



Invariant odor recognition with ON–OFF neural ensembles

Srinath Nizampatnam^{a,1} , Lijun Zhang^{a,1}, Rishabh Chandak^{b,1} , James Li^b , and Baranidharan Raman^{a,b,2}

^aDepartment of Electrical and Systems Engineering, Washington University in St. Louis, St. Louis, MO 63130; and ^bDepartment of Biomedical Engineering, Washington University in St. Louis, St. Louis, MO 63130

Edited by John Hildebrand, Department of Neuroscience, The University of Arizona, Tucson, AZ; received November 9, 2020; accepted November 18, 2021

Invariant stimulus recognition is a challenging pattern-recognition problem that must be dealt with by all sensory systems. Since neural responses evoked by a stimulus are perturbed in a multitude of ways, how can this computational capability be achieved? We examine this issue in the locust olfactory system. We find that locusts trained in an appetitive-conditioning assay robustly recognize the trained odorant independent of variations in stimulus durations, dynamics, or history, or changes in background and ambient conditions. However, individual- and population-level neural responses vary unpredictably with many of these variations. Our results indicate that linear statistical decoding schemes, which assign positive weights to ON neurons and negative weights to OFF neurons, resolve this apparent confound between neural variability and behavioral stability. Furthermore, simplification of the decoder using only ternary weights ($\{+1, 0, -1\}$) (i.e., an “ON-minus-OFF” approach) does not compromise performance, thereby striking a fine balance between simplicity and robustness.

olfaction | sensory invariance | behavioral recognition | computational neuroscience | antennal lobe

Robustly recognizing a sensory stimulus is a necessity for the survival and propagation of all animals. Since this capability is demonstrated in all sensory systems, it raises the following question: what is the neural basis that underlies this feat of pattern recognition? Most stimuli are encountered in a multitude of ways in natural environments. Often, stimulus features such as intensity, duration, and recurrence could vary. In addition, external perturbances due to changes in environmental conditions (such as changes in humidity or temperature), the presence of other competing cues, or the temporal context (i.e., when it is received in a stimulus sequence) could also change independently of the variation in stimulus-specific features. An additional degree of interference can arise from changes in the sensory circuit due to plastic changes arising either from prior exposures or co-occurrence with other sensory cues. Given the complexity in carrying out the basic task of recognizing a stimulus, we wondered whether there exists a computational framework that can compensate for all these disparate sources of variation and allow robust recognition of a stimulus. In particular, we sought to examine this issue in the well-studied locust olfactory system (1–13).

In the locust olfactory system, odorants activate olfactory receptor neurons in the antenna. This signal is transmitted downstream to the antennal lobe (analogous to the vertebrate olfactory bulb) where it drives responses in cholinergic projection neurons (PNs) and GABAergic local neurons (LNs). The interaction between PNs and LNs transforms the sensory input received into complex patterns of spiking activities distributed across ensembles of PNs that become the output of the antennal lobe circuit. Prior work has shown that information about the identity and intensity of an odorant is encoded by spatiotemporal PN activity patterns (5). While individual PN responses were perturbed by manipulating stimulus dynamics (11, 14), stimulus history (15, 16), and presence of background chemicals (7), the ensemble neural patterns still allowed recognition of

odorants. Behavioral evidences also support this interpretation and reveal that odorants can be recognized independent of background cues (7) and stimulus history (15).

It is worth noting that prior studies examined neural response variabilities that arose due to each of these perturbations in isolation. In natural contexts, such interferences could occur independently or in conjunction with one another. Could robust odor recognition still be achieved? Would an array of schemes that extract information from a variety of response features be necessary for compensating changes associated with each perturbation? Alternately, can the variable neural responses be decoded in a manner that can simultaneously allow invariant odor recognition independent of all these perturbations? If so, what neural response features would be important for achieving this result? We sought to examine these issues in this study.

Results

Robust Odor Recognition in a Behavioral Assay. We began by testing our core hypothesis that locusts could indeed recognize an odorant in an invariant fashion. To do this, we used an appetitive-conditioning assay (Fig. 1A). In this assay, starved locusts were presented an odorant (conditioned stimulus; CST) followed by a food reward (unconditioned stimulus; UST). The food reward is alone sufficient to evoke an innate extension/opening of sensory appendages close to the mouth of the locusts (called the maxillary palps). After training with six trials when CST and UST were delivered in an overlapping sequence, the ability of the locusts to

Significance

The smell of coffee is the same whether it is smelled in a coffee shop or grocery shop (different backgrounds), on a hot day or a cold day (different ambient conditions), after lunch or dinner (different temporal contexts), or using a deep inhalation or normal inhalation (different stimulus dynamics). This feat of pattern recognition that is still difficult to achieve in artificial chemical sensing systems is performed by most sensory systems for their survival. How is this capability achieved? We explored this issue. We found that there are two orthogonal ensembles of neurons, one activated during stimulus presence (ON neurons) and one activated after its termination (OFF neurons), and both contribute to this important computation in a complementary fashion.

Author contributions: B.R. designed research; S.N., L.Z., R.C., and J.L. performed research; S.N., L.Z., R.C., and B.R. analyzed data; and S.N., L.Z., R.C., and B.R. 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/).

¹S.N., L.Z., and R.C. contributed equally to this work.

²To whom correspondence may be addressed. Email: barani@wustl.edu.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2023340118/-/DCSupplemental>.

Published January 7, 2022.

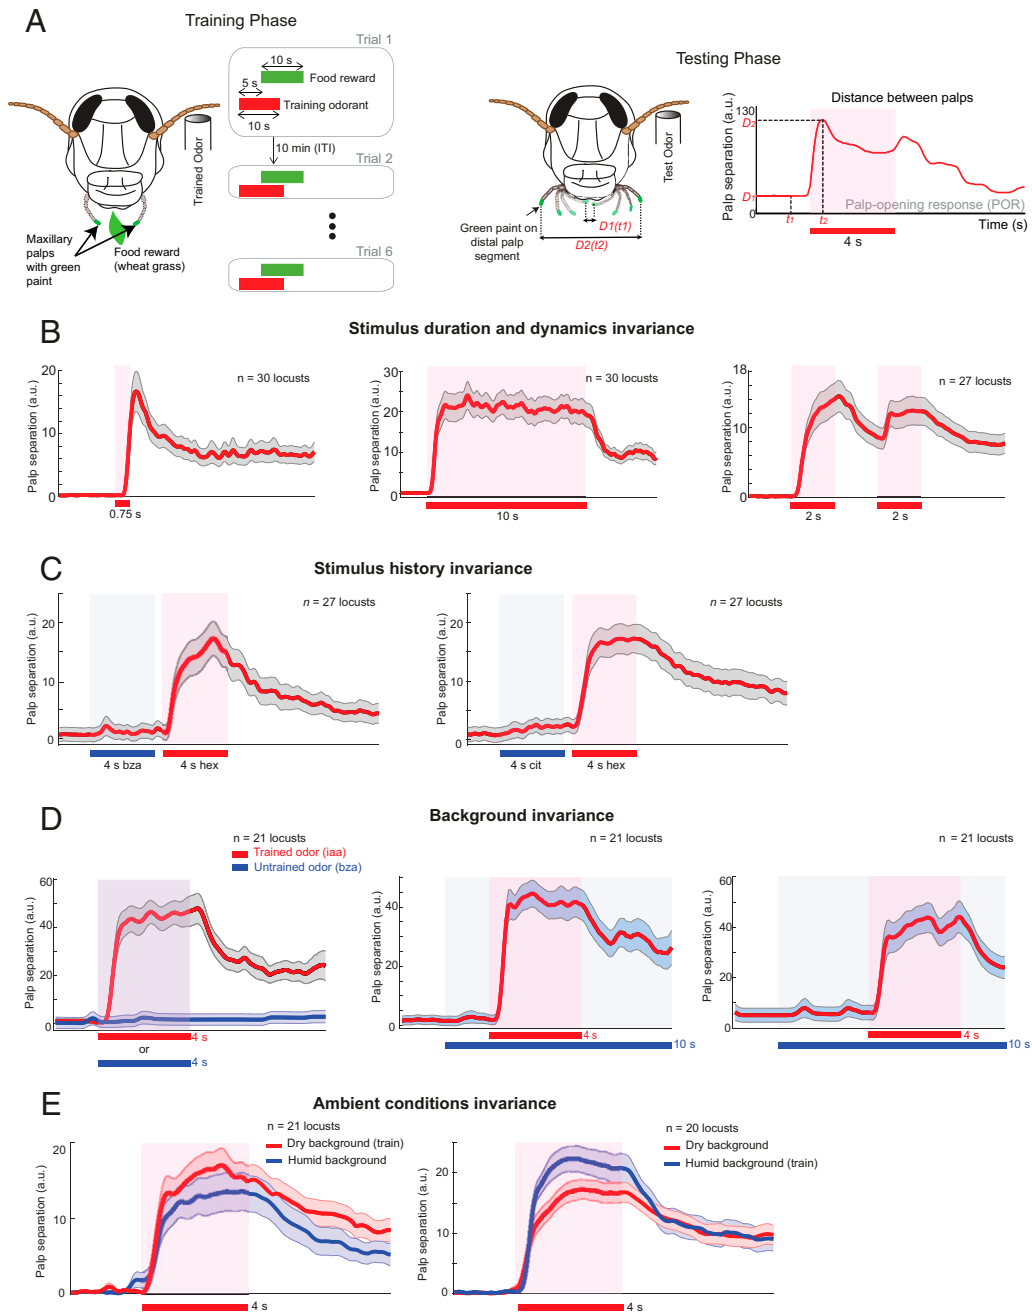


Fig. 1. Invariant odor recognition in an appetitive-conditioning assay. (A) Schematic showing the training and testing protocols followed for the POR assays. Briefly, starved locusts were presented six trials of the training odorant (CST) followed by a food reward (UST). The odorant pulse was 10 s in duration, and the food reward was presented 5 s following the onset of the conditioning stimulus. Locusts that accepted food reward in four of the six training trials were regarded as successfully “trained locusts,” and their PORs were evaluated in an unrewarded testing phase. Selective opening of their maxillary palps (sensory appendage close to the mouth) upon onset of the conditioning stimulus was regarded as successful recognition of the training odorant. The PORs were quantified by painting and tracking the distal ends of the palps (*Methods*). A sample POR trajectory is shown on the *Right* where the training odor presentation (color bar) led to increased palp separation, indicating a POR. (B) Locusts were trained using hex as the CST and were tested using hex pulses that varied in their duration. The mean PORs (\pm SEM) of all trained locusts are shown for test pulses of 0.75-s ($n = 30$ locusts) and 10-s ($n = 30$ locusts) durations, and a sequence of two pulses that was 2 s ON–2 s gap–2 s ON ($n = 27$ locusts). The color bar indicates when odor was presented. (C) Locusts were trained using hex as the CST and tested by presenting hex pulses in a nonoverlapping sequence following the termination of a distractor odorant (i.e., introducing variations in stimulus history). Trained locusts first encountered a distractor odor pulse for 4 s, followed by a 0.5-s gap, that was then followed by a 4-s pulse of hex during the testing phase. The mean responses (\pm SEM) of locusts ($n = 27$) tested with hex presentations following two distractor odors (bza and cit) are shown. The color bars indicate when odors were presented. Blue bars indicate the time periods when the distractor was presented, and the red bars indicate duration of exposure to the trained odorant. (D) Locusts were trained using iaa as the CST and tested by presenting iaa atop a background odorant (bza). (*Left*) The mean PORs to solitary presentations of iaa and bza are shown. Note that only the CST evokes POR responses, whereas the bza introductions did not elicit any detectable POR responses. (*Middle and Right*) PORs (mean \pm SEM; $n = 21$ locusts) are shown during presentations of iaa that was introduced 2 s and 4 s after the onset of a sustained bza pulse (i.e., the background odorant). The color bars indicate when odors were presented and how they overlapped. (E, *Left*) Locusts trained with hex in a dry background (0% RH) were tested for hex responses in dry (red) and humid (100% RH, blue) conditions (*Methods*). The mean PORs of trained locusts ($n = 21$) are shown. The color bar at the bottom indicates when odor was presented. The shaded regions indicate the SEM. (*Right*) Similar plots as in the *Left* panel, but locusts trained for hex in humid (100% RH) conditions are shown ($n = 20$ locusts).

recognize the CST was examined in an unrewarded testing phase. Opening of the maxillary palps following presentation of the CST was regarded as an indicator of successful recognition of the trained odorant. Note that the palp opening response (POR) was selective to the CST (with a few caveats; see Fig. 2*A*). Further, to make the read-out quantitative, locust palps were painted with a nonodorous green paint, and the distance between the palps was tracked as a function of time (Fig. 1*A, Right*). Notably, the PORs remained consistent when probed multiple times with the CST in the unrewarded test phase, thereby allowing us to examine the POR when we made a battery of perturbations.

First, we examined how PORs changed as we varied the duration of the stimulus. We found that the PORs initiated rapidly, and the palps were kept open for the duration of the odor pulse and terminated following cessation of the trained odorant. Although we trained locusts using a particular duration of CST pulse (4-s pulse of hexanol [hex] for results shown in Fig. 1*B*), we found that POR duration was briefer for a shorter CST pulse, and the palps remained open for the entire duration of a longer CST pulse. These results indicate that locusts recognized the CST and maintained their responses to the trained odorant independent of the stimulus duration during training. Further, when two short, nonoverlapping pulses of the CST were presented in quick succession, the locust palps opened, began closing, and again opened matching the dynamics of the stimulus delivery (Fig. 1*B, Right*). The fact that the second CST pulse in the stimulus sequence elicited a POR comparable to the first pulse indicates that locusts could respond to the trained odorant robustly independent of variations in the stimulus delivery/encounters.

Could response to a stimulus change depending on what other cues were encountered recently? To understand this, we presented the CST in different nonoverlapping sequences with a number of distractor cues. Note that the distractor cues terminate before the onset of CST and are only used to determine whether stimulus history can alter recognition performance. Our results indicate that locusts could robustly recognize the trained odorant irrespective of the stimulus that was encountered before (Fig. 1*C*).

Next, we wondered whether locusts trained to recognize a particular odorant could do so independent of the presence of other competing cues (Fig. 1*D*). Note that locusts trained with a CST had a POR only when tested with the trained odorant (isoamyl acetate or iaa) and had no detectable POR response following the presentation of an untrained odorant (benzaldehyde or bzald). Presentation of the trained odorant (iaa) atop the untrained background cue (bzald) with different latencies did not alter the locust POR response to the trained odorant. In all the cases, a rapid and vigorous POR response was observed following the introductions of the CST, and the palps started closing upon termination of the CST. Similar results were also reported when locusts trained with hex were tested by presenting hex alone or atop a background cue (7). Taken together, these results indicate that the locusts could recognize the trained odorant in a background-invariant manner.

Could changes in ambient conditions impact recognition performance? To understand that, we trained locusts in dry conditions (0% relative humidity [RH]). In the testing phase, we examined the ability of locusts to recognize the CST presented either in dry or humid (100% RH) conditions. Our results show that locusts opened their palps to all the introductions of the CST in both dry and humid conditions (Fig. 1*E, Left*). The performance was near identical, indicating a robust odor recognition that was invariant with respect to changes in ambient conditions. Similar results were also obtained when locusts were trained in humid conditions and tested in both dry and humid conditions (Fig. 1*E, Right*). These results indicate that locusts can recognize trained odorants independent of changes in ambient conditions.

The locust recognition performance under the battery of perturbations discussed is summarized in Fig. 2. Taken together, these results support the idea that locusts could recognize an odorant independent of variations in stimulus features such as its duration and dynamics and extrinsic features such as encounters with other distractor cues, presence of other competing cues, or changes in ambient conditions (Fig. 2*B–E*). Furthermore, while locust responses were selective and the CST evoked the strongest response, locusts trained with one odorant also showed PORs to a select few other odorants (generalization; Fig. 2*A*). These observations raise several questions regarding whether certain odor-evoked neural response features remain robust to such perturbations to allow invariant odor recognition and whether achieving robust odor recognition also causes behavioral responses to generalize across odorants.

Stimulus Dynamics, History, and Competing Cues Induce Variations in PN Responses. How were odor-evoked responses of individual PNs in the locust antennal lobe perturbed? To understand this, we designed a stimulation protocol that presented two “target odorants,” hex and iaa, in a pulsatile fashion (Fig. 3*A* and *B*). Note that the two target odorants (hex and iaa) were also used as CST in the behavioral experiments. These target odor introductions were of different durations with varying interstimulus intervals or presented atop a background cue (bzald) or following a distractor odorant (citral; cit). A photoionization detector was used to characterize stimulus dynamics achieved using this delivery protocol (*SI Appendix, Fig. 1*).

We recorded responses of 89 PNs in the locust antennal lobe ($n = 25$ locusts). First, we examined the ability of individual PNs to robustly encode the identity of two “target” odorants. In general, we found that most individual PNs had robust and reliable responses during certain exposures of the target odorant but not all. In the entire ensemble of PNs that was recorded ($n = 89$ in total), we found that four PNs had a detectable response to all encounters of the target odorants (Fig. 3*A* and *B*; PN6). Since these “reliable” PNs were activated by both the target odorants (hex and iaa) and their response intensity and spiking patterns varied considerably across pulses (*SI Appendix, Fig. 2*), we note that these PNs individually did not provide discrimination between these two odorants.

To quantify the response variability observed at the level of individual PNs, we computed correlations between the PN response to the first pulse of the target odorant and all the other introductions of the same chemical (Fig. 3*C*). A high correlation value would indicate that PN firing-rate patterns remained consistent across different pulses. However, the computed distribution of the PN response correlations in our dataset revealed that spiking activities during different encounters of the target odorant had only a weak pattern match with the responses elicited during the very first encounter of an odorant (Fig. 3*D* and *E*). Note that even for those PNs that had detectable responses across different pulses (PN6 in Fig. 3*A*), the mean correlations were low, as the spike patterns across the different target odor pulses were not consistent.

Furthermore, we computed the ratio of mean spike counts across the 11 pulses of the same odorant (and across the 10 trials) with variance in spike counts (i.e., Fano factor; Fig. 3*F* and *G*). A Fano factor of one indicates Poisson variability. As can be noted, most PNs had a supra-Poisson variability. Taken together, these results indicate that individual PN responses vary considerably and may not provide a reliable read-out of odor identity across diverse conditions.

Variations Due to Changes in Ambient Conditions. Next, we examined whether changes in humidity conditions would further exacerbate the problem of robustly encoding odorant identity. For this purpose, we used the same stimulus delivery protocol

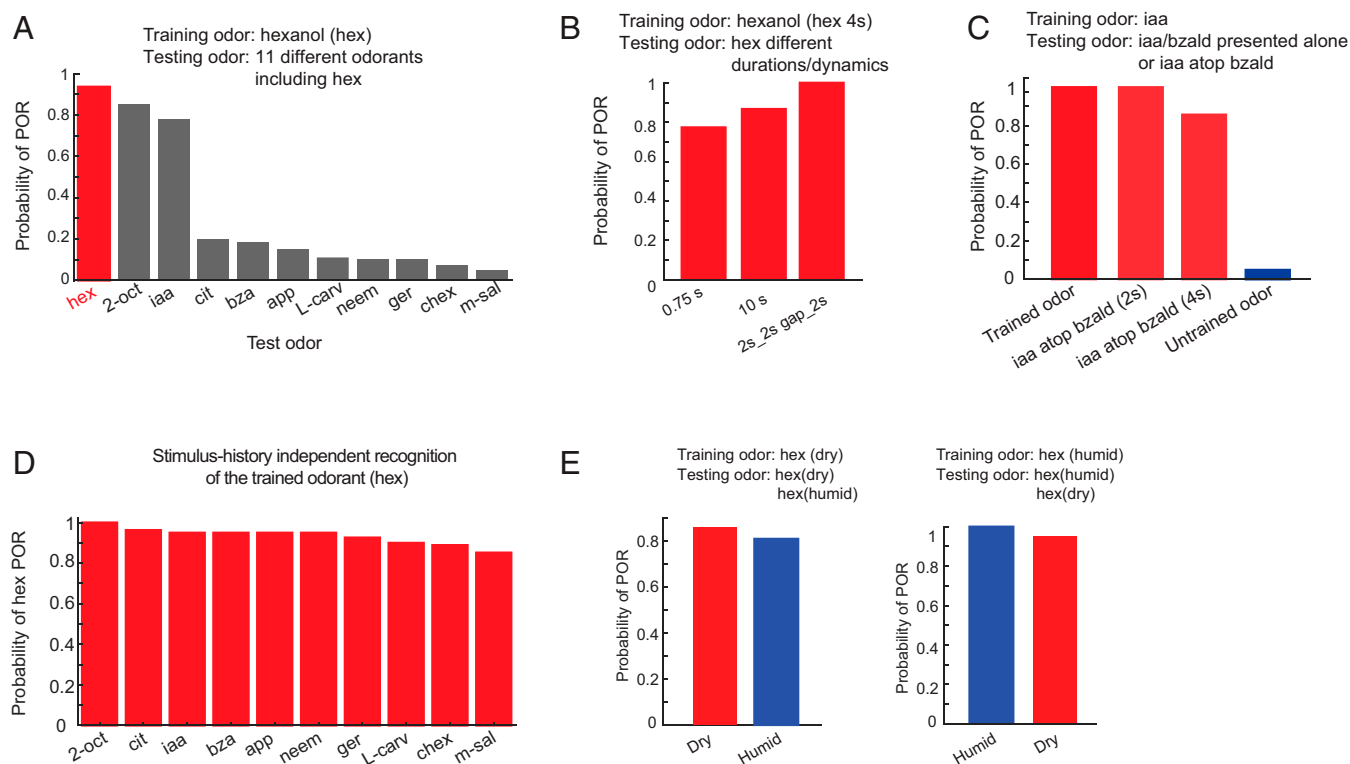


Fig. 2. Summary of odor recognition performance in the behavioral assay. (A) The probability of PORs for locusts trained using hex as conditioning stimulus is shown as a bar plot. PORs to the trained odorant (shown in red) and to a diverse odor panel (nontrained odors shown in gray) are shown to allow comparison (*Methods*). A higher probability indicates that a larger proportion of trained locusts performed significant PORs during the testing phase when presented with that corresponding odorant (identified along x-axis). As can be seen, locusts had the highest POR probability to the trained odorant. Interestingly, locusts trained with hex also had significant PORs to 2-octanol (2-oct) and iaa. Other odorants, cit, bzald, apple (app), L-carvone (L-carv), neem, geraniol (ger), cyclohexanone (chex), and methyl salicylate (m-sal), did not evoke strong PORs in hex-trained locusts. These results were obtained by combining two datasets to yield $n = 47$ locusts for hex, $n = 27$ locusts tested with random presentations of hex, 2-oct, iaa, cit, bza, and app, and $n = 20$ locusts tested with random presentations of hex, L-carv, neem, ger, chex, and m-sal. (B) POR response probability for locusts trained and tested with the same CST (hex) but presented for different durations (0.75 s or 10 s; $n = 30$ locusts) or in a pulsatile fashion (2 s ON–2 s gap–2 s ON; $n = 27$ locusts). Reference Fig. 1B for representative POR traces. As can be noted, locusts have a high probability of response for all test pulses indicating robust recognition invariant of the stimulus duration or dynamics. (C) POR response probability following iaa introductions either solitary or atop a background odorant (bza). Note that iaa was presented with two different latencies (2 s and 4 s) following the onset of the sustained bza pulse. Representative POR traces for this stimulation protocol are shown in Fig. 1D. Iaa-trained locusts ($n = 27$ locusts) showed a very low probability of response to the untrained distractor odor (bza). All introductions of iaa, either solitary or atop bza background, evoked strong POR responses with a high probability of response across locusts. These results indicate that locusts could recognize a trained odorant in a background-invariant fashion. (D) POR probabilities during hex introductions following 10 different nonoverlapping distractor pulses are shown. Hex was used as the conditioning odorant. Representative PORs are shown in Fig. 1C. As mentioned earlier, $n = 27$ locusts were tested with 2-oct, cit, iaa, bza, and app as distractor stimuli, and a different set ($n = 20$ locusts) were tested neem, ger, L-carv, chex, and m-sal as distractors. Distractor odorants, the first pulse used in the sequence, are identified along the x-axis. These results indicate that hex-trained locusts have a high probability of response to the trained odorant irrespective of which distractor odorant was encountered prior to their onset (i.e., invariance with respect to stimulus history). (E) The POR probability for locusts trained with hex in a dry background (Left; $n = 21$ locusts) and for locusts trained with hex in a humid background (Right; $n = 20$ locusts) are shown. As can be seen, for both training paradigms, locusts have a high probability of response to hex under both dry and humid testing conditions. The results indicate that a trained stimulus could be recognized independent of changes in ambient humidity conditions. The difference in POR responses, though appear to be slightly stronger in training conditions, are not statistically significant ($P = 0.405$ for dry training, $P = 0.054$ for humid training; t test).

but using either dry air (0% RH) or humid air (100% RH) as the carrier stream. We found that again, most PNs in our dataset had responses that were variable in both dry and humid conditions (Fig. 4A and B). The overall distribution of response correlation between the first pulse and the subsequent encounters of the same odorant was low but comparable between dry and humid conditions (Fig. 4C).

Next, we sought to examine whether the ensemble neural responses across all PNs could reliably represent information about the identity of the target stimulus. To understand this, we visualized data after a principal component analysis dimensionality reduction. Briefly, spikes were binned in nonoverlapping 50-ms time bins, and the average spike counts (across 10 trials) of all 89 neurons became a high-dimensional snapshot of odor-evoked neural response in a particular time bin. The high-dimensional

vectors were projected onto the three eigenvectors of the data covariance matrix (corresponding to the largest eigenvalues) for the purpose of visualization. The three-dimensional representations of the ensemble PN activity were color coded to differentiate among the 11 target odor pulses (Fig. 4D, Left). As can be noted, the ensemble responses were variable and did not form a single well-defined cluster (i.e., not a unimodal distribution). More importantly, the PN response vectors in dry and humid conditions clearly formed distinct response clusters (Fig. 4D, Right), indicating that changes in humidity might also pose an additional challenge for achieving invariant odor recognition.

A Decoding Scheme for Robust Odor Recognition. Our results indicate that the behavioral recognition (i.e., the POR) remained invariant with respect to stimulus dynamics, history, the presence

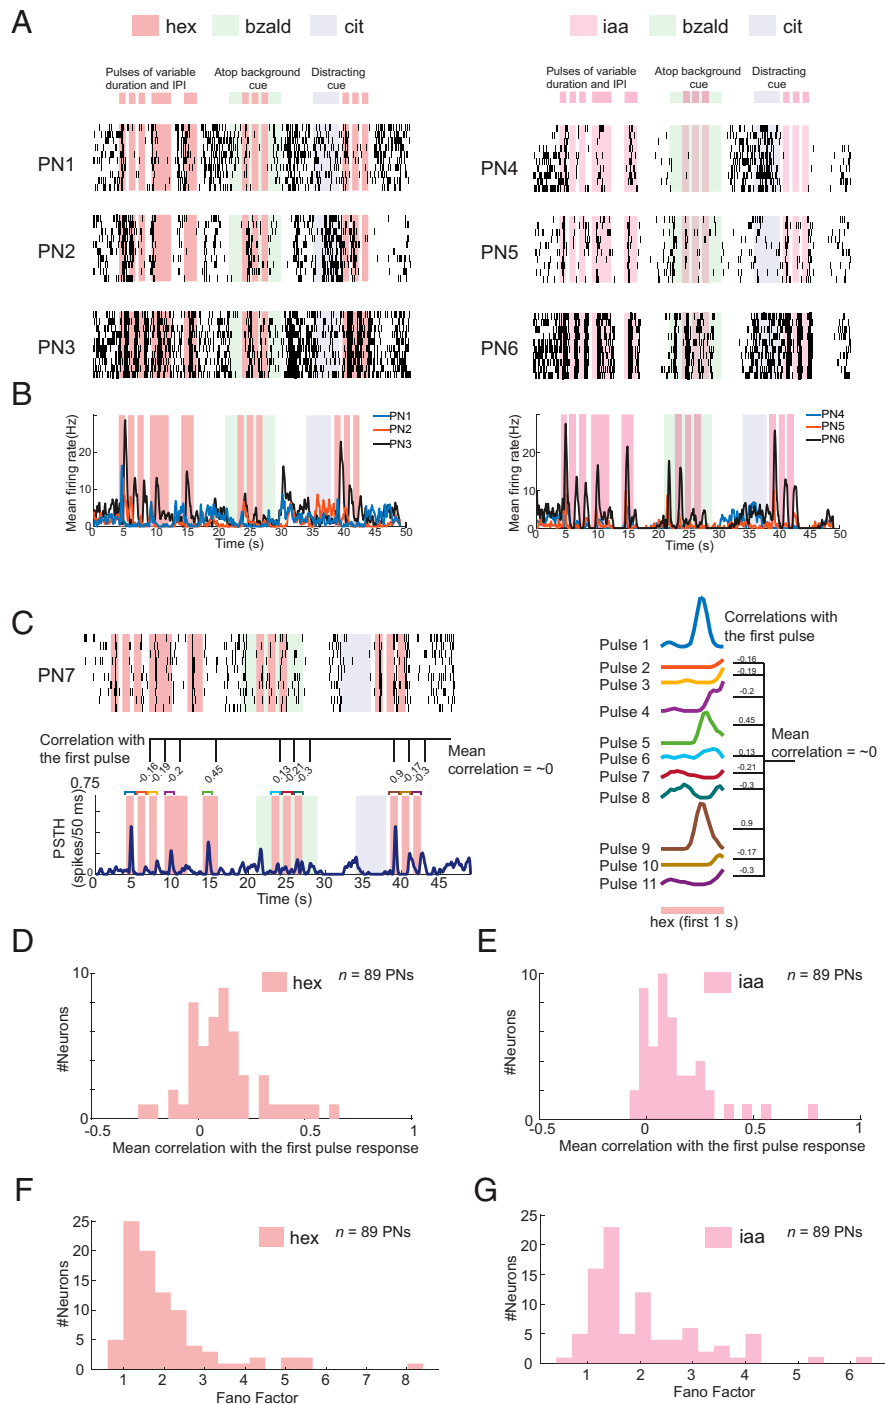


Fig. 3. Individual PN responses are highly variable. (A, Left) Raster plots showing PN responses during a pulsatile presentation of a target stimulus (hex) in back-to-back sequences of variable duration and inter-pulse intervals atop a background cue (bzald) and following a distracting stimulus (cit). Each black tick represents an action potential fired by the PN. PN responses are shown for 10 consecutive trials (10 rows). (Right) Similar plots as in the Left, but the target stimulus was iaa. Notice the responses evoked by the target odorant in these six PNs were highly variable. (B) Firing rates of the PNs (50-ms time bins; trial-averaged) shown in A are now plotted as a function of time. While both the PNs responded strongly to the first pulse of the target odorant, the response diminished during later encounters of the same stimuli. (C, Left) Raster plot and peristimulus time histogram (PSTH) of a PN are shown. Average spike rates across 10 trials are plotted as a function of time. Color line shown on the top of each hex presentation indicates the first 1-s response that was used to compute correlations. (Right) A schematic showing how the correlations were computed. Neural response to the first pulse was used as the response template to be pattern matched. Pairwise correlations with the first stimulus pulse were computed and averaged to obtain a mean correlation value for each PN. Note that a higher mean correlation value indicates a consistent response across all odor pulses. (D and E) Similarities between PN responses evoked during the first pulse of hex (target odorant) with all other encounters were computed. For this quantification, PN response was first binned into 50-ms time bins and averaged across 10 trials. The first 1-s response following onset of each target odorant pulse was used to compare response similarity between different target odor encounters (i.e., 20-dimensional response vectors). For each PN, the mean similarity across odor pulses was determined, and the response similarity across PNs were then plotted as a distribution. (E) Similar distribution as plotted in D but for iaa. (F and G) Distributions of Fano factor are shown across 89 PNs for hex and iaa. (F) Distributions are shown for hex pulses. Mean of the first 1-s neural response across 11 presentations of hex and 10 repeated trials was used to compute Fano factor (σ^2/μ) for each PN. (G) Similar plot as in F, but Fano factors were computed for iaa trials.

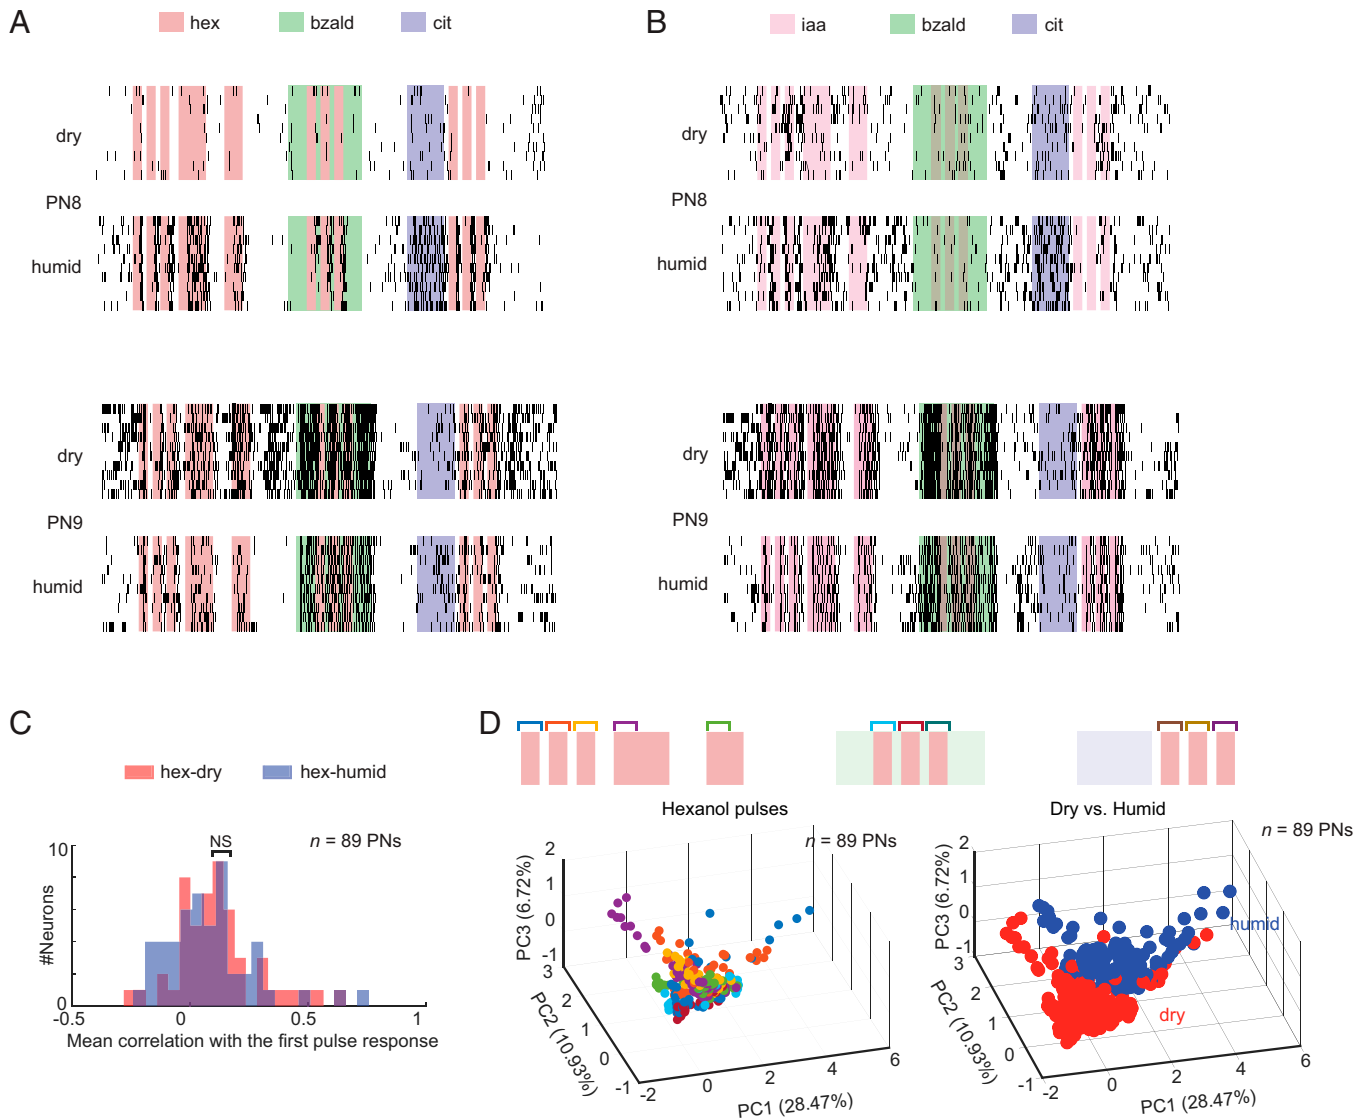


Fig. 4. Changing humidity conditions alter both individual- and ensemble-level PN activity. (A) Similar raster plot showing PN responses to the stimulation protocol used in Fig. 3. For each PN shown, the *Top* and *Bottom* plots reveal the spiking activity of the same PN between dry (carrier stream: 0% RH) and humid (carrier stream: 100% RH) conditions. Note that changes in humidity levels of the carrier stream resulted in increase or decrease in spiking activity in individual PNs. (B) Similar plots as in A but showing PN responses to a different target stimulus (iaa). Note that the same set of PNs are shown. (C) Similar plot as in Fig. 3D but comparing response similarity between PN responses ($n = 89$ PNs) observed in dry and humid conditions. “NS” indicates that the two distributions are not significantly different (two-sample Kolmogorov–Smirnov test; $P = 0.05$). (D, *Left*) Spiking activities of all 89 PNs during different pulses of the target odorant are shown after principal component analysis (PCA) dimensionality reduction. Spikes in individual PNs were binned in 50-ms nonoverlapping windows and treated as vector components. Each colored sphere represents an 89-dimensional PN spike count vector along the first three principal components. Different time segments are indicated by different color codes. The *Right* panel shows a similar plot but comparing odor-evoked responses in dry and humid conditions. Red and blue colored spheres are used to indicate the differences observed in the ensemble PN spiking activity in dry and humid conditions, respectively.


of competing cues, and changes in ambient conditions [Fig. 2; also refer to prior results on background (7) and history invariance (15)]. However, odor-evoked neural responses at single-cell and ensemble level vary with most of these perturbations (Figs. 3 and 4). Given this discrepancy between variability in neural encoding and robustness in behavioral output, we sought to determine whether a neural decoder could be designed for achieving robust odor recognition.

To investigate this issue, we again regarded the ensemble activity across the 89 PNs recorded in a 50-ms time bin as a high-dimensional neural activity (i.e., 89-dimensional firing-rate vector). To detect the presence of the target odorant and selectively recognize the identity of the target odorant, we trained a linear support vector machine classifier (SVM). Note that the


SVM classifier was trained using a separate set of training trials in which only a solitary 4-s pulse of the target odorant was presented. Two different classifiers, one for recognizing hex (hex-SVM; Fig. 5) and another for recognizing iaa (iaa-SVM; *SI Appendix, Fig. 3*) were trained. The probability that hex is present in any particular 50-ms time segment, as predicted using the hex-SVM, is plotted as a function of time (Fig. 5 A–C; similar plots for iaa-SVM shown in *SI Appendix, Fig. 3 A–C*). As can be noted, all introductions of the target odorant are detected and selectively recognized by both hex-SVM and iaa-SVM classifiers.

How did the linear classification approach manage to robustly recognize the target odorant from highly variable neural responses? To understand this, we examined what weights were assigned to different PNs in the dataset (Fig. 5D and *SI Appendix, Fig. 3D*).

A Support Vector Machine Classification results



B Recognition of target odor exposures in humid conditions



C Classifier trained on Hex and tested on IAA


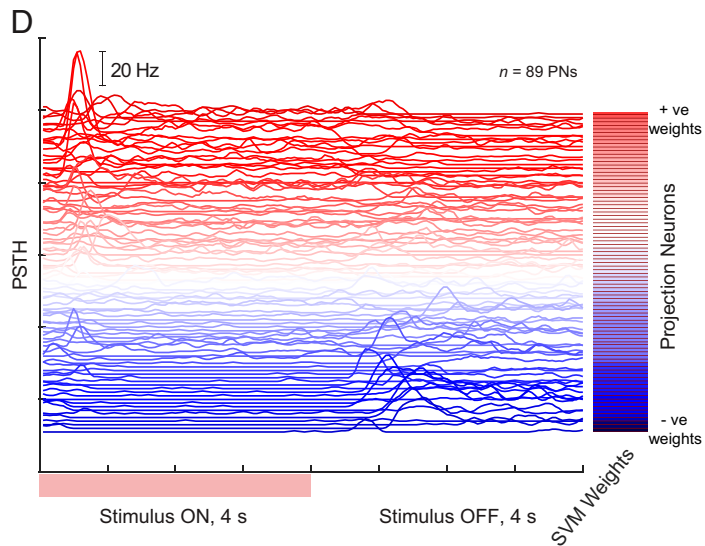



Fig. 5. Support Vector classification for robust hex recognition. (A) Classification probabilities for the target stimulus, hex, under different conditions using a SVM classifier are shown. The classifier was trained using 10 trials in which only a solitary pulse of hex was presented. For each time bin, the probability was obtained by averaging classification results across 10 trials. Time segments when a stimulus was presented and the identity of the stimulus that was presented are indicated using colored bars. The dotted line indicates the classification probability of 0.5 (an arbitrary threshold) that was used to detect the presence of hex. The segments when the trial-averaged classification probability exceeded this threshold are indicated using an asterisk. (B) Performance of the hex classifier but tested during those trials in humid condition. Note that humidity did not perturb target order classification. (C) Performance of the hex classifier but tested during those trials when iaa was used as the target odorant. Note that none of the iaa presentations were classified as the target stimulus being present. (D) The PSTHs of each PN during 4-s solitary hex presentation and 4-s after stimulus termination are shown. Each trace shows PSTH of one PN. The PSTHs are sorted and color coded based on the hex-SVM classifier weights assigned to each PN ($n = 89$ PNs). The neurons that were assigned the most-positive weights are at the top, and the most negatively weighted neurons are near the bottom. SVM classifier weights that were assigned to each PN are also schematically shown as a heatmap to the right. (E) The fraction of ON and OFF neurons that received '+ve' or '-ve' weights are shown as a bar plot for comparison.

We sorted the PNs based on the weights they were assigned by the SVM classifier and plotted their average spike counts to the solitary target odor pulse (i.e., training data that was used). Note that PNs with strong ON responses are at the top and PNs with stronger OFF responses are at the bottom. The weights assigned to each PN by the hex-SVM classifier are shown next to the firing-rate plots to allow comparison. Our results indicate that ON-responsive PNs received mostly positive weights and the OFF-responsive PNs were assigned negative weights (Fig. 5 D and E). A similar weighting scheme was also used by the iaa-SVM classifier (SI Appendix, Fig. 3 D and E). This result suggests that although individual and population neuron responses vary, both detection and recognition of the odor identity can still be achieved using a linear coding scheme, and the weight assignment to individual neurons had a simple structure (i.e., positive weights to ON neurons and negative weights to OFF neurons).

Discrete Classifier with Ternary Weights Allows Robust Recognition.

Finally, we wondered how stable this classification approach was. The variations that we examined here are but only a few of the many intrinsic or extrinsic perturbations feasible. Therefore, we particularly wondered how important the analog weights assigned to each neuron were. The PNs that had the strongest ON and OFF responses to solitary pulses of the target odorant might not necessarily have similarly intense responses during other encounters of the same stimulus. However, the set of PNs that get excited or inhibited can be expected to be maintained across different encounters of the same odorant. Therefore, we wondered whether the classification approach could be simplified by replacing the analog weights assigned to each neuron with a simpler scheme.

The first simplification that we examined was converting the analog SVM weights into just three values: $\{-1, 0, +1\}$. This discretization was done by comparing the analog SVM weights with two arbitrarily chosen thresholds. We found that this simpler approach was still able to allow robust detection as well as discrimination of both the target odorants (Fig. 6A and SI Appendix, Fig. 4A).

Since the thresholding of analog SVM weights might not provide an optimal approach to derive a ternary classifier, we sought to directly determine the optimal ternary weights (Methods). We found that the ternary weights learned using this constrained approach also provided robust recognition of the target odorants (Fig. 6B and SI Appendix, Fig. 4B). Further simplification of this approach by just using Boolean weights, $\{1 \text{ or } 0\}$, allowed detection of the target odorant, but discrimination between the two target odorants was compromised (Fig. 6C and SI Appendix, Fig. 4C). Furthermore, as a control, random assignment of ternary weights to individual PNs failed to provide decent recognition performance (SI Appendix, Fig. 5). A comparison of the weights assigned by the different classifiers are shown in Fig. 7. These results indicate that the ternary classifier provides a sort of lower bound on the recognition performance that could be achieved. In other words, it strikes a fine balance between decoding scheme complexity and recognition performance.

In sum, our results indicate that a very simple classification scheme using just ternary weights could provide robust odor recognition.

Discussion

We examined how invariant recognition of odorants can be achieved in a relatively simple locust olfactory system. Our results indicate that individual PN responses can vary with one or several of the perturbations that we studied, including stimulus dynamics, repetition, stimulus history, presence of background odorants, and changes in humidity conditions. Nevertheless, a simple linear classification scheme was sufficient to extract the relevant information. The classifier essentially boiled down to adding the

contribution of PNs that were strongly activated when the odorant was presented (ON neurons) and subtracting the contribution of PNs that were activated after the termination of the odorant (OFF neurons). Notably, such a classifier was not only able to robustly detect all introductions of a target stimulus (i.e., solve the detection problem) but also provided sufficient discrimination between the two target odorants (i.e., solve the recognition problem as well).

We found that not all neurons were perturbed, and only a small subset (four of 89 PNs) of them responded reliably to all introductions of both the target odorants. While these neurons allowed robust detection of the target odorants, they were not specific and responded to both the target odorants examined in this study. Furthermore, various response features such as spike counts, response latency, etc. varied across different encounters of the same stimulus, thereby making discrimination between target odorants based on just these reliable PNs not feasible (SI Appendix, Fig. 2). Therefore, an approach based on a single or on a small subset of neurons encoding for a stimulus under all conditions cannot be expected to be fault tolerant.

Prior publications (7, 11, 14–16) have also found individual neurons to be unreliable but found that robustness emerged at an ensemble level. However, our prior results indicated that odorants delivered atop different background cues generated ensemble responses that only partially overlap across conditions (7). When additional perturbations were introduced, our results indicate that the ensemble response features also tend to vary unreliably. In other words, even at the ensemble level, there was not a single feature that could remain consistent when the odor-evoked responses were minimally perturbed. How then could sensory invariance be achieved?

Given the complexity of individual PN response changes, we did not expect that a linear classifier could provide robust recognition. Nevertheless, both a linear SVM and logistic regression classifier were able to decode target odor identity independent of the perturbations made. Surprisingly, simplification of this approach by constraining weights to assume ternary values $\{-1 \text{ or } 0 \text{ or } 1\}$ also provided robust recognition. The goal for constraining the decoding scheme in this fashion was twofold: interpretability of the approach taken and determining the simplest possible approach (i.e., a sort of a lower bound in recognition performance). Our results indicate that not only did such a scheme exist, but it exploited a simple stratagem. We found that successful decoding schemes (discrete SVM or a logistic regression approach) assigned “+1” weights to ON neurons, “0” weight to nonresponders, and “-1” weight to OFF neurons (i.e., an “ON minus OFF” classifier). Notably, robust recognition was achieved using this simple approach in both dry (Fig. 6) and humid conditions (SI Appendix, Fig. 6A and B).

If ensemble response features varied across conditions, how did this “ON minus OFF” classification approach achieve invariance? It is worth breaking this classification scheme into its two components: the ON component and the OFF component. Assigning “+1” to the most-strongly responding ON neurons and setting a recognition threshold that is less than this sum allows the classification scheme to be flexible. Interpreted differently, this indicates that an odor can be recognized as long as a subset of strongly responding ON neurons are activated so that their sum reaches the threshold. The composition of this subset can change across conditions, thereby allowing this approach to be more flexible.

What then is the contribution of the OFF component of the classifier? In an earlier study (8), we found that OFF responses were better at predicting when the behavioral response to a conditioned odorant terminated. In this study, we found that the OFF component increased separability between activation patterns of different odorants. This effect was particularly noticeable when the ternary weights were further simplified to a Boolean classifier with binary weights. While the Boolean classifier allowed

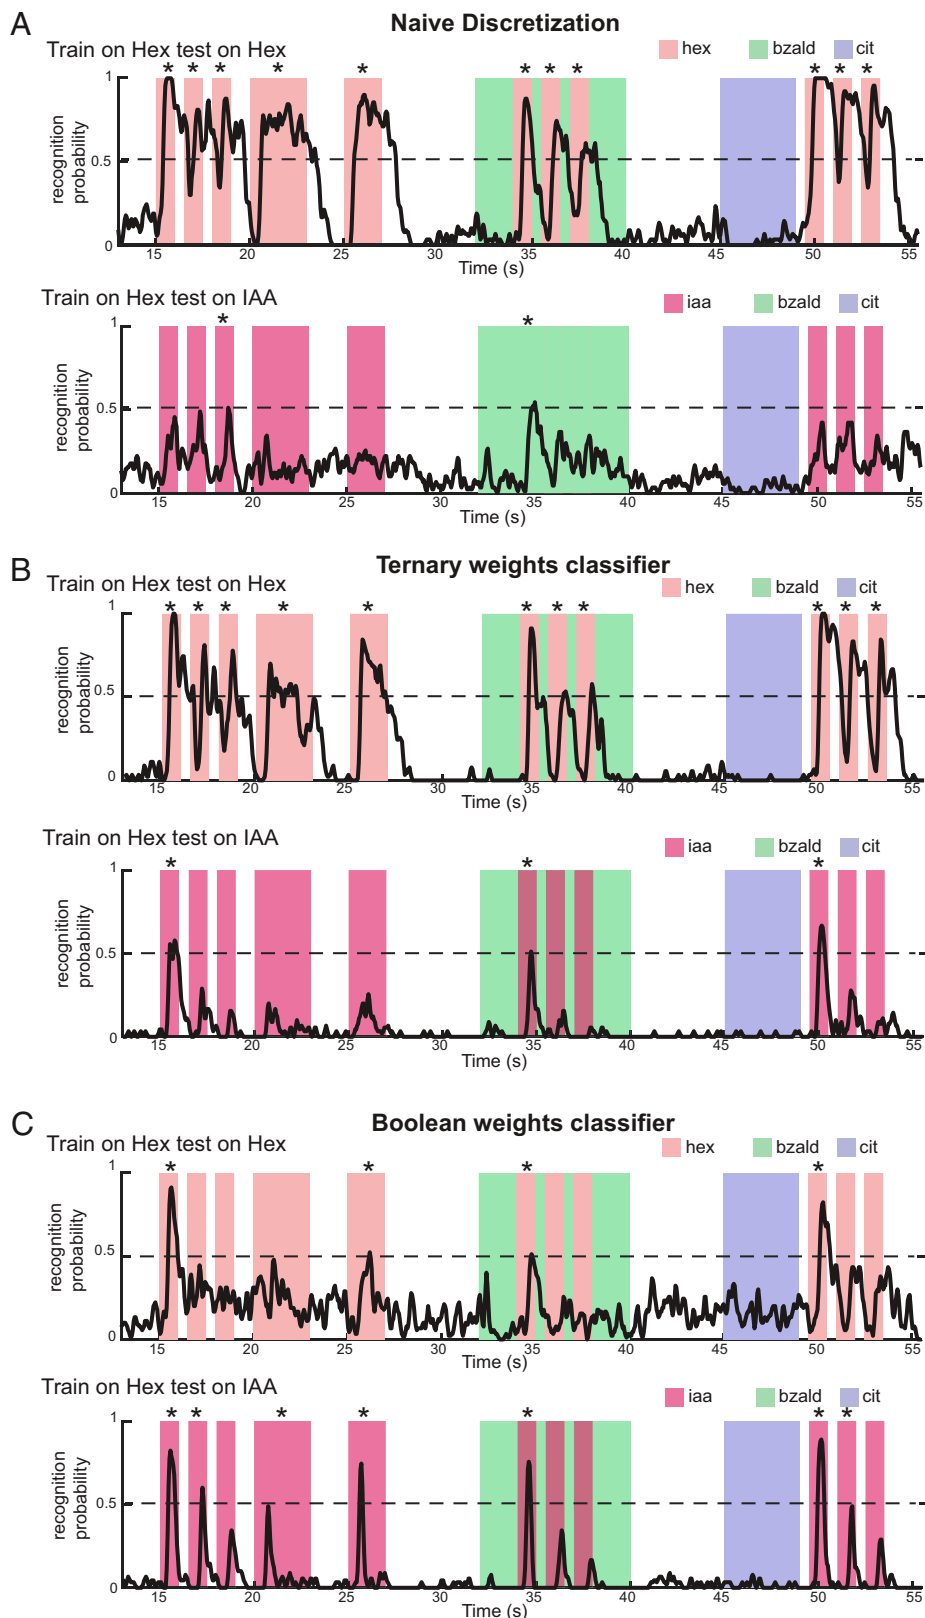


Fig. 6. Performance comparison between discrete SVM versus ternary versus Boolean classifiers. (A) Classification probabilities for an SVM classifier that used discrete set of weights, $\{-1, 0, 1\}$, are shown. The ternary weights were obtained from SVM classifier by discretizing SVM weights using two preset thresholds. Other conventions are same as Fig. 5. (Bottom) Performance of the Hex-SVM classifier with discrete weights on recognizing iaa introductions. (B) Similar plots as earlier but showing performance of a Bayesian logistic regression approach in which each neuron was directly assigned a ternary weight (“-1” or “0” or “1”; *Methods*). In total, Ten trials of solitary exposures to hex were used to train the classifier. (C) Similar plots as earlier but showing performance of a Bayesian logistic regression approach in which each neuron was directly assigned a Boolean weight {“1” or “0”}. Ten trials of solitary exposures to hex were used to train the classifier.

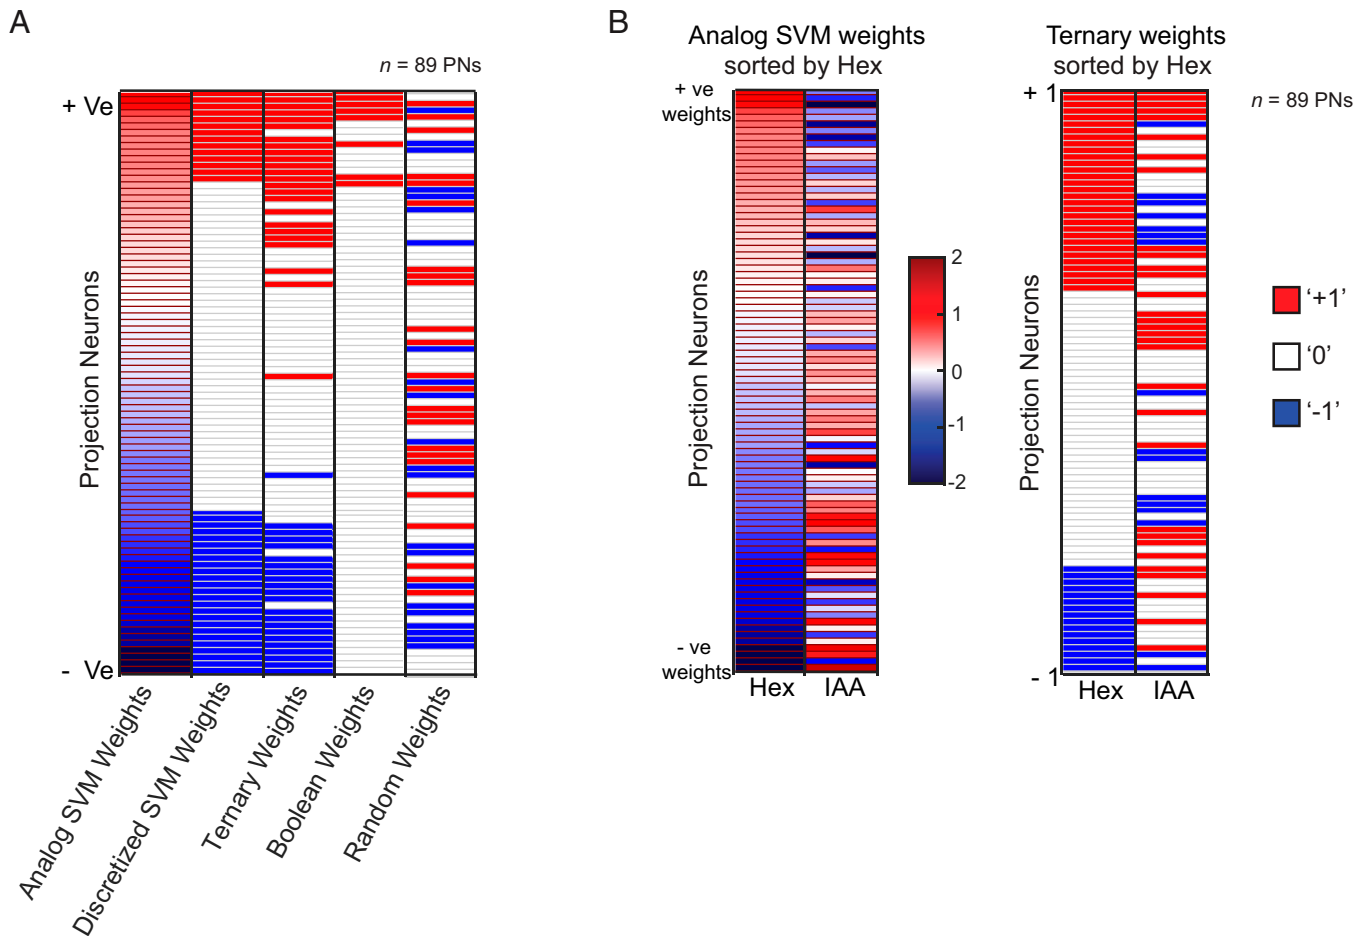


Fig. 7. Comparison of weights assigned to PNs by different classification schemes. (A) The neurons ($n = 89$ PNs) are sorted such that the most-positive weights are at the top and the most negatively weighted neurons are near the bottom. The first column shows analog SVM weights learned. The ternary weights after discretization that were assigned to each PN are shown in the second column: red indicates a weight of “+1,” white indicates a weight of “0,” and blue indicates a weight of “-1.” The third column shows the ternary weights learned using Bayesian framework as a heatmap. The fourth column shows the Boolean weights learned, and the last column shows random weights generated. All columns use the same index as column one. (B) The weight vector learned by the analog SVM trained with hex exposures (i.e., hex-SVM) and its ternarized version are shown along with analog and ternary weights learned by the SVM classifier trained with iaa exposures (i.e., iaa-SVM). Note that each weight vector has the same number of components as the number of neurons in the dataset. The weights are sorted based on the values assigned by the analog Hex-SVM classifier.

detection of the target odor pulses, there was a significant increase in the false-positive rates. Therefore, we conclude that assigning a negative weight to the OFF neurons enhanced discrimination between odorants and thereby reduced false positives.

There were other considerations that went behind considering this simpler classification scheme. As we have noted, the response evoked by an odorant in each neuron (i.e., changes in firing rates) can vary with most of the perturbations that we introduced. In addition, repeated presentation of a stimulus will also lead to adaptation that can further attenuate the neural responses. Therefore, weighing individual neurons based on their response to solitary exposures of the target odorant might make the classification scheme unstable when recognition under other conditions/perturbations are required. Could the combination of neurons activated alone be a more-robust indicator than the firing-rate distribution across the same set of neurons (i.e., vector direction being more important than the length of the vector in any particular direction)? If this is the case, a ternary weight vector should allow robust decoding of odor identity. Our results confirm this expectation. Furthermore, it can be shown that a ternary version of a high-dimensional weight vector in a classifier is highly aligned with the analog version of the same weight vector

(i.e., angular distance is low compared to pairs of random vectors; *SI Appendix, Fig. 7*).

Earlier studies have argued that the antennal lobe neural network can be viewed as a nonlinear dynamical system (6, 17, 18). Under this perspective, both the initial conditions and odor-evoked response dynamics become important for recognizing the identity of the encountered stimulus. Our data indicate that both the odor-evoked responses and the spontaneous activity can vary across conditions. Yet, at direct odds with our neural data, we find that the behavioral recognition is robust even during these drastic changes. No detectable differences in response latency, intensity, or duration were found. Hence, our results indicate that the rules for translating neural responses and their dynamics to generate appropriate behavioral output needs further investigation.

The behavioral data also indicates that locusts trained with one odorant as CST generalized their responses to a few other odorants (Fig. 2A). Both analog SVM and the ternary classification schemes were able to generate prediction results from neural data that matched with the observed trends in behavioral data (*SI Appendix, Fig. 8*). These results further support the proposed scheme for translating the variable neural responses to robust behavioral outcomes.

Finally, we wondered whether the internal state of an organism such as its hunger level and whether it was trained to associate a particular target odorant with a reward would alter the stability of neural responses. To examine this issue, we compared the neural responses in locusts with different hunger levels (unstarved versus starved) and/or training state (untrained versus trained). Our results indicate that irrespective of the internal state, the neural responses to target odorants were highly variable (supra-Poisson Fano factors), and the odor-evoked responses were just as inconsistent across different encounters of the same target odorant (*SI Appendix, Figs. 9–11*). While hunger level and odor-reward pairing have been suggested to alter certain odor-evoked response features in this neural circuit (19–21), our data indicate that such state-dependent changes may still not compensate for variations induced by the battery of extrinsic perturbations such as those explored in this study. Therefore, we conclude that the decoding scheme proposed by our results would still be relevant and necessary for robust odor recognition.

Methods

For comprehensive description of 1) odor stimulation, 2) electrophysiology, 3) spike sorting, 4) behavior experiments, 5) principal component analysis, 6) neural response characterization, 7) linear SVM classification, and 8) binary/ternary classification neural networks, refer to *SI Appendix*.

Data Availability. All data presented in this paper are publicly available in Figshare (DOI: [10.6084/m9.figshare.17162570](https://doi.org/10.6084/m9.figshare.17162570)) (22). Previously published data were used for this work (some behavioral data used in Figs. 1 and 2 and *SI Appendix, Fig. 8* were reanalyzed from our earlier publications; this is clearly described in the main text and *Methods*, and appropriate citations are included in the manuscript).

ACKNOWLEDGMENTS. We thank members of the B.R. laboratory (Washington University in St. Louis) for feedback on the manuscript. We thank Pearl Olsen for insect care. This research was supported by NSF CAREER Award (No. 1453022), NSF Collaborative Research in Computational Neuroscience Grant (No. 1724218), NSF Biology Integration Institute Grant (No. 2021795), and Office of Naval Research Grant (No. N000141912049) to B.R.

1. M. Stopfer, S. Bhagavan, B. H. Smith, G. Laurent, Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* **390**, 70–74 (1997).
2. I. Ito, R. C. Y. Ong, B. Raman, M. Stopfer, Sparse odor representation and olfactory learning. *Nat. Neurosci.* **11**, 1177–1184 (2008).
3. S. A. Kreher, D. Mathew, J. Kim, J. R. Carlson, Translation of sensory input into behavioral output via an olfactory system. *Neuron* **59**, 110–124 (2008).
4. R. Fdez Galán, S. Sachse, C. G. Galizia, A. V. M. Herz, Odor-driven attractor dynamics in the antennal lobe allow for simple and rapid olfactory pattern classification. *Neural Comput.* **16**, 999–1012 (2004).
5. M. Stopfer, V. Jayaraman, G. Laurent, Intensity versus identity coding in an olfactory system. *Neuron* **39**, 991–1004 (2003).
6. O. Mazor, G. Laurent, Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* **48**, 661–673 (2005).
7. D. Saha *et al.*, A spatiotemporal coding mechanism for background-invariant odor recognition. *Nat. Neurosci.* **16**, 1830–1839 (2013).
8. D. Saha *et al.*, Engaging and disengaging recurrent inhibition coincides with sensing and unsensing of a sensory stimulus. *Nat. Commun.* **8**, 15413 (2017).
9. D. Saha *et al.*, Behavioural correlates of combinatorial versus temporal features of odour codes. *Nat. Commun.* **6**, 6953 (2015).
10. J. A. Riffell, H. Lei, J. G. Hildebrand, Neural correlates of behavior in the moth *Manduca sexta* in response to complex odors. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 19219–19226 (2009).
11. S. L. Brown, J. Joseph, M. Stopfer, Encoding a temporally structured stimulus with a temporally structured neural representation. *Nat. Neurosci.* **8**, 1568–1576 (2005).
12. N. J. Vickers, T. A. Christensen, T. C. Baker, J. G. Hildebrand, Odour-plume dynamics influence the brain's olfactory code. *Nature* **410**, 466–470 (2001).
13. M. N. Geffen, B. M. Broome, G. Laurent, M. Meister, Neural encoding of rapidly fluctuating odors. *Neuron* **61**, 570–586 (2009).
14. Z. N. Aldworth, M. A. Stopfer, Trade-off between information format and capacity in the olfactory system. *J. Neurosci.* **35**, 1521–1529 (2015).
15. S. Nizampatnam, D. Saha, R. Chandak, B. Raman, Dynamic contrast enhancement and flexible odor codes. *Nat. Commun.* **9**, 3062 (2018).
16. B. M. Broome, V. Jayaraman, G. Laurent, Encoding and decoding of overlapping odor sequences. *Neuron* **51**, 467–482 (2006).
17. M. Rabinovich *et al.*, Dynamical encoding by networks of competing neuron groups: Winnerless competition. *Phys. Rev. Lett.* **87**, 068102 (2001).
18. G. Laurent *et al.*, Odor encoding as an active, dynamical process: Experiments, computation, and theory. *Annu. Rev. Neurosci.* **24**, 263–297 (2001).
19. J. Beshel, Y. Zhong, Graded encoding of food odor value in the *Drosophila* brain. *J. Neurosci.* **33**, 15693–15704 (2013).
20. C. Martelli *et al.*, SIFamide translates hunger signals into appetitive and feeding behavior in *Drosophila*. *Cell Rep.* **20**, 464–478 (2017).
21. T. Faber, J. Joerges, R. Menzel, Associative learning modifies neural representations of odors in the insect brain. *Nat. Neurosci.* **2**, 74–78 (1999).
22. S. Nizampatnam, L. Zhang, R. Chandak, J. Li, B. Raman, Invariant odor recognition with ON-OFF ensembles: Data and Matlab scripts. Figshare. https://figshare.com/articles/dataset/Data_and_Matlab_Scripts/17162570. Deposited 12 December 2021.