the mixture, $I(s''; s)$. After habituation, $I(s''; s)$ and $I(s')$ strongly overlap (96% Jaccard similarity). Thus, habituation to the dominant component of the odor allowed the minor component (s') to “pop out” above the background, making it easier to detect.

We expanded this test from one example odor pair to all $\sim 6,000$ pairs of odors and their corresponding mixtures. The similarity between the tag of the mixture after habituation and the tag of odor s' (the minor component) is nearly sixfold higher than their similarity before habituation (Fig. 4B). We repeated this analysis with a nonlinear model to form mixtures (*Methods*), and observed similar trends (*SI Appendix*, Fig. S4). Finally, we also used this simple odor habituation algorithm, derived from biological mechanisms, to perform background subtraction in streaming imaging data (*SI Appendix*, Fig. S5).




Fig. 3. Habituation and fine discrimination. (A) Odors s_A , s_B and s_C are highly correlated with each other. Shaded circles illustrate ORN firing rates. (B) (Left) Before habituation to odor s_A , the tags for odors s_B and s_C highly overlap. (Right) After habituation to odor s_A , overlap between the tags of odors s_B and s_C decreased. Odor s_A : dimethyl sulfide. Odor s_B : acetone. Odor s_C : 2-butanone. (C) Firing rates of PNs for each odor before and after habituation to odor s_A . After habituation, PNs that respond highly to odor s_A are mostly silent for all three odors, while some activity of other PNs remains intact for odors s_B and s_C . (D) Comparing Jaccard similarity between odors s_B and s_C before and after habituation to odor s_A , for all odor triplets (s_A, s_B, s_C) with pairwise correlation $r > 0.80$. After habituation, the overlap between the tags of odors s_B and s_C generally decreases, making them easier to discriminate.

Analysis of mixture representations posthabituation. Our empirical results are supported by theoretical results. Recall that the mixture $s'' = \xi s + (1 - \xi)s'$. Detecting the minor component of a mixture depends on the concentrations of the two components (determined by ξ) and the values of α (habituation rate) and β (dishabituation rate). When the concentration of the habituated odor s within the mixture surpasses a threshold, KC responses to the mixture are suppressed.

Lemma 2: If $\xi \geq \alpha/(\alpha + \beta)$ and

$$\|s\|_{\infty}, \|s'\|_{\infty} \leq \frac{\alpha + \beta}{c\beta} \cdot \tau_o,$$

then $I(s''); s = \emptyset$, where $I(s''); s$ is the tag associated with s'' after habituation to s .

Proof: See [SI Appendix, Lemma S1](#). $\square$

On the other hand, the tag for the mixture resembles the tag for the unhabituated component when $\xi < \alpha/(\alpha + \beta)$, which is 83% for $\alpha = 0.05$, $\beta = 0.01$. See [SI Appendix, Lemma S3](#).

Application to Online Similarity Search. The two problems above exemplify why only encoding the change of an input with respect to data observed in the recent past may better highlight outliers or novel components of the input in streaming data. We use this motivation to present a twist on the classic similarity search problem (76–81) ([SI Appendix](#)).

We calculated the tags $I(s)$ and $I(s')$ of two odors, and the tags of their mixture $I(s'')$ and $I(s''); s$ before and after habituation, respectively. We then calculated the top N nearest neighbors for each tag ([Methods](#)). Before habituation, $I(s'')$ and $I(s)$ share a large portion of nearest neighbors (median mAP = 0.81 ± 0.08 , over all pairs s, s'), as expected since s'' con-

tains 80% of s . However, after habituation, this overlap decreased (median mAP = 0.19 ± 0.16), and, instead, the mixture shares almost the same nearest neighbors as odor s' (median mAP = 0.84 ± 0.22 ; [Fig. 4C](#)). These results suggest that, after habituation, the representation assigned for the mixture allows the retrieval of nearest neighbors that are relevant to pure odor s' , even though s' represents only a small component of the input. In [SI Appendix](#), we discuss additional applications where online similarity searches may be beneficial.

Discussion

We developed and analyzed an online unsupervised learning algorithm for short-term odor habituation based on the “negative image” mechanism in the fruit fly olfactory circuit (16, 40). Our model can replicate three important aspects of habituation: decreased activity to repeated, neutral stimulus; stimulus specificity; and reversibility (dishabituation). The algorithm demonstrates how habituation can filter out background signals, making it easier to discriminate between similar odors and detect novel components in odor mixtures, even if these components have relatively low concentration. Our results are consistent with previous experimental studies of habituation and potentially lay out a framework for understanding the effects of habituation on downstream sensory coding, behavior, and cognition in other systems and species (24, 82, 83).

A critical component of the habituation algorithm is the use of rectilinear thresholding by KCs, which filters noise prior to the “winner take all” calculation. After habituating to odor A , there will likely still be small nonzero PN activity in response to A . Without a rectilinear function, this noise will persist in the KCs, and may not be eliminated by APL because the winner-take-all mechanism adapts the amount of feed-back




Fig. 4. Habituation, foreground discrimination, and improved similarity search. After habituation to the dominant component in a binary mixture, the tag of the mixture becomes more similar to the minor component. (A) Each odor mixture contains two odors, a dominant odor ($\xi = 80\%s$) and a minor odor ($1 - \xi = 20\%s'$). Before habituation, the KC tag of the mixture has higher overlap with odor s than with odor s' . After habituation, the tag of the mixture almost completely overlaps with that of odor s' , suggesting that information of odor s has been “subtracted” from the mixture. Odor s : acetic acid. Odor s' : ethyl hexanoate. Shaded circles illustrate ORN firing rates. (B) Bar plot comparing the similarity of the tags between odor s' and the mixture s'' , before and after habituation to s . Approximately 6,000 mixtures are tested. Before habituation, the mixture has low overlap with s' , but, after habituation, the mixture has high overlap due to the s' component “popping out” above the background (P value < 0.001 , paired t test). (C) *Left* compares the similarity of nearest neighbors of odor s with the mixture before and after habituation. *Right* shows similar results of odor s' with the mixture. After habituation, the predicted nearest neighbors of the mixture are more similar to those of odor s' than those of s . The height of the bars shows the median value, and the error bars show median absolute deviation (both P values < 0.001 , paired t test).

inhibition it provides based on the amount of feed-forward excitation it receives (28). In contrast, a rectilinear KC function will apply hard thresholding to suppress KC noise prior to the winner-take-all mechanism and the selection of the tag. This is important to reduce the chance that noise due to habituation activates KCs and triggers downstream behavioral responses.

Our analysis raises an interesting circuit design question: Why not have the locus of habituation be directly at the ORN $\rightarrow$ PN synapse, implemented via anti-Hebbian plasticity (84), as opposed to using an intermediary inhibitory neuron (LN1)? Two reasons come to mind. First, dishabituation can quickly occur by inhibiting the LN1, without requiring any change in the LN1 $\rightarrow$ PN synaptic weight (40). This could be important if, for example, the background suddenly becomes important and needs to be attended to. Second, habituation should mostly occur for neutral stimuli, as opposed to odors with a strong valence (e.g., danger). Inhibiting LN1 for important odors provides a simple mechanism to avoid undesirable habituation. Indeed, there are several neuromodulators targeting the antennae lobe (16), and some of them may serve these functions.

In addition to PN habituation, there are various mechanisms of ORN adaptation (e.g., the Weber–Fechner law) that help preserve information about odor timing and odor identity, especially in turbulent environments (85–88). Experimental results have shown that 1) ORN responses are reduced after repeated odor exposure, 2) ORN responses to an odor can change after prior exposure to other odors, and 3) ORNs use the timing of firing to encode complex odor environments (89, 90). ORN adaptation occurs at a much faster time scale (~ 100 ms) than LN1–PN plasticity (~ 30 min), which we study here. While ORN adaptation certainly improves the robustness of odor coding in noisy environments, it does not fully explain the effects of short-term habituation, for three reasons. First, Das et al.

(16) showed that GABAergic local interneurons and GABA-A receptors in PNs are required for short-term habituation (91). Second, ORN adaptation, which initiates after milliseconds, also recovers on a much faster time scale than short-term habituation (92). And third, in a simple computational model, we show that a basic model of ORN adaptation is not as odor specific as PN adaptation (*SI Appendix, Fig. S6*), again making ORN adaptation unlikely to be the sole mechanism driving short-term habituation. Nonetheless, combining ORN and PN adaptation to form a more comprehensive habituation algorithm is a natural next step.

The algorithm proposed here can be improved in several ways, while potentially capturing more biological detail. First, we assumed a binary KC tag, whereas, in reality, active KCs fire at varied (albeit similar) rates (47, 66); including firing rates for KCs in our model may further increase the ability to discriminate odors. Second, we used a simplified model of the antennae lobe, without considering, for example, multiple LN1 neurons, more elaborate LN1–PN interactions, or timing aspects of LN inhibition (93). Including these complexities, once mapped anatomically and physiologically, may lead to a better understanding of cross-odor habituation. Third, including mechanisms for both short-term and long-term habituation (25) may reveal how habituation can operate on multiple time scales simultaneously. Fourth, the antennae lobe sends signals to the mushroom body as well as to the lateral horn, a region of the *Drosophila* brain believed to be responsible for innate behaviors (94). As a result, habituation that originates in the antennal lobe may affect both learned behavior and innate behavior in fruit flies.

While not explicitly our goal here, translating principles of neural computation to improve machine learning is a long-sought goal at the interface of computer science and

neuroscience (39, 95–97). Specifically, these insights may be applicable for deep learning (98) [e.g., to enhance attention-modulated artificial neural networks (99, 100)], for online similarity search problems, and in electronic noses for event-based smell processing. Moreover, recent theoretical work has argued that deep neural networks particularly excel in overcoming real-world “nuisances”, that is, task-irrelevant variation in inputs such as translations, rotations, and deformations (101, 102). Habituation may further support this advance by reducing background nuisances.

Data Availability

Code for the habituation algorithms is available at GitHub, <https://github.com/aspen-shen/Habituation-as-a-neural-algorithm-for-online-odor-discrimination>.

ACKNOWLEDGMENTS. We thank Mani Ramaswami, Jaiyam Sharma, and Shyam Srinivasan for helpful discussions and comments on the manuscript. S.N. was supported by the Pew Charitable Trusts, the National Institutes of Health under Awards 1R01DC017695 and 1U1NS111692, and funding from the Simons Center for Quantitative Biology at Cold Spring Harbor Laboratory. Y.S. was supported by a Swartz Foundation Fellowship.

1. D. A. Wilson, C. Linster, Neurobiology of a simple memory. *J. Neurophysiol.* **100**, 2–7 (2008).
2. G. Vivanti *et al.*, Attention to novelty versus repetition: Contrasting habituation profiles in Autism and Williams syndrome. *Dev. Cogn. Neurosci.* **29**, 54–60 (2018).
3. J. F. Cavanagh, P. Kumar, A. A. Mueller, S. P. Richardson, A. Mueen, Diminished EEG habituation to novel events effectively classifies Parkinson's patients. *Clin. Neurophysiol.* **129**, 409–418 (2018).
4. S. Marsland, Using habituation in machine learning. *Neurobiol. Learn. Mem.* **92**, 260–266 (2009).
5. M. Sha'abani *et al.*, A habituation based approach for detection of visual changes in surveillance camera *AIP Conf. Proc.* **1883**, 020038.
6. S. Marsland, U. Nehmzow, J. Shapiro, Novelty detection on a mobile robot using habituation. arXiv preprint cs/0006007 (2 June 2000).
7. E. Özbilge, On-line expectation-based novelty detection for mobile robots. *Robot. Autonom. Syst.* **81**, 33–47 (2016).
8. Y. Hu, K. Sirlantzis, G. Howells, N. Ragot, P. Rodriguez, An online background subtraction algorithm deployed on a NAO humanoid robot based monitoring system. *Robot. Autonom. Syst.* **85**, 37–47 (2016).
9. C. Kim, J. Lee, T. Han, Y. M. Kim, A hybrid framework combining background subtraction and deep neural networks for rapid person detection. *J. Big Data* **5**, 22 (2018).
10. M. Markou, S. Singh, Novelty detection: A review—part 2: Neural network based approaches. *Signal Process.* **83**, 2499–2521 (2003).
11. D. L. Wang, M. A. Arbib, How does the toad's visual system discriminate different worm-like stimuli? *Biol. Cybern.* **64**, 251–261 (1991).
12. D. Wang, M. A. Arbib, Modeling the dishabituation hierarchy: The role of the primordial hippocampus. *Biol. Cybern.* **67**, 535–544 (1992).
13. T. Kohonen, *Self-Organization and Associative Memory* (Springer-Verlag, Berlin, Germany, ed. 3, 1989).
14. B. W. Stiles, J. Ghosh, Habituation based neural networks for spatio-temporal classification. *Neurocomputing* **15**, 273–307 (1997).
15. A. I. Weber, K. Krishnamurthy, A. L. Fairhall, Coding principles in adaptation. *Annu. Rev. Vis. Sci.* **5**, 427–449 (2019).
16. S. Das *et al.*, Plasticity of local GABAergic interneurons drives olfactory habituation. *Proc. Natl. Acad. Sci. U.S.A.* **108**, E646–E654 (2011).
17. P. Szyszka, J. S. Stierle, S. Biergens, C. G. Galizia, The speed of smell: Odor-object segregation within milliseconds. *PLoS One* **7**, e36096 (2012).
18. S. Haney, D. Saha, B. Raman, M. Bazhenov, Differential effects of adaptation on odor discrimination. *J. Neurophysiol.* **120**, 171–185 (2018).
19. C. Linster, L. Henry, M. Kadohisa, D. A. Wilson, Synaptic adaptation and odor-background segmentation. *Neurobiol. Learn. Mem.* **87**, 352–360 (2007).
20. C. Linster, A. V. Menon, C. Y. Singh, D. A. Wilson, Odor-specific habituation arises from interaction of afferent synaptic adaptation and intrinsic synaptic potentiation in olfactory cortex. *Learn. Mem.* **16**, 452–459 (2009).
21. J. Y. Chen *et al.*, Learning modifies odor mixture processing to improve detection of relevant components. *J. Neurosci.* **35**, 179–197 (2015).
22. F. F. Locatelli *et al.*, Nonassociative plasticity alters competitive interactions among mixture components in early olfactory processing. *Eur. J. Neurosci.* **37**, 63–79 (2013).
23. D. Rokni, V. Hemmelder, V. Kapoor, V. N. Murthy, An olfactory cocktail party: Figure-ground segregation of odorants in rodents. *Nat. Neurosci.* **17**, 1225–1232 (2014).
24. H. K. Kato, M. W. Chu, J. S. Isaacson, T. Komiyama, Dynamic sensory representations in the olfactory bulb: Modulation by wakefulness and experience. *Neuron* **76**, 962–975 (2012).
25. D. Chaudhury *et al.*, Olfactory bulb habituation to odor stimuli. *Behav. Neurosci.* **124**, 490–499 (2010).
26. G. C. Turner, M. Bazhenov, G. Laurent, Olfactory representations by *Drosophila* mushroom body neurons. *J. Neurophysiol.* **99**, 734–746 (2008).
27. T. Hige, Y. Aso, M. N. Modi, G. M. Rubin, G. C. Turner, Heterosynaptic plasticity underlies aversive olfactory learning in *Drosophila*. *Neuron* **88**, 985–998 (2015).
28. 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).
29. R. A. Campbell *et al.*, Imaging a population code for odor identity in the *Drosophila* mushroom body. *J. Neurosci.* **33**, 10568–10581 (2013).
30. S. J. Caron, V. Ruta, L. F. Abbott, R. Axel, Random convergence of olfactory inputs in the *Drosophila* mushroom body. *Nature* **497**, 113–117 (2013).
31. Y. Aso *et al.*, Mushroom body output neurons encode valence and guide memory-based action selection in *Drosophila*. *Elife* **3**, e04580 (2014).
32. Y. Aso *et al.*, The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife* **3**, e04577 (2014).
33. S. Y. Takemura *et al.*, A connectome of a learning and memory center in the adult *Drosophila* brain. *Elife* **6**, e26975 (2017).
34. Z. Zheng *et al.*, A complete electron microscopy volume of the brain of adult *Drosophila melanogaster*. *Cell* **174**, 730–743 (2018).
35. R. F. Thompson, W. A. Spencer, Habituation: A model phenomenon for the study of neuronal substrates of behavior. *Psychol. Rev.* **73**, 16–43 (1966).
36. D. L. Glanzman, Habituation in *Aplysia*: The Cheshire cat of neurobiology. *Neurobiol. Learn. Mem.* **92**, 147–154 (2009).
37. T. D. Gover, T. W. Abrams, Insights into a molecular switch that gates sensory neuron synapses during habituation in *Aplysia*. *Neurobiol. Learn. Mem.* **92**, 155–165 (2009).
38. A. C. Giles, C. H. Rankin, Behavioral and genetic characterization of habituation using *Caenorhabditis elegans*. *Neurobiol. Learn. Mem.* **92**, 139–146 (2009).
39. C. Pehlevan, D. B. Chklovskii, Neuroscience-inspired online unsupervised learning algorithms: Artificial neural networks. *IEEE Signal Process. Mag.* **36**, 88–96 (2019).
40. M. Ramaswami, Network plasticity in adaptive filtering and behavioral habituation. *Neuron* **82**, 1216–1229 (2014).
41. S. Dasgupta, C. F. Stevens, S. Navlakha, A neural algorithm for a fundamental computing problem. *Science* **358**, 793–796 (2017).
42. T. Bouwmans, S. Javed, M. Sultana, S. K. Jung, Deep neural network concepts for background subtraction: A systematic review and comparative evaluation. *Neural Network* **117**, 8–66 (2019).
43. K. Goyal, J. Singhai, Review of background subtraction methods using Gaussian mixture model for video surveillance systems. *Artif. Intell. Rev.* **50**, 241–259 (2018).
44. B. Malnic, J. Hirono, T. Sato, L. B. Buck, Combinatorial receptor codes for odors. *Cell* **96**, 713–723 (1999).
45. C. F. Stevens, What the fly's nose tells the fly's brain. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 9460–9465 (2015).
46. E. A. Hallem, J. R. Carlson, Coding of odors by a receptor repertoire. *Cell* **125**, 143–160 (2006).
47. S. R. Olsen, V. Bhandawat, R. I. Wilson, Divisive normalization in olfactory population codes. *Neuron* **66**, 287–299 (2010).
48. K. Asahina, M. Louis, S. Piccinotti, L. B. Vosshall, A circuit supporting concentration-invariant odor perception in *Drosophila*. *J. Biol.* **8**, 9 (2009).
49. C. M. Root *et al.*, A presynaptic gain control mechanism fine-tunes olfactory behavior. *Neuron* **59**, 311–321 (2008).
50. V. Bhandawat, S. R. Olsen, N. W. Gouwens, M. L. Schlieff, R. I. Wilson, Sensory processing in the *Drosophila* antennal lobe increases reliability and separability of ensemble odor representations. *Nat. Neurosci.* **10**, 1474–1482 (2007).
51. E. Gruntman, G. C. Turner, Integration of the olfactory code across dendritic claws of single mushroom body neurons. *Nat. Neurosci.* **16**, 1821–1829 (2013).
52. H. Li, Y. Li, Z. Lei, K. Wang, A. Guo, Transformation of odor selectivity from projection neurons to single mushroom body neurons mapped with dual-color calcium imaging. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 12084–12089 (2013).
53. R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, H. S. Seung, Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature* **405**, 947–951 (2000).
54. X. Glorot, A. Bordes, Y. Bengio, “Deep sparse rectifier neural networks” in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, G. Gordon, D. Dunson, M. Dudik, Eds. (Proceedings of Machine Learning Research, Fort Lauderdale, FL, 2011), vol. 15, pp. 315–323.
55. X. Liu, R. L. Davis, The GABAergic anterior paired lateral neuron suppresses and is suppressed by olfactory learning. *Nat. Neurosci.* **12**, 53–59 (2009).
56. N. Y. Masse, G. C. Turner, G. S. Jefferis, Olfactory information processing in *Drosophila*. *Curr. Biol.* **19**, R700–R713 (2009).
57. L. N. Groschner, G. Miesenböck, Mechanisms of sensory discrimination: Insights from *Drosophila* olfaction. *Annu. Rev. Biophys.* **48**, 209–229 (2019).
58. H. Amin, A. C. Lin, Neuronal mechanisms underlying innate and learned olfactory processing in *Drosophila*. *Curr. Opin. Insect Sci.* **36**, 9–17 (2019).
59. H. C. Barron, T. P. Vogels, T. E. Behrens, M. Ramaswami, Inhibitory engrams in perception and memory. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 6666–6674 (2017).
60. R. Okada, T. Awasaki, K. Ito, Gamma-aminobutyric acid (GABA)-mediated neural connections in the *Drosophila* antennal lobe. *J. Comp. Neurol.* **514**, 74–91 (2009).
61. N. K. Tanaka, K. Endo, K. Ito, Organization of antennal lobe-associated neurons in adult *Drosophila melanogaster* brain. *J. Comp. Neurol.* **520**, 4067–4130 (2012).
62. Y. H. Chou *et al.*, Diversity and wiring variability of olfactory local interneurons in the *Drosophila* antennal lobe. *Nat. Neurosci.* **13**, 439–449 (2010).
63. I. Twick, J. A. Lee, M. Ramaswami, Olfactory habituation in *Drosophila*-odor encoding and its plasticity in the antennal lobe. *Prog. Brain Res.* **208**, 3–38 (2014).

64. T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, W. Gerstner, Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* **334**, 1569–1573 (2011).
65. W. Li, E. Luxenberg, T. Parrish, J. A. Gottfried, Learning to smell the roses: Experience-dependent neural plasticity in human piriform and orbitofrontal cortices. *Neuron* **52**, 1097–1108 (2006).
66. Dasgupta S., Sheehan T. C., Stevens C. F., Navlakha S. (2018) A neural data structure for novelty detection. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 13093–13098.
67. D. Oswald, S. Waddell, Olfactory learning skews mushroom body output pathways to steer behavioral choice in *Drosophila*. *Curr. Opin. Neurobiol.* **35**, 178–184 (2015).
68. P. Cognigni, J. Felsenberg, S. Waddell, Do the right thing: Neural network mechanisms of memory formation, expression and update in *Drosophila*. *Curr. Opin. Neurobiol.* **49**, 51–58 (2018).
69. K. Shen, S. Tootoonian, G. Laurent, Encoding of mixtures in a simple olfactory system. *Neuron* **80**, 1246–1262 (2013).
70. A. A. M. Mohamed et al., Odor mixtures of opposing valence unveil inter-glomerular crosstalk in the *Drosophila* antennal lobe. *Nat. Commun.* **10**, 1201 (2019).
71. P. Giraudet, F. Berthommier, M. Chaput, Mitral cell temporal response patterns evoked by odor mixtures in the rat olfactory bulb. *J. Neurophysiol.* **88**, 829–838 (2002).
72. R. Tabor, E. Yaksi, J. M. Weislogel, R. W. Friedrich, Processing of odor mixtures in the zebrafish olfactory bulb. *J. Neurosci.* **24**, 6611–6620 (2004).
73. V. Singh, N. R. Murphy, V. Balasubramanian, J. D. Mainland, Competitive binding predicts nonlinear responses of olfactory receptors to complex mixtures. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 9598–9603 (2019).
74. Y. Lin, R. Jin, D. Cai, S. Yan, X. Li, “Compressed hashing” in 2013 IEEE Conference on Computer Vision and Pattern Recognition (Institute of Electrical and Electronics Engineers, 2013), pp. 446–451.
75. K. S. Honegger, 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).
76. P. Indyk, R. Motwani, Approximate nearest neighbors: Towards removing the curse of dimensionality in *Proceedings of the 30th Annual ACM Symposium on Theory of Computing, STOC’98* (Association for Computing Machinery, New York, NY 1998), pp. 604–613.
77. A. Andoni, P. Indyk, Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM* **51**, 117–122 (2008).
78. M. Datar, N. Immorlica, P. Indyk, V. S. Mirrokni, “Locality-sensitive hashing scheme based on p-stable distributions” in *Proceedings of the Twentieth Annual Symposium on Computational Geometry, SCG ’04* (Association for Computing Machinery, New York, NY, 2004), pp. 253–262.
79. A. Gionis, P. Indyk, R. Motwani, “Similarity search in high dimensions via hashing” in *Proceedings of the 25th International Conference on Very Large Data Bases, VLDB’99*, M. P. Atkinson, M. E. Orlowska, P. Valduriez, S. B. Zdonik, M. L. Brodie, Eds. (Morgan Kaufmann, San Francisco, CA, 1999), pp. 518–529.
80. J. Wang, H. T. Shen, J. Song, J. Ji, Hashing for similarity search: A survey. *arXiv:1408.2927* (13 August 2014).
81. J. Wang, T. Zhang, J. Song, N. Sebe, H. Shen, A survey on learning to hash. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**, 769–790 (2018).
82. M. C. Ogg, J. M. Ross, M. Bendahmane, M. L. Fletcher, Olfactory bulb acetylcholine release dishabituates odor responses and reinstates odor investigation. *Nat. Commun.* **9**, 1868 (2018).
83. S. Nizampatnam, D. Saha, R. Chandak, B. Raman, Dynamic contrast enhancement and flexible odor codes. *Nat. Commun.* **9**, 3062 (2018).
84. C. Dempsey, L. F. Abbott, N. B. Sawtell, Generalization of learned responses in the mormyrid electrosensory lobe. *Elife* **8**, e44032 (2019).
85. S. Gorur-Shandilya, M. Demir, J. Long, D. A. Clark, T. Emonet, Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli. *Elife* **6**, e27670 (2017).
86. N. Kadakia, T. Emonet, Front-end Weber-Fechner gain control enhances the fidelity of combinatorial odor coding. *eLife* **8**, e45293 (2019).
87. S. L. Brown, J. Joseph, M. Stopfer, Encoding a temporally structured stimulus with a temporally structured neural representation. *Nat. Neurosci.* **8**, 1568–1576 (2005).
88. B. Raman, J. Joseph, J. Tang, M. Stopfer, Temporally diverse firing patterns in olfactory receptor neurons underlie spatiotemporal neural codes for odors. *J. Neurosci.* **30**, 1994–2006 (2010).
89. D. Saha et al., A spatiotemporal coding mechanism for background-invariant odor recognition. *Nat. Neurosci.* **16**, 1830–1839 (2013).
90. C. D. Wilson, G. O. Serrano, A. A. Koulakov, D. Rinberg, A primacy code for odor identity. *Nat. Commun.* **8**, 1477 (2017).
91. A. Larkin et al., Central synaptic mechanisms underlie short-term olfactory habituation in *Drosophila* larvae. *Learn. Mem.* **17**, 645–653 (2010).
92. F. Zufall, T. Leinders-Zufall, The cellular and molecular basis of odor adaptation. *Chem. Sens.* **25**, 473–481 (2000).
93. K. I. Nagel, R. I. Wilson, Mechanisms underlying population response dynamics in inhibitory interneurons of the *Drosophila* antennal lobe. *J. Neurosci.* **36**, 4325–4338 (2016).
94. M. J. Dolan et al., Neurogenetic dissection of the *Drosophila* lateral horn reveals major outputs, diverse behavioural functions, and interactions with the mushroom body. *eLife* **8**, e43079 (2019).
95. M. Helmstaedter, The mutual inspirations of machine learning and neuroscience. *Neuron* **86**, 25–28 (2015).
96. D. Hassabis, D. Kumaran, C. Summerfield, M. Botvinick, Neuroscience-inspired artificial intelligence. *Neuron* **95**, 245–258 (2017).
97. Y. Hitron, N. Lynch, C. Musco, M. Parter, “Random sketching, clustering, and short-term memory in spiking neural networks” in *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, Leibniz International Proceedings in Informatics (LIPIcs), T. Vidick Ed. (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2020), vol. 151, pp. 23:1–23:31.
98. W. Liu et al., Fast retinomorphic event stream for video recognition and reinforcement learning. *arXiv:1805.06374v2* (19 May 2018).
99. W. Yin, H. Schütze, B. Xiang, B. Zhou, Abcnn: Attention-based convolutional neural network for modeling sentence pairs. *Trans. Assoc. Comput. Linguist.* **4**, 259–272 (2016).
100. O. Vinyals, A. Toshev, S. Bengio, D. Erhan, Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**, 652–663 (2016).
101. A. Achille, S. Soatto, Emergence of invariance and disentanglement in deep representations. *J. Mach. Learn. Res.* **19**, 1947–1980 (2018).
102. A. B. Patel, T. Nguyen, R. G. Baraniuk, “A probabilistic framework for deep learning” in *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, D. D. Lee, U. von Luxburg, R. Garnett, M. Sugiyama, I. Guyon, Eds. (Curran Associates Inc., 2016), pp. 2558–2566.