

Neural manifolds for odor-driven innate and acquired appetitive preferences

Rishabh Chandak and Barani Raman[#]

Department of Biomedical Engineering, Washington University in St. Louis

[#]Corresponding author. Tel: 1-314-935-8538 Fax- 1-314-935-7448 Email: barani@wustl.edu

Abstract

Sensory stimuli evoke spiking neural responses that innately or after learning drive suitable behavioral outputs. How are these spiking activities intrinsically patterned to encode for innate preferences, and could the neural response organization impose constraints on learning? We examined this issue in the locust olfactory system. Using a diverse odor panel, we found that ensemble activities both during ('ON response') and after stimulus presentations ('OFF response') could be linearly mapped onto overall appetitive preference indices. Although diverse, ON and OFF response patterns generated by innately appetitive odorants were still limited to a low-dimensional subspace (a 'neural manifold'). Similarly, innately non-appetitive odorants evoked responses that were separable yet confined to another neural manifold. Notably, only odorants that evoked neural response excursions in the appetitive manifold were conducive for learning. In sum, these results provide insights on how encoding for innate preferences can also set limits on associative learning.

Introduction

In many organisms, the olfactory system serves as the primary sensory modality that guides a plethora of behaviors such as foraging for food, finding mates, and evading predators. The genetic makeup of these organisms determines the innate preference, or valence, associated with different olfactory stimuli¹⁻⁶. Consequently, neural responses evoked by these stimuli have to be patterned to drive motor neurons to perform appropriate behaviors (i.e., move towards or away) that are key for survival. Given the importance of rapid and robust decision-making⁷⁻¹⁰, we wondered how information regarding the valence of a chemical cue is encoded in the olfactory system. Particularly, we examined whether and how neural responses are spatiotemporally structured to represent odor valence in the early locust olfactory system.

In insects, odor stimuli are detected by olfactory sensory neurons in the antenna that transduce chemical cues to electrical signals and relay them to the antennal lobe. A network of cholinergic projection neurons (PNs, excitatory) and GABAergic local neurons (inhibitory) in the antennal lobe fire in unique spatiotemporal combinations to encode for stimulus identity¹¹⁻¹⁷. The PN responses are patterned over space and time to encode for different odorants near the insects and relay this information to higher centers responsible for learning, memory, and overall behavioral preferences¹⁸⁻²⁰. The odor-evoked PN response patterns are elaborate and continue well after the stimulus is terminated. Since the behavioral responses initiated by an odorant are often rapid^{8,9,21}, the relevance of neural activity that occurs well after the stimulus onset remains to be understood. However, what has been reported is that the behavioral responses elicited by an odorant last the duration of the stimulus exposure⁵⁻⁷.

Recent studies have shown that the temporal patterns of neural responses change most dramatically after the stimulus is terminated. The set of PNs activated during the stimulus presence (i.e., the ON responders) and those that get activated after stimulus termination (i.e., the OFF responders) have minimal overlap^{23,24}. Intriguingly, these OFF responses also tend to be odor-specific and appear to contain as much information as the ON responses elicited by the stimuli. Whether the combination of neurons and the temporal order in which they are activated could be used to control and shape the behavioral response

dynamics remains to be investigated. In addition to encoding for innate odor preferences, the neural responses evoked by an odorant must also underlie how and when it is associated with other stimuli through learning^{25,26}. The importance of timing between a stimulus and reward, and how it controls learning and the rate of learning is also well documented^{12,18,27–31}. Given that most odorants elicit spatiotemporally varying activity, and the relative timing of the reinforcing stimuli can be controlled, can any response segment be reinforced with a reward? In this study, we examined the spatiotemporal coding logic that constrains neural representation for odorants and how the organizational logic of odor-evoked responses also impacts behavioral preferences and sensory memory.

Results

Innate appetitive preferences of locusts to an odor panel

We began by assaying the innate appetitive preferences of starved locusts to a large, diverse panel of odorants (1% v/v unless stated otherwise). Each odor in the panel was presented to every locust once using a pseudorandomized order. The palp-opening responses (POR) evoked by all odorants in the panel were recorded (**Fig. 1a, b**). We used a binary metric to quantify whether each locust responded to an odor by opening its palps (a score of 1 to indicate a palp-opening response (white-colored boxes), and a score of 0 to indicate no response (gray-colored boxes)). For visualization, the odors were sorted based on the number of PORs they elicited across locusts.

We converted these results to a preference index for each odor (see **Methods**). As can be seen from **Fig. 1c**, we obtained a broad range of preferences for the odor panel. Hexanol (at 10% v/v; leftmost odorant; x-axis), a green-leaf volatile, had the highest preference, whereas linalool (rightmost odorant; x-axis), a pesticide, had the lowest preference. We categorized odorants as being appetitive, neutral, and unappetitive (one-sided binomial test comparison; neutral and unappetitive odors are jointly referred to as ‘non-appetitive’). Prior studies have found that preferences for certain odorants can vary between males and females of the same species^{1,32,33}. To examine this possibility, we also compared behavioral responses between male and female locusts ($n=13$ for each gender, **Supplementary Fig. 1a**). While appetitive preferences for certain odorants did vary between males and females in our dataset (e.g., hexanal and garlic), these differences were not significant (t-test, $p>0.1$ for all odors).

Is there a simple stimulus feature that could account for these diverse appetitive preferences? Since the odorants were diluted to the same concentration (1% v/v) and delivered identically (except hexanol which was alone delivered at 10%, 1%, and 0.1% dilutions), the vapor pressure of the chemicals directly determined how much of each stimulus was delivered. We wondered then if locusts were simply behaving more frequently for more volatile odors (higher vapor pressure). However, as can be seen in **Fig. 1d** (and **Supplementary Fig. 1b**), a regression between the vapor pressure of the stimuli against the behavioral responses poorly explained the observed POR trend.

Another potential confound that could impact the observed trends could arise from fatigue/loss of motivation that could potentially diminish the locust PORs in the later trials of the experiment. To eliminate this possibility, we plotted the observed number of PORs as a function of the trial number (**Fig. 1e**). As can be noted, our results indicate that locust performance remains robust and even slightly increased as the experiment progressed ($R^2 = 0.23$; **Supplementary Fig. 1c**). Finally, we performed Monte Carlo simulations to verify that population-level responses were not biased by a handful of individuals. Our results confirmed that this is indeed the case and the results converged when any random subset of eighteen or more locusts was used to calculate behavioral preference indices for different odorants (**Fig. 1f**). These results, combined with the pseudorandom presentation of odorants, indicate that

the behavioral preferences obtained are a strong indicator of the innate appetitive preference of the locusts, and the sample size used was sufficient to get a stable readout.

Individual projection neuron responses to appetitive and non-appetitive odors

Next, we sought to understand the neural basis of this behavioral readout. To examine this, we recorded odor-evoked responses from projection neurons (PNs) in the locust antennal lobe (**Fig. 2a**). We stimulated the antenna with the same odor panel used in the behavioral experiments. The stimulus dynamics of each odorant were quantified using a photo-ionization detector (PID) and the mean voltage responses for all odors are shown in **Fig. 2b** (left panel; see **Methods**). The right panel shows the peak PID response for each odorant arranged in order of innate appetitive preferences (cues that evoked the highest behavioral responses are on the left and lowest on the right).

We presented each odorant for ten repetitions in a pseudorandomized order. A total of 89 PNs (~10% of the total number of PNs in a single antennal lobe) were recorded using this approach and used for all subsequent analyses. Consistent with prior data, we found that odor-evoked responses had two prominent epochs: an ON response that occurred during the 4 s when the stimulus was presented, and an OFF response that occurred during a 4 s window immediately following stimulus termination. We found a PN that had an ON response for most of the odorants (**Fig. 2c**, PN A), whereas many PNs responded to a subset of odorants either with an ON response or an OFF response. A small fraction of neurons were OFF-responders to a few appetitive odors but switched to ON-responses for some of the non-appetitive odorants (**Fig. 2c**; PN B; 8/89 PNs with similar tuning). Complementing these responses, we also found a small fraction of PNs that was ON-responsive to all five appetitive odorants but was OFF responsive to one or more non-appetitive odorants (**Fig. 2c**, PN C; 11/89 PNs with similar tuning). On average, odorants with higher valence elicited stronger ON and OFF responses across more PNs than those with lower valence, while inhibition increased as the odorants became less appetitive (**Fig. 2d**; see **Methods**).

We computed the correlation between the individual PN responses to different odorants with the overall behavioral preferences to the same panel (**Fig. 2e**). Notably, we found a small subset of neurons that had either a strong positive or negative correlation with the POR responses observed. Furthermore, our results indicate that such correlations could be found when either the ON or OFF responses were used. Although, it would be worth noting that different subsets of PNs had a high correlation with appetitive preference during the ON and the OFF periods.

How selective are individual PN responses? To answer this, we computed a tuning curve for each PN during both the odor ON and OFF periods (**Fig. 2f**). We found that most PNs responded to at least two odorants or more during the ON period (84/89 PNs) and a small fraction of neurons (11/89 PNs) responded to ten or more odorants (**Fig. 2f**, bar plots along the y-axis). The odor-evoked responses were more selective during the OFF period, with 70/89 PNs responding to two or more odors and only three PNs responding to more than ten odorants. In sum, these results indicate that individual PNs responded to the odor panel with great diversity.

Ensemble projection neuron responses to appetitive and non-appetitive odors

Next, we examined how the odor-evoked responses vary at an ensemble level. To visualize the ensemble neural responses and how they change as a function of time, we used a linear dimensionality reduction technique (Principal Component Analysis, PCA; see **Methods**). PCA neural response trajectories for the ON period are shown for all odorants (**Fig. 3a**). Consistent with prior findings^{15,24,34,35}, our data also reveal that each odorant produced a distinct looped response trajectory. Interestingly, we observed that neural response trajectories evoked by odorants that were labeled as innately appetitive in the behavioral assay

evolved in a similar direction (blue trajectories). This indicates that the combination of PNs excited by these odors had overlap and hence the PN ensemble vectors were near one another in the state space. Similarly, the trajectories for odors labeled as unappetitive also evolved in a similar direction (red trajectories) and occupied a different region of the state space. Note that the sets of red and blue trajectories did not overlap, indicating that odors within different groups (appetitive and unappetitive) were being encoded by distinct subsets of PNs.

We confirmed these dimensionality reduction results with a high-dimensional clustering analysis (**Fig. 3b**). We found that the spiking profiles for odors that belonged to the same group (appetitive or unappetitive) were similar, and hence clustered within the same branch when visualized using a dendrogram. These results support our interpretation that unique subsets of PNs in the antennal lobe are activated in a manner that is representative of the innate appetitiveness of the stimulus.

Predicting behavioral preferences from odor-evoked neural responses

How well do the neural responses map onto the behavioral preferences for different odorants? To examine this, we used linear regression to predict the probability of generating a POR given the ensemble PN activity elicited by that odorant. (**Fig. 4a**). Note that for these predictions, we used the normalized behavioral responses for each odor (see **Methods**), which could also be interpreted as the probability of a palp-opening response to a given odorant (across locusts). The regression weights were trained using all but one odorant and used to predict the probability of POR for the left-out odorant (i.e., a leave-one-odorant-out-cross-validation approach; 22 different linear regression models were used). We found that this simple approach yielded robust predictions for all odorants (**Fig. 4b, c**).

Note that we made predictions using the mean ensemble PN activity during 4 s of odor exposure (i.e., an ‘ON-regressor’), and using 4 s of odor-evoked activity after the termination of the odorant (i.e., an ‘OFF-regressor’). Both the regressors performed relatively well with the ON-regressor performance being better than the OFF-regressor. Further, the performance of the linear regression approach with shuffled prediction probabilities for different odorants (i.e., ‘shuffled control’ for both ON and OFF cases) predicted values around the mean POR probability for all odorants (**Fig. 4b, c**; mean = ~0.4), and was significantly inferior compared to the ON- and OFF- regression approaches. The poor performance of the shuffled control approach compared to the ON- and OFF- regressors suggests that the spiking activity across PNs is indeed organized to enable mapping between neural and behavioral responses spaces.

How consistent were the different regression models? Our results indicate that the weights assigned to each PN remained stable irrespective of the odor that was left out to train the regression model (**Fig. 4d**). This consistency of the assigned weights across regressors indicates that no particular odorant disproportionately influenced the regression model used to transform neural responses into POR probabilities. Additionally, Monte Carlo simulations (see **Methods**) revealed that both the ON- and OFF-regressors’ performance improved as the number of PNs used in the analyses was increased (**Supplementary Fig. 2**).

We wondered whether the same set of PNs contributed during both ON and OFF periods to predict the preference index for different odorants. To understand this, we calculated the correlation coefficient between the weights assigned by both these regression approaches (**Fig. 4e**). Our results indicate that there was only a weak correlation between weights assigned by the ON- and OFF-regressors. These results indicate that information regarding the overall appetitive preference is distributed across different sets of PNs during the ON vs OFF epochs. In sum, we conclude that the ensemble neural responses during odor presentations and after their termination are odor-specific, and contain information about the overall innate behavioral response generated by that odorant.

Innate versus acquired appetitive preferences for odorants

Next, we wondered if innate appetitive preferences to odorants and the neural responses they evoke can inform regarding other behavioral dimensions such as learning and memory. To examine this, we used an appetitive-conditioning assay (**Fig. 5a**). Locusts were starved for 24 hours and pre-screened for innate responses to the odorants used in the assay. Only those that did not have innate responses were used for the appetitive-conditioning experiments (see **Methods**).

We trained locusts with four chemically and behaviorally diverse odorants as conditioned stimuli in an ‘ON-training paradigm’ (**Fig. 5a, Supplementary Fig. 3**). Following training, we examined the ability of the trained locusts to respond to the conditioned stimulus in an unrewarded test phase. Opening of maxillary palps (palp-opening response) was regarded as a read-out of successful stimulus recognition. We found that locusts trained with hexanol or isoamyl acetate as conditioned stimulus robustly responded to the presentation of these odorants in the test trials. However, we found that locusts trained with citral and benzaldehyde showed no palp-opening response during the testing phase (**Fig. 5b, c**).

Next, we examined whether locusts could be conditioned when the reward was delayed until half a second after the termination of the conditioned stimulus (i.e., ‘OFF-training paradigm’). For this set of experiments, we only used hexanol and benzaldehyde as the conditioned stimuli (**Fig. 5d**). Once again, our results indicated that only locusts trained with hexanol robustly responded with PORs to the trained odorant in the testing phase. However, the POR dynamics observed in OFF-paradigm trained locusts were noticeably different from those we noted in the ON-training paradigm case. In the ON-training case, we found that locust PORs began immediately after the onset of the CS, lasted the duration of the stimulus, and the palps began to close following the termination of the stimulus. The peak of the PORs always occurred during the CS presentations. In contrast, for the OFF-training case, locust PORs were significantly slower (**Supplementary Fig. 4**), and the peak of the PORs in many locusts occurred after the termination of the stimulus.

In sum, these results indicate that only some odorants can successfully be associated with the food reward. Furthermore, both presentations during and after the termination of the stimulus can lead to odor-reward association but the behavioral response dynamics are significantly different between the two cases.

A linear model predicts behavioral response dynamics and cross-learning

Next, we wondered how locusts conditioned with a particular odorant (i.e., ‘the training odor’) respond when tested using other untrained odorants i.e., how olfactory learning generalizes between odorants. Our results indicate that locusts trained with hexanol responded robustly to presentations of isoamyl acetate (another odorant with a positive valence; **Fig. 6a**). Exposures to citral and benzaldehyde evoked no responses in hexanol-trained locusts. Surprisingly, locusts trained with citral and benzaldehyde showed no responses to the trained odorant, but a significant fraction of them showed PORs to hexanol and isoamyl acetate (**Fig. 6b**). For the OFF-training paradigm, we found that learning/cross-learning was observed only in those locusts that received rewards within 2 s of the termination of the conditioned stimulus. Interestingly, a large fraction of locusts (~60%) that received reward immediately after the termination of benzaldehyde (0.5 s after cessation) responded to hexanol and isoamyl acetate.

How predictable are these behavioral response dynamics and memory cross-talks given the neural responses evoked by these four odorants? To understand this, we set up determining the neural-behavioral transformation as a regression problem with sparsity constraints. For each training paradigm, the goal was to predict the POR responses to all four odorants examined given the time-varying ensemble neural

responses evoked by each odorant. Six such regression problems were set up, one for each training paradigm used in our study. We found that POR responses to all four odorants could be predicted reliably for all cases (red curves, **Fig. 6a**). We found that a linear mapping could indeed be found where the POR dynamics predicted from the neural responses were in good agreement with those observed in behavioral experiments (**Fig. 6a**; black (actual) vs. red (predicted); **Supplementary Fig. 5**). Notably, the regression weights assigned to different PNs to predict the POR for each training paradigm were highly similar (**Fig. 6c**). This result indicated that the mapping between neural responses and the PORs are highly consistent since the main trend observed in all cases were PORs to positive valence odorants (hex and iaa) and a lack of response to those with negative valence (citral and bzald). Consistent with this interpretation, we found that those PNs that received the most positive weights in the linear regression responded strongly to both positive valence odorants and had little to no responses to exposures of benzaldehyde and citral (**Fig. 6d**). On the other hand, PNs that responded strongly to the negative valence odorants and had transient responses at the onset and offset of both positive valence odorants received the most negative weights. More importantly, the negatively weighted PNs showed stronger spiking activities to the non-appetitive odorants, which allowed the suppression of POR responses (**Fig. 6d**; gray traces taller than black traces for benzaldehyde and citral).

In sum, these results indicate that the behavioral responses' strength and dynamics evoked by different odorants could be predicted from time-varying ensemble neural responses observed in the antennal lobe, and that a robust linear mapping involving ~50% of the total neurons was sufficient to transform neural activity into POR output.

A spatiotemporal coding logic for encoding appetitive odor preferences

Are the neural responses to appetitive and non-appetitive odorants organized in an interpretable fashion to explain the diverse set of neural and behavioral observations? To understand this, we visualized the ensemble neural activities of different odorants during both the ON and OFF periods. As can be observed, the odor-evoked ensemble responses were organized into four well-defined subspaces/clusters: appetitive ON, appetitive OFF, non-appetitive ON, and non-appetitive OFF (**Fig. 7a, b**; non-appetitive cluster includes odorants with both neutral and negative valences). Note that the different directions in this coding space indicate different combinations of PN responses, and nearby regions indicate pattern-matched neural responses. Therefore, these results indicate that while the neural activities during appetitive odorant exposures varied from one odorant to another (**Fig. 7a, b** – ensemble 1), they were still constrained to exploit only a limited combination of PN responses and therefore restricted to a particular subspace/region in this coding space. Extending this logic, these results also indicate that ensemble activities after the termination of appetitive odorants (**Fig. 7a, b** – ensemble 2), during exposures to non-appetitive odorants (**Fig. 7a, b** – ensemble 3), and after cessation of the non-appetitive stimuli (**Fig. 7a, b** – ensemble 4) all employed restricted combinations of ensemble neural responses that were different from each other.

Notably, the variance in neural responses evoked by appetitive odorants primarily spanned a low-dimensional space (i.e., a 'neural manifold') that contained ensembles 1 and 2. Only odorants that evoked neural responses limited to this manifold could be associated with food rewards (therefore referred to as the 'learning manifold'; **Fig. 7a**). Presenting the reward during activation of either neural ensemble 1 or ensemble 2 led to learning. However, the behavioral response dynamics significantly varied depending on whether the reward overlapped with ensemble 1 or ensemble 2 (**Supplementary Fig. 4**). In contrast, the variance in neural responses evoked by non-appetitive odorants spanned a different manifold that contained ensembles 3 and 4. Presenting reward during the activation of either of these ensembles did not result in successful conditioned stimulus-reward associations (therefore referred to as the 'non-learning

manifold'). Remarkably, our results indicate that similar odor-evoked response manifolds were also observed when neural responses were monitored in behaving preparations (**Supplementary Fig. 6**).

In sum, these results reveal an organizational logic for patterning spatiotemporal ensemble neural responses to mediate both innate and acquired odor-driven appetitive preferences.

Discussion

In this study, we examined the neural correlates of innate and acquired olfactory preferences. Our results indicate that while the neural responses evoked by an odorant are patterned over combinations of neurons activated and over time, the ensemble neural responses are still constrained by the overall behavioral relevance of the chemical cue. Odorants that have a positive appetitive preference, or valence, evoked ensemble neural responses that overlapped during odor presentations (i.e., ON responses) and after their terminations (i.e., OFF responses). Similarly, odorants with a neutral or negative appetitive preference evoked spiking activities that formed similar ON and OFF response clusters that were distinct from the appetitive response clusters. As a direct consequence of this spatiotemporal organization of neural responses, the innate behavioral responses were entirely predictable from neural responses during both these epochs but using distinct subsets of neurons.

Furthermore, our results indicate that delivering gustatory rewards during ON and OFF response epochs of odorants with positive appetitive valences alone resulted in successful Pavlovian conditioning. Reinforcing non-appetitive odorants did not generate successful odor-reward associations, but resulted in an increase in behavioral responses to other odorants with a positive valence. Notably, a linear model could map neural responses evoked by the odorants onto behavioral response dynamics and cross-associations learned. In sum, our results reveal a spatiotemporal coding logic that supports encoding both innate and acquired odor-driven appetitive preferences.

Chemical (input) vs. neural vs. behavioral (output) spaces

Could the observed appetitive preferences to different odorants be predicted directly from the stimulus/chemical space^{33,36,37}? We found that chemical features such as those extracted by nuclear magnetic resonance spectra or infrared spectra did not have good correlations with the overall appetitive preferences for different chemicals on the odor panel (**Supplementary Fig. 7**). Our results indicate that chemically similar odorants evoked divergent neural responses (isoamyl acetate and ethyl acetate – both esters but opposite valences). Conversely, we found odorants that had different chemical features mapped onto similar appetitive preferences (benzaldehyde and cyclohexanone). Even features such as the vapor pressure that controls the number of molecules reaching the antenna did not have a good correlation with the overall behavioral preference. While this is not an exhaustive list of chemical features that can be extracted, these results appear to indicate that it would be difficult to find a simple linear mapping of the chemical space onto the behavioral space. Similar results have recently been reported in the mouse olfactory bulb³⁸. Contrasting the non-linearity between the chemical – neural transformations, a linear mapping was indeed found between neural and behavioral spaces. These results support the idea that neural responses, even in those circuits very early in the olfactory pathway, are organized to generate appropriate behavioral outcomes rather than faithfully represent the chemical features of the odorants.

Surprisingly, at the individual neuron level, we found that responses in a small subset of PNs had a strong correlation with the overall innate preference for different odorants (**Fig. 2e**; correlations > 0.75 for 4/89 PNs for ON responses and 2/89 PNs for OFF responses). Such encoding of overall odor valence by individual neurons so early in the olfactory pathway has indeed been reported in other invertebrate models²⁻⁴. While the simplest model to predict the behavioral outcomes from the neural activity would be

to just use a few of these neurons, whether such a model would be robust is unclear. Earlier studies have indeed shown that individual projection neurons responses change unpredictably with changes in stimulus dynamics, intensity, competing cues, stimulus history, and ambient conditions^{16,22,23,39,40}. Notably, the behavioral recognition of odorants was found to remain invariant under a battery of these perturbations⁴¹. Therefore, a more robust and fault-tolerant model to overcome such variations in neural responses that arise due to natural perturbations would involve a combinatorial read-out of the ensemble activity as proposed in our regression analyses.

Innate vs. acquired appetitive preferences

To understand the appetitive preferences of locusts to different odorants, we used the palp-opening responses that locusts use to grab food. While preferences of individual locusts to the odor panel were idiosyncratic (i.e., varied from one locust to another; see **Fig. 1b**), as a group they tended to have similar behavioral preferences (**Fig. 1f**). This simple readout provided a one-dimensional quantitative summary of the innate appetitive preferences for the different odorants used in our panel. We found that a simple linear regression was sufficient to map ensemble neural responses during both stimulus presentation and after termination onto this behavioral dimension. Therefore, we concluded that the neural responses were spatiotemporally formatted to support the generation of innate behavioral outcomes.

Prior studies have shown that the palp-opening responses to an odorant could also be learned through associative conditioning^{22,42}. To understand the rules that constrain learning in this paradigm we screened and identified locusts that did not have any innate responses. We were concerned that repeated exposures to an odorant may induce PORs in these locusts. In this scenario, the PORs observed in the testing phase may not arise from conditioning but rather from sensitization due to repeated exposures to a stimulus. However, our results indicate that when the introductions of the reward were delayed to occur well after the termination of the odorant (hexanol OFF 4 s and benzaldehyde OFF 4 s paradigms), locusts did not show PORs and maintained their lack of responses to the conditioning odorants (**Fig. 6b**). We interpreted this result as an appropriate control indicating that locusts did not become sensitized to generate PORs to the conditioned stimulus and that PORs in these locusts were observed only in certain scenarios that suited associative learning.

Our conditioning experiments revealed that only two of the four odorants (hex and iaa) used resulted in a successful association between the odorant and the reward. As a result, locusts responded with PORs to the presentation of these odorants during the testing phase. We also observed generalization of the learned PORs to other odorants. Locusts trained with hexanol also showed responses to isoamyl acetate and vice versa (generalization to similar odors). Intriguingly, locusts trained with citral and benzaldehyde also increased PORs to hexanol and isoamyl acetate (cross-learning could also alter behavioral responses to unrelated odorants). We again found that linear mapping between neural and behavioral responses existed and captured all the important trends in our data (**Fig. 6a**).

We found that delaying reward such that it was delivered either during the presentation of hexanol (ON-training paradigm) or immediately after its termination (OFF-training paradigm) both resulted in associative learning. However, we found that the POR dynamics were different between these two training paradigms. We note that locusts in the ON-training paradigm had PORs that were significantly different from those observed in locusts trained using the OFF-paradigm. This result suggests that the timing of the reward could be controlled to coincide during different phases of neural response dynamics and such manipulations result in predictable changes in behavioral responses.

What potential mechanism could provide a neural correlate for cross-learning observed in our conditioning experiments? As can be noted from our regression analyses (**Fig. 6**), the two subsets of

neurons weighted opposingly (i.e., one group (black) received positive and the other (gray) was assigned negative weights) were sufficient to predict the behavioral responses observed to different odorants. Depending on the identity of the odorant, the set of neurons activated altered the overall balance between these two groups of neurons thereby generating or suppressing PORs. Therefore, if initially broad responses to appetitive and non-appetitive odorants were refined as a result of learning, such changes would alter the balance between these two competing neural ensembles and generate PORs selectively to appetitive odorants. One mechanism that was recently reported to remap neural responses and enhance the representation of stimulus valence was by increasing global inhibitory interactions⁴³. Whether such a mechanism could also result in our behavioral observations remains to be explored.

Neural manifolds for generating and patterning behavioral outcomes

In this study, our datasets comprised of neural responses evoked by a panel of diverse odorants, and their innate and acquired appetitive preferences. Surprisingly, we found that there exists a theoretical framework that would allow us to integrate these observations and understand the neural underpinnings of behavior. We regarded the ensemble neural activity as a high-dimensional neural response trajectory. Each odor-evoked response trajectory consisted of two non-overlapping segments, one during odor presentation (i.e., ON response), and the other after its terminations (i.e., OFF response). Notably, we found that ON responses and OFF responses evoked by innately appetitive odorants were on or near a low dimensional sub-space or ‘manifold’ (**Fig. 7a**). Similarly, we found that ON and OFF responses evoked by odorants with negative appetitive valence were on or near a separate low-dimensional manifold in the coding space (**Fig. 7a**).

We note that neuronal manifolds that encode for different behavioral response motifs have been reported in other model organisms^{36,44–46}. In *C. elegans*, these neuronal manifolds appear to arise globally and engage several circuits throughout the entire brain. Importantly, even those neuronal circuits that are directly downstream of sensory neurons were incorporated in these brain-wide dynamics to orchestrate the innate behavioral outcomes⁴⁴. If this is indeed a generic phenomenon, we would expect the spiking response patterns in the early olfactory circuits such as invertebrate antennal lobe or vertebrate olfactory bulb would be organized into behaviorally relevant neural manifolds. Our results indeed reveal that this is the case at least in the locust olfactory system.

Results from our conditioning experiments indicated that delivering rewards while the odor-driven neural activities were in the ‘appetitive manifold’ resulted in successful conditioning, whereas no associative learning occurred while delivering rewards during responses excursion in the ‘non-appetitive manifold’. Interpreted differently, this result suggests that neural activity patterns on some manifolds are conducive for learning, while activity patterns outside this manifold could be harder to learn. Similar results have been reported in the context of motor control in the primate motor cortex⁴⁷. While the motor cortex result arose from constraints imposed by the neural circuitry making certain neural activity patterns difficult to generate, here the antennal lobe network could generate neural response excursions in both learnable and non-learnable manifolds depending on the identity of the stimuli.

Acknowledgments

We thank members of the Raman Lab (Washington University in St. Louis) and Dr. Debajit Saha (Michigan State University) for feedback on the manuscript. We thank Pearl Olsen for insect care. This research was supported by NSF (1453022, 2021795) and ONR (N00014-19-1-2049, N00014-21-1-2343) grants to B.R.

Author contributions

RC and BR conceived the study and designed the experiments/analyses. RC performed all the experiments and analyzed the data. RC and BR wrote the paper. BR supervised all aspects of the work.

Methods

Odor stimulation

All odorants were delivered at a 1% v/v dilution in mineral oil and placed in dark 60-ml bottles. A constant background air stream (desiccated and filtered) at 0.75 L/min was used as the carrier stream for 0.1 L/min pulses of odorants. A large vacuum funnel placed directly behind the antenna allowed for the constant clearing of the odorants delivered.

For behavioral experiments to quantify innate appetitive preferences, each odorant in the panel was presented for one trial in a pseudorandomized order (**Fig. 1a**). Odorants were delivered by displacing a 0.1L/min of headspace in the odor bottles using a pneumatic picopump (WPI Inc., PV-820). Each odor pulse was 4 s and the inter-trial interval was 60 s.

For electrophysiology experiments, each odorant was presented for ten trials in a pseudorandomized order. To minimize interference during the experiment, we designed and built a custom olfactometer (SMC valves, NI-DAQ controller) that was automated and triggered using MATLAB. Each odor pulse was 4 s in duration, and the inter-pulse interval was 60 s.

Behavior experiments to characterize innate palp-opening responses

Young adult locusts of either sex were starved for 24 hours before the experiment. Locusts were immobilized within a plastic tube and their compound eyes were covered using black tape. All twenty odorants were diluted to 1% v/v as previously described. Hexanol alone was additionally diluted to 0.1% and 10% dilutions (i.e., a total of 22 odorants in the panel). Each locust was presented with all 22 odorants in a pseudorandomized order for 4 s pulses separated by 56 s inter-pulse intervals (60 s between the starts of two consecutive pulses). The experiments were recorded using a video camera (Microsoft). An LED was used to track stimulus onset/offset. The POR responses were scored offline in a blind fashion with no odorant information to remove any experimenter biases. Responses to each odorant were scored a 0 or 1 depending on if the palps remain closed or opened (**Fig. 1b**). A successful POR was defined as an opening of the maxillary palps beyond the facial ridges as shown on the locust schematic (**Fig. 1a**).

Preference Index

As noted above, locust responses to each odorant were binarized. The responses of all locusts to an odor were then summed to obtain a Total Score. A normalized score for each odorant was then calculated as follows:

$$Norm_score_{odor} = \frac{Total\ Score_{odor}}{Total\ \# \ locusts}$$

The preference index (**Fig. 1c**) was then calculated for each odorant by performing a median subtraction from the Norm_score as follows –

$$Preference\ index_{odor} = Norm_score_{odor} - Norm_score_{median}$$

$Norm_score_{median}$ was obtained by calculating the median across all odorants.

Vapor Pressure Analysis

Vapor pressure data for 18 odorants were obtained from an online database (The Good Scents Company)⁴⁸. Data for neem and garlic could not be obtained and these odors were omitted from our

analyses in **Fig. 1d**. Regression analysis was performed between vapor pressure values and the POR Total Scores. An R^2 value was obtained using the 'fitlm' function in MATLAB (**Fig. 1d**). One of the odorants in the panel (ethyl acetate) had a vapor pressure much higher than all other chemicals, and hence the weak correlations in **Fig. 1d** could be driven by this potential outlier. To control for this, a similar analysis was performed in **Supplementary Fig. 1b**, but using only seventeen odorants (i.e., excluding ethyl acetate).

Monte Carlo simulations for evaluating behavioral stability

We performed Monte Carlo simulations on the data shown in **Fig. 1b**. We randomly sampled locusts (' n ' ranging from 1 to 26) and calculated preference indices for all odors using POR scores using the selected subsets of locusts. For each n , we performed 100 such simulations and computed an average preference index, which was then compared with the preferences obtained using all twenty-two locusts. The mean correlation for each n is shown in **Fig. 1f**. Error bars indicate standard error of the mean (s.e.m.).

Electrophysiology experiments

Young adult locusts of either sex were used for these experiments⁴⁹. The legs and wings were removed, and they were immobilized on a custom platform. The head was fixed into place by a wax cup and the antennae were held in place inside a thin tube using epoxy glue. The cuticle above the brain was cut open, the air sacs covering the brain were removed, and the locusts were degutted to minimize any internal movements. A metal wire platform was then inserted underneath the brain to lift and stabilize it. Finally, the transparent sheath covering the brain was removed after applying protease enzyme.

Locust brains prepared this way were super-fused with artificial saline buffer and a reference electrode (Ag/Ag-Cl) was inserted into the saline. Multi-unit recordings were made from the antennal lobe projection neurons (PNs) using a 4x4 silicon probe (NeuroNexus) with impedance in the 200-300 k Ω range (**Fig. 2a**). Data were acquired at a 15 kHz sampling rate using a custom MATLAB program and filtered between 0.3-6 kHz using an amplifier system (Caltech) that provided a 10,000 gain.

Offline spike-sorting (IgorPro) was performed using the best 4 channels recorded⁵⁰. To identify single units (PNs), the following previously published criteria were used: unit cluster separation >5 noise s.d., number of spikes within 20 ms $<6.5\%$, and spike waveform variance <6.5 noise s.d. To account for baseline drift and loss of neurons during an experiment, we only included PNs with consistent baseline spiking activity in all 220 trials (22 odors, 10 trials each). We defined a PN as being consistent if its baseline firing rate (during 4 s period before odor presentation) in all trials was no less than 15% of the maximum baseline firing rate for that PN. A total of 89 PNs were identified using these criteria (originally acquired 131 PNs from 26 locusts).

PID experiment

We used a fast-photoionization diode (miniPID, Aurora Scientific) to characterize the stimulus delivery dynamics of all odors used in the electrophysiology experiments. Each odor was presented for 5 trials and PID signals were acquired at 15 kHz using a custom MATLAB program. The mean signals for all odors are shown in **Fig. 2b**.

Projection neuron response classification

We defined 4 s of odor presentation as an ON period, and the 4 s immediately following odor termination as an OFF period. PNs were classified as ON-responsive if the firing activity was 6.5 s.d. above mean baseline (2 s preceding the stimulus) firing activity in at least 5 of the 10 trials during the ON period. Similarly, PNs were classified as being OFF-responsive using a similar metric applied to the OFF period.

PNs were classified as ‘Inhibited’ if their firing activity did not exceed 2 s.d. of baseline in any time bin during odor presentation and the mean firing rate during the entire stimulus duration (4 s) was lower than mean baseline activity (in at least 5 out of 10 trials). These classifications are summarized for all odors in **Fig. 2d**.

Dimensionality reduction analysis

We used Principal Component Analysis (PCA) to visualize ensemble PN activity (**Fig. 3a; Fig. 7a**). The spiking activity for each PN during 4 s of odor presentation was averaged across all 10 trials and binned in 50 ms non-overlapping time bins. In this manner, we obtained an 89 PN x 80 time-bin matrix for each odorant. We concatenated these data matrices obtained for each odor to obtain an 89x1760 data matrix (80 bins * 22 odors). We then computed a covariance matrix (89x89) for this data matrix.

Each 89-dimensional response vector was then projected onto the top three eigenvectors (that captured the highest variance). For visualization, the first time-bin was subtracted from each odor to obtain a similar pre-stimulus baseline for all odors. The odor trajectories were smoothed using a three-point moving average low-pass filter.

Hierarchical clustering analysis

The spiking activity of each PN during 4 s of odor presentation was summed to obtain an 89x1 (89 PNs) vector per odorant. Agglomerative hierarchical clustering was performed on vectors for all 22 odors using the ‘linkage’ function in MATLAB. The odors were clustered based on a correlation distance metric, and the farthest pairwise distance between clusters was minimized. The clustering was visualized using the ‘dendrogram’ function (**Fig. 3b**) after obtaining a leaf ordering using the ‘optimalleaforder’ function.

Linear regression to predict valence from PN activity

Mean odor-evoked activity for each PN (n_i) was used as the input for the linear regressor and the behavioral Norm_score for each odor was used as the output. A softmax layer was added to ensure that the final prediction was always between 0 and 1. A leave-one-out-cross-validation (LOOCV) approach was used, where the model weights were trained using data for 21 odors using gradient descent, and then the neural response for the test odorant was used to predict the behavioral POR preference index. The mean squared error cost function was minimized.

$$Predicted\ POR = softmax \left(\sum_{i=1}^{89} w_i * n_i + bias \right)$$

Where n_i is the number of spikes evoked during odor exposure in PN_i , and w_i is the weight assigned by the linear regressor for PN_i .

As controls for the regressors, the POR preference indices of different odorants were shuffled randomly before training. We used the entire 4 s of PN activities during odor presentation for the ON-regressor, and 4 s of OFF activity immediately following odor termination for the OFF-regressor (**Fig 4**).

Monte Carlo simulations for electrophysiology

We performed Monte Carlo simulations to gauge the performance of the linear regressors as a function of the number of PNs used for the analysis was varied. To achieve this, we randomly sub-sampled n (where n ranged from 1 to 89) PNs and quantified the predictive performance using mean squared error (MSE).

For each n , we performed 1000 simulations and reported the average MSE (**Supplementary Fig. 2**). We performed these simulations for both the ON- and OFF-regressors.

Behavior experiments – Classical conditioning

Appetitive classical conditioning experiments were performed on young adult locusts of either sex starved for 24 hours before the experiment. Locusts were immobilized within a plastic tube, their eyes were closed using black tape, and their maxillary palps were painted using a zero-volatile-organic-chemical green paint (Valspar ultra). A brief 20-minute buffer period was allowed for paint to dry and the locust to acclimatize back to baseline activity levels.

Prior to conditioning, each locust was presented with a 4 s pulse of all four odorants used in the experiment (hexanol, isoamyl acetate, benzaldehyde, and citral). If a locust had a palp-opening response to any of these odorants, it was deemed ‘pre-conditioned’ and was discarded from the experiment. A 15-minute buffer was allowed between this pre-test and the training phase.

During the training phase, locusts were presented the training odorant diluted at 1% v/v at a rate of 0.1 L/min diluted in a constant background air stream (desiccated and filtered) of 0.75 L/min. A vacuum funnel placed behind the locust allowed for odor clearance. The odor was presented for 10 s and a food reward (wheat grass) was presented at 5 s post-odor onset for ON-conditioning. The odor was presented for 10 s and a food reward (wheat grass) was presented at 0.5 s, 2 s, or 4 s post-odor termination for OFF-conditioning. Six such training trials were performed with an inter-trial interval of 10 minutes. Locusts that met the training criteria (>3 food reward acceptances out of 6) were then evaluated in the testing phase.

During the testing phase, locusts were presented with 4 s pulses of various odorants (at 1% dilution) in a pseudorandomized manner with a minimum interval of 20 minutes between successive tests. The palp-opening responses of the locusts were recorded using a video camera (Microsoft) at 30 fps. The odor delivery and video acquisition were synced using a custom LabView program.

Locusts were kept on a 12 h day – 12 h night cycle (7 am – 7 pm day). All behavioral experiments were performed between 10 am – 3 pm to ensure that the training phase coincided with the daily feeding time for the locusts.

Palp-tracking algorithm

To accurately track maxillary palp separation, we trained a UNet convolutional neural network using randomized initialization of weights in Keras and Tensorflow⁵¹. During the training phase, the input into this network was a single channel (green) 128x128 image cropped around the palps. The outputs were manually labeled palps (as binarized 128x128 matrices with 1’s indicating palps and 0’s indicating no palps). We trained the network using the Adam optimizer and binary cross-entropy loss function. We performed image augmentation using the ‘imgaug’ Python library and trained the network on approximately 2000 labeled frames.

Videos were input into the trained network frame-by-frame and the output was thresholded and binarized using a combination of Otsu, mean, and triangle filters from the ‘skimage’ library. Palp distance for each frame was calculated as the distance between the centroids of the two predicted palps using the ‘regionprops’ function.

Responsive Locusts

Locusts were considered ‘responsive’ to a particular odor if they had a palp-opening response that was >6.5 s.d. above pre-stimulus baseline (2 s) for at least 30 time-frames (1 s) with palp separation > 1.5 a.u. (which was the noise threshold of the tracking algorithm) (**Fig. 5b, e; Fig. 6b**).

Individual locust responses

For the normalized POR traces shown in **Fig. 5c, d**, we scaled each locust’s response such that 0 corresponded to the minimum palp separation and 1 corresponded to the maximum palp separation the locust had across all test odors. Note that after each training paradigm, we tested locusts on four odors – hexanol, isoamyl acetate, benzaldehyde, and citral.

Mapping neural responses onto palp-opening response dynamics

PN activity and POR responses (distance between palps) for hexanol, isoamyl acetate, benzaldehyde, and citral were averaged across trials and down-sampled to 10 Hz. For each odor, we used 2 s baseline, 4 s of odor presentation, and 4 s after odor termination to obtain a 10 s vector (100 elements at 10 Hz). We then concatenated responses from all 4 odors to obtain 400-dimensional vectors. The input data was hence 89×400 (89 PNs; spiking activity at each time point) and the output was 400×1 (palp-separation at each time point). A regularized model was fitted using ‘lasso’ (sklearn in Python) with an ‘alpha’ value of 0.01. The learned 89×1 weights were then used with the input data to generate predicted POR responses shown in red in **Fig. 6a**.

We trained 6 such models for each training condition shown in **Fig. 6a**. The weights obtained for all 6 models were sorted using the weights from the hexanol-ON model and are shown in **Fig. 6c**. The inset plot shows pair-wise correlations between each weight vector pair. The weights across all six models were averaged for each PN. 21/89 PNs had a weight > 0 and 19 PNs had a weight < 0 , with the remainder of PNs assigned a weight of 0 due to regularization. The PSTH’s of the PNs assigned positive and negative weights are shown for all 4 odors in **Fig. 6d**.

Monitoring neural responses in behaving locusts

We developed a minimally invasive preparation to facilitate the monitoring of projection neuron responses in locusts while they were classically conditioned. In brief, the locusts were immobilized identically to the procedure followed for the prior classical conditioning experiments (see above). A small cut was made in their cuticle to allow access to the antennal lobe, which was stabilized using a metal-wire platform. Finally, the antennal lobe was de-sheathed to allow electrode implantation. The neural recordings were performed similarly to the previous set of electrophysiology experiments.

Before conditioning, we recorded 5 trials of responses to each of the 6 odors used (appetitive – hexanol, isoamyl acetate, 2-octanol; non-appetitive – cyclohexanone, benzaldehyde, citral). After a 15-minute gap, we performed the conditioning as follows - locusts were presented with 6 trials of trained odor (hexanol or benzaldehyde) with overlapping presentations of a food reward (sucrose in water 1g/10ml concentration) similar to conditioning methods described above. To minimize movement of the locust and conserve neural stability, we switched from solid food reward (grass) to liquid food reward (sucrose in water) and presented it in an automated manner using a pneumatic pump (WPI Inc., PV-820). The inter-trial interval was set to 3 minutes for the training phase. Post-training, we waited for 15 minutes and then repeated the presentations of all 6 odors for 5 trials each. In all blocks of neural recordings, we pseudorandomized the order of odor presentation.

The neural data acquired in these experiments could not be reliably spike sorted using the approach mentioned above. As a result, we used an alternative approach for processing this dataset⁵². The raw data signals (acquired at 15 kHz) were de-noised using a band-pass between 300 Hz and 6000 Hz followed by clipping of signals 5 s.d. above or below the baseline level. These were then passed through a continuous moving root-mean-squared (RMS) filter with a 20 ms window (DSP toolbox on MATLAB), down-sampled by a factor of 150, smoothed by a 10-point moving average filter, and finally down-sampled by a factor of 5 to produce a temporal resolution of 20 Hz (50 ms, similar to spike sorted PN responses). The samples were finally baseline subtracted using the mean of 1 s baseline prior to odor presentation (two sample recordings shown in **Supplementary Fig. 6a**) to obtain the Δ RMS signal. For the PCA analysis shown in **Supplementary Fig. 6b**, we followed a similar approach as mentioned above. We used the mean of 4 s of odor presentation and 4 s of responses immediately after odor termination to obtain a 160-dimension vector for each odor (8 seconds x 20 samples per second) for each locust. We recorded from 10 locusts each for hexanol and benzaldehyde training experiments and concatenated these neural responses to obtain a final 20 locust x 160 bin response matrix for each odor during both the pre- and post-training periods.

References

1. Grosjean Y, Rytz R, Farine J-P, et al. An olfactory receptor for food-derived odours promotes male courtship in *Drosophila*. *Nature*. 2011;478(7368):236-240. doi:10.1038/nature10428
2. Semmelhack JL, Wang JW. Select *Drosophila* glomeruli mediate innate olfactory attraction and aversion. *Nature*. 2009;459:218-223. doi:10.1038/nature07983
3. Knaden M, Strutz A, Ahsan J, Sachse S, Hansson BS. Spatial representation of odorant valence in an insect brain. *Cell Rep*. 2012;1(4):392-399. doi:10.1016/j.celrep.2012.03.002
4. Kreher SA, Mathew D, Kim J, Carlson JR. Translation of sensory input into behavioral output via an olfactory system. *Neuron*. 2008;59(1):110-124. doi:10.1016/j.neuron.2008.06.010
5. Kermen F, Midroit M, Kuczewski N, et al. Topographical representation of odor hedonics in the olfactory bulb. *Nat Neurosci*. 2016;19(7):876-878. doi:10.1038/nn.4317
6. Rengarajan S, Hallem EA. Olfactory circuits and behaviors of nematodes. *Curr Opin Neurobiol*. 2016;41:136-148. doi:10.1016/j.conb.2016.09.002
7. Wilson CD, Serrano GO, Koulakov AA, Rinberg D. A primacy code for odor identity. *Nat Commun*. 2017;8(1):1477. doi:10.1038/s41467-017-01432-4
8. Abraham NM, Spors H, Carleton A, Margrie TW, Kuner T, Schaefer AT. Maintaining Accuracy at the Expense of Speed: Stimulus Similarity Defines Odor Discrimination Time in Mice. *Neuron*. 2004;44(5):865-876. doi:10.1016/j.neuron.2004.11.017
9. Uchida N, Mainen ZF. Speed and accuracy of olfactory discrimination in the rat. *Nat Neurosci*. 2003;6(11):1224-1229. doi:10.1038/nn1142
10. Szyszka P, Gerkin RC, Galizia CG, Smith BH. High-speed odor transduction and pulse tracking by insect olfactory receptor neurons. *Proc Natl Acad Sci*. 2014;111(47):16925-16930. doi:10.1073/pnas.1412051111
11. Cassenaer S, Laurent G. Conditional modulation of spike-timing-dependent plasticity for olfactory learning. *Nature*. 2012;482(7383):47-52. doi:10.1038/nature10776
12. Ito I, Ong RC, Raman B, Stopfer M. Sparse odor representation and olfactory learning. *Nat Neurosci*. 2008;11(10):1177-1184. doi:10.1038/nn.2192
13. Laurent G, Davidowitz H. Encoding of olfactory information with oscillating neural assemblies. *Science*. 1994;265(5180):1872-1875. doi:10.1126/science.265.5180.1872
14. Raman B, Joseph J, Tang J, Stopfer M. Temporally Diverse Firing Patterns in Olfactory Receptor Neurons Underlie Spatiotemporal Neural Codes for Odors. *J Neurosci*. 2010;30(6):1994-2006. doi:10.1523/JNEUROSCI.5639-09.2010
15. Stopfer M, Jayaraman V, Laurent G. Intensity versus identity coding in an olfactory system. *Neuron*. 2003;39(6):991-1004. doi:10.1016/j.neuron.2003.08.011
16. Saha D, Leong K, Li C, Peterson S, Siegel G, Raman B. A spatiotemporal coding mechanism for background-invariant odor recognition. *Nat Neurosci*. 2013;16(12):1830-1839. doi:10.1038/nn.3570

17. Wilson RI. Early Olfactory Processing in *Drosophila*: Mechanisms and Principles. *Annu Rev Neurosci.* 2013;36:217-241. doi:10.1146/annurev-neuro-062111-150533
18. Faber T, Joerges J, Menzel R. Associative learning modifies neural representations of odors in the insect brain. *Nat Neurosci.* 1999;2(1):74-78. doi:10.1038/4576
19. Cassenaer S, Laurent G. Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature.* 2007;448(7154):709-713. doi:10.1038/nature05973
20. Gupta N, Stopfer M. Functional Analysis of a Higher Olfactory Center, the Lateral Horn. *J Neurosci.* 2012;32(24):8138-8148. doi:10.1523/JNEUROSCI.1066-12.2012
21. Saha D, Li C, Peterson S, Padovano W, Katta N, Raman B. Behavioural correlates of combinatorial versus temporal features of odour codes. *Nat Commun.* 2015;6(1):6953. doi:10.1038/ncomms7953
22. Nizampatnam S, Saha D, Chandak R, Raman B. Dynamic contrast enhancement and flexible odor codes. *Nat Commun.* 2018;9(1):3062. doi:10.1038/s41467-018-05533-6
23. Saha D, Sun W, Li C, et al. Engaging and disengaging recurrent inhibition coincides with sensing and unsensing of a sensory stimulus. *Nat Commun.* 2017;8(1):15413. doi:10.1038/ncomms15413
24. Fdez Galán R, Sachse S, Galizia CG, Herz AVM. Odor-driven attractor dynamics in the antennal lobe allow for simple and rapid olfactory pattern classification. *Neural Comput.* 2004;16(5):999-1012. doi:10.1162/089976604773135078
25. Yamazaki D, Hiroi M, Abe T, et al. Two Parallel Pathways Assign Opposing Odor Valences during *Drosophila* Memory Formation. *Cell Rep.* 2018;22(9):2346-2358. doi:10.1016/j.celrep.2018.02.012
26. Vickers NJ, Christensen TA, Baker TC, Hildebrand JG. Odour-plume dynamics influence the brain's olfactory code. *Nature.* 2001;410(6827):466-470. doi:10.1038/35068559
27. Menzel R. Associative learning in honey bees. *Apidologie.* 1993;24(3):157-168. doi:10.1051/apido:19930301
28. Giurfa M. Cognition with few neurons: higher-order learning in insects. *Trends Neurosci.* 2013;36(5):285-294. doi:10.1016/j.tins.2012.12.011
29. Hall JF. Backward conditioning in Pavlovian type studies. Reevaluation and present status. *Pavlov J Biol Sci.* 1984;19(4):163-168. doi:10.1007/BF03004514
30. Hansche WJ, Grant DA. Onset versus termination of a stimulus as the CS in eyelid conditioning. *J Exp Psychol.* 1960;59(1):19-26. doi:10.1037/h0041407
31. Giurfa M, Fabre E, Flaven-Pouchon J, et al. Olfactory conditioning of the sting extension reflex in honeybees: Memory dependence on trial number, interstimulus interval, intertrial interval, and protein synthesis. *Learn Mem Cold Spring Harb N.* 2009;16(12):761-765. doi:10.1101/lm.1603009
32. Zhang L, Guo M, Zhuo F, Xu H, Zheng N, Zhang L. An odorant-binding protein mediates sexually dimorphic behaviors via binding male-specific 2-heptanone in migratory locust. *J Insect Physiol.* 2019;118:103933. doi:10.1016/j.jinsphys.2019.103933

33. Keller A, Gerkin RC, Guan Y, et al. Predicting human olfactory perception from chemical features of odor molecules. *Science*. 2017;355(6327):820-826. doi:10.1126/science.aal2014
34. Mazor O, Laurent G. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*. 2005;48(4):661-673. doi:10.1016/j.neuron.2005.09.032
35. Bathellier B, Buhl DL, Accolla R, Carleton A. Dynamic Ensemble Odor Coding in the Mammalian Olfactory Bulb: Sensory Information at Different Timescales. *Neuron*. 2008;57(4):586-598. doi:10.1016/j.neuron.2008.02.011
36. Haddad R, Weiss T, Khan R, et al. Global Features of Neural Activity in the Olfactory System Form a Parallel Code That Predicts Olfactory Behavior and Perception. *J Neurosci*. 2010;30(27):9017-9026. doi:10.1523/JNEUROSCI.0398-10.2010
37. Kepple D, Koulakov A. Constructing an olfactory perceptual space and predicting percepts from molecular structure. *ArXiv170805774 Q-Bio*. Published online June 6, 2018. Accessed July 27, 2021. <http://arxiv.org/abs/1708.05774>
38. Chae H, Kepple DR, Bast WG, Murthy VN, Koulakov AA, Albeanu DF. Mosaic representations of odors in the input and output layers of the mouse olfactory bulb. *Nat Neurosci*. 2019;22(8):1306-1317. doi:10.1038/s41593-019-0442-z
39. Brown SL, Joseph J, Stopfer M. Encoding a temporally structured stimulus with a temporally structured neural representation. *Nat Neurosci*. 2005;8(11):1568-1576. doi:10.1038/nn1559
40. Broome BM, Jayaraman V, Laurent G. Encoding and decoding of overlapping odor sequences. *Neuron*. 2006;51(4):467-482. doi:10.1016/j.neuron.2006.07.018
41. Nizampatnam S, Zhang L, Chandak R, Katta N, Raman B. Invariant Odor Recognition with ON-OFF Neural Ensembles. *bioRxiv*. Published online November 8, 2020:2020.11.07.372870. doi:10.1101/2020.11.07.372870
42. Simões P, Ott SR, Niven JE. Associative olfactory learning in the desert locust, *Schistocerca gregaria*. *J Exp Biol*. 2011;214(15):2495-2503. doi:10.1242/jeb.055806
43. Frank T, Mönig NR, Satou C, Higashijima S, Friedrich RW. Associative conditioning remaps odor representations and modifies inhibition in a higher olfactory brain area. *Nat Neurosci*. 2019;22(11):1844-1856. doi:10.1038/s41593-019-0495-z
44. Kato S, Kaplan HS, Schrödel T, et al. Global Brain Dynamics Embed the Motor Command Sequence of *Caenorhabditis elegans*. *Cell*. 2015;163(3):656-669. doi:10.1016/j.cell.2015.09.034
45. Briggman KL, Abarbanel HDI, Kristan WB. Optical Imaging of Neuronal Populations During Decision-Making. *Science*. 2005;307(5711):896-901. doi:10.1126/science.1103736
46. Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. *J Neurophysiol*. 2009;102(1):614-635. doi:10.1152/jn.90941.2008
47. Sadtler PT, Quick KM, Golub MD, et al. Neural constraints on learning. *Nature*. 2014;512(7515):423-426. doi:10.1038/nature13665

48. The Good Scents Company - Aromatic/Hydrocarbon/Inorganic Ingredients Catalog information. Accessed July 22, 2021. <http://www.thegoodscentcompany.com/data/rw1002711.html>
49. Saha D, Leong K, Katta N, Raman B. Multi-unit Recording Methods to Characterize Neural Activity in the Locust (*Schistocerca Americana*) Olfactory Circuits. *JoVE J Vis Exp*. 2013;(71):e50139. doi:10.3791/50139
50. Pouzat C, Mazor O, Laurent G. Using noise signature to optimize spike-sorting and to assess neuronal classification quality. *J Neurosci Methods*. 2002;122(1):43-57. doi:10.1016/S0165-0270(02)00276-5
51. Ronneberger O, Fischer P, Brox T. U-Net: Convolutional Networks for Biomedical Image Segmentation. *ArXiv150504597 Cs*. Published online May 18, 2015. Accessed July 22, 2021. <http://arxiv.org/abs/1505.04597>
52. Saha D, Mehta D, Altan E, et al. Explosive sensing with insect-based biorobots. *Biosens Bioelectron X*. 2020;6:100050. doi:10.1016/j.biosx.2020.100050
53. Chemical Name Search. Accessed July 22, 2021. <https://webbook.nist.gov/chemistry/name-ser/>

Figures

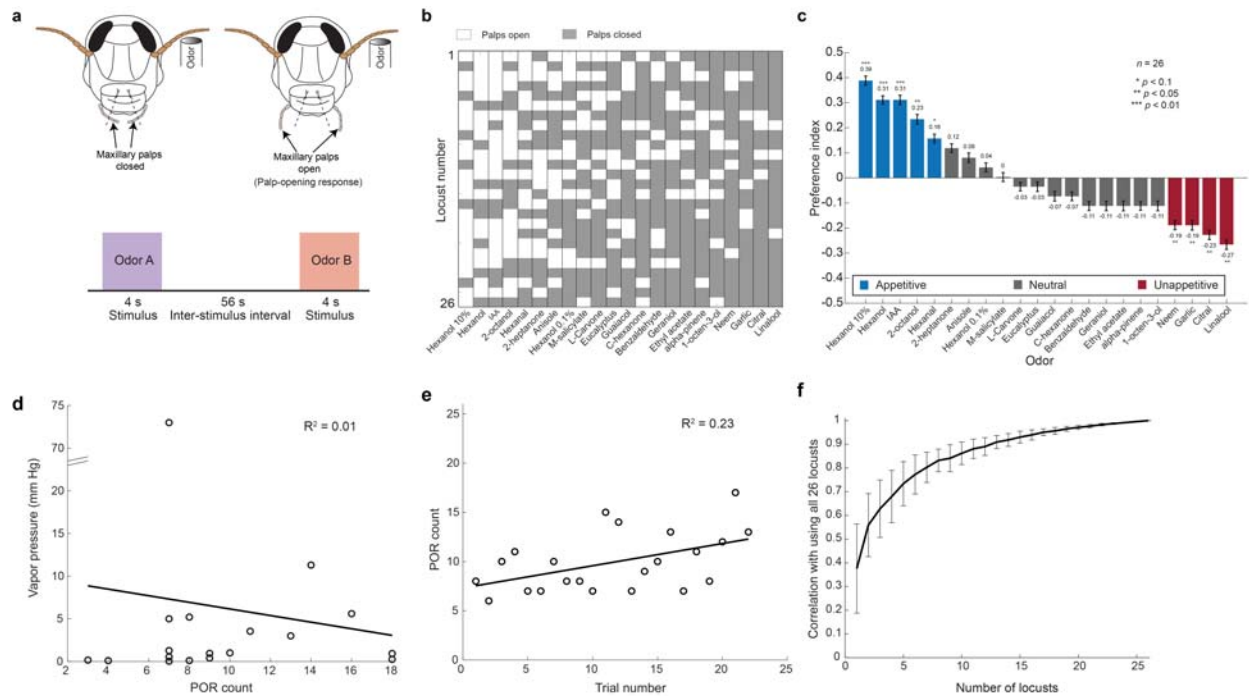


Figure 1

Figure 1: Innate appetitive preferences of locusts to a diverse odor panel

a) A schematic showing a palp-opening response (POR) and experimental protocol. A successful POR was defined as an opening of the maxillary palps beyond the facial ridges shown on the locust. Odors were delivered in a pseudorandomized order onto the locust antenna. The stimulus delivery was 4 s in duration, and the inter-stimulus interval was set to 56 s.

b) Innate preferences of twenty-six locusts for the twenty-two odorants tested are shown. Each row shows the POR responses of a locust to the odor panel. White boxes indicate a successful POR to an odor and gray boxes indicate no POR. Note that odors are sorted from those that elicited the highest number of PORs across locusts (leftmost) to the lowest (rightmost). Note that this ordering was just to facilitate the readability of data and does not represent the actual order in which each locust was tested.

c) Preference indices were calculated for all odors tested and are shown as a bar plot ($n = 26$ locusts). Blue bars indicate odors classified as appetitive, gray bars indicate neutral odors, and red bars indicate unappetitive odors. Locusts with a significant deviation from the median response (one-sided binomial test, $p < 0.1$), were classified as either being appetitive or unappetitive; * indicates $p < 0.1$, ** indicates $p < 0.05$, *** indicates $p < 0.01$). Error bars indicate s.e.m.

d) Regression analysis of odor vapor pressure versus number of PORs generated (across 26 locusts) is shown for all odorants in the panel. Each open circle indicates values (vapor pressure vs POR count) for one odorant. Only odorants with available vapor pressure data were considered for this analysis (18 out of 22 odors at 1% v/v concentration). Best fit line using a linear regression model is shown in black. The calculated R^2 value for this model is 0.01.

e) Regression analysis of POR counts versus trial number in the experiment is shown. Each circle indicates the number of locusts with successful PORs in that particular trial. Best fit line using a linear regression model is shown in black. The R^2 value calculated for this model is 0.23.

f) Results from Monte Carlo simulations are shown (see **Methods**). Valence predictions were obtained by using a random subset of locusts of a particular size (i.e. any n -locusts-out-of-26) and were compared with overall valence obtained using all 26 locusts using a correlation metric. For each number of locusts, 100 such simulations were performed with random subsets of locusts chosen in each simulation. The mean correlation and s.e.m. across the simulations are shown. An R^2 value above 0.95 was obtained for simulations with $n > 18$ locusts.

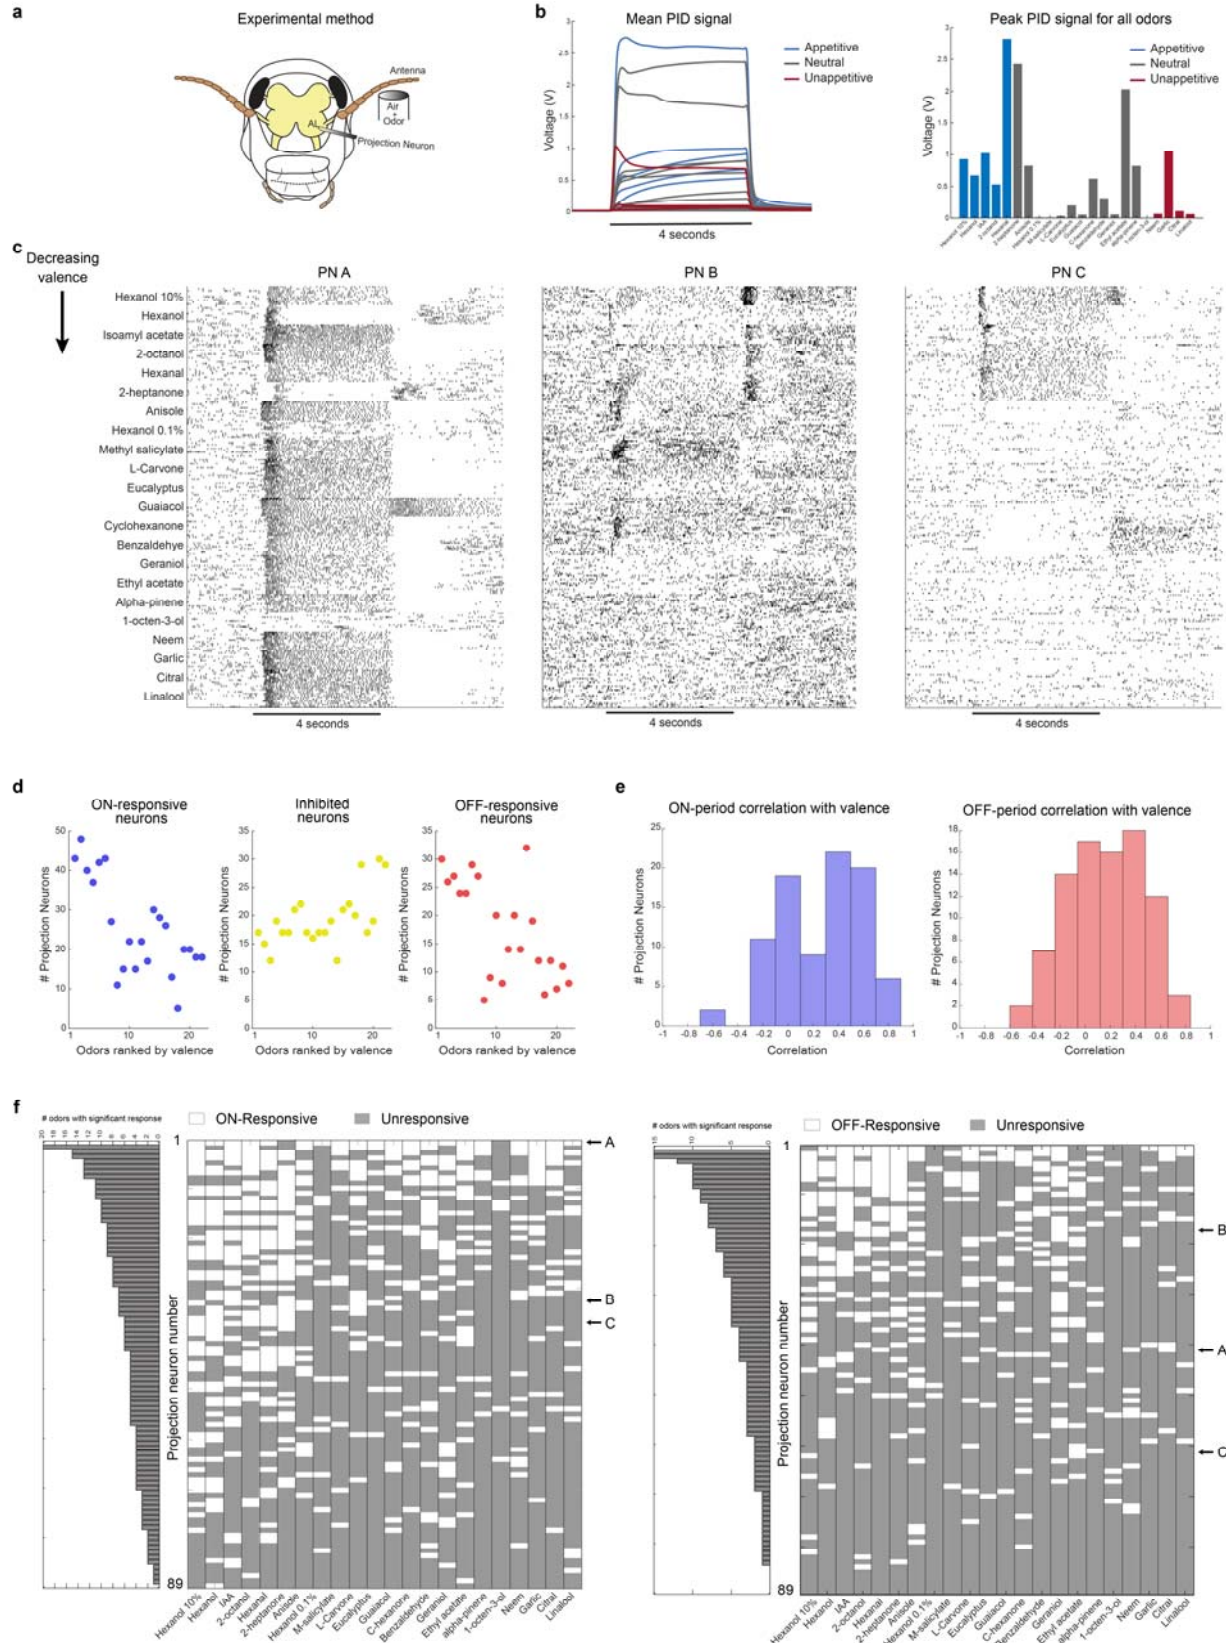


Figure 2

Figure 2: Individual PN responses to appetitive and non-appetitive odorants.

a) A schematic of the experimental setup is shown. Extracellular projection neuron (PN) recordings were made from the locust antennal lobe. A constant stream of background air was presented to the ipsilateral antenna at all times, with odor pulses being presented atop the air stream for 4 s in every 60 s trial. A vacuum behind the prep kept a constant flux across the antenna. Each odorant was presented for ten repeated trials and the order of odorants was pseudorandomized in each experiment.

b) Left: Mean voltage signals acquired from a photoionization detector are shown for all twenty-two odorants in the panel. Each odorant was presented for 4 s and repeated for 5 trials. Each trace is colored using preference indices obtained in **Fig. 1c**. Blue traces indicate appetitive odorants, gray indicates neutral odorants, and red indicates unappetitive odorants. **Right:** The peak voltage signal obtained from the photoionization detector is shown for all twenty-two odorants. Odorants are sorted from highest (leftmost) to lowest (rightmost) appetitive preference or valence. Same color convention as in the left panel.

c) Representative PN responses to all twenty-two odorants in the panel are shown. Each tick indicates an action potential, each row corresponds to one trial, and ten trial blocks are shown for each odorant. Odors are sorted based on their behavioral preferences, with the highest appetitive preferences shown as the top block of ten trials, and the lowest shown at the bottom (**Fig. 1**). A black bar along the x-axis indicates the four seconds odor presentation window.

d) Left panel: Number of PNs that were activated during the odor presentation window (ON-responsive) is plotted for twenty-two different odorants in the panel. Odorants are again sorted based on their appetitive valence (highest – leftmost to lowest – rightmost). **Middle and right panels:** Similar plots but showing the number of PNs that were inhibited during odor presentation, and that number of PNs activated after odor termination (OFF-responsive) are shown for different odorants on the panel. The odorants are again arranged based on appetitive valence (same as left panel).

e) Left panel: For each PN, we took the mean of the spiking activity across 4 s of odor presentation and across all 10 trials for each odor to obtain a 22-dimensional vector. Next, we computed the correlation between this vector and the appetitive preferences obtained for each odor (also a 22-dimensional vector; **Fig. 1**). The distribution of correlations obtained using this approach is shown for all 89 PNs. **Right panel:** Similar plot as the left panel, but the OFF-period PN activity (4 s immediately following odor termination) was now correlated with the overall odor valences.

f) Left: Responses of individual PNs to all twenty-two odors during the ON-period are shown. Each row corresponds to a single PN, and the odorants (columns) were organized from highest valence to lowest (from left to right). PNs were classified as ON responsive (white box) or unresponsive (gray box). Bar plot on the left indicates the number of odorants that activated each PN. PNs are sorted such that those that responded to most odorants are at the top (i.e., least selective). Note that individual PNs whose rasters are shown in panel **c** are identified. **Right:** Similar plot as the left panel, but characterizing OFF-responses across all eighty-nine PNs to all odorants in the panel.

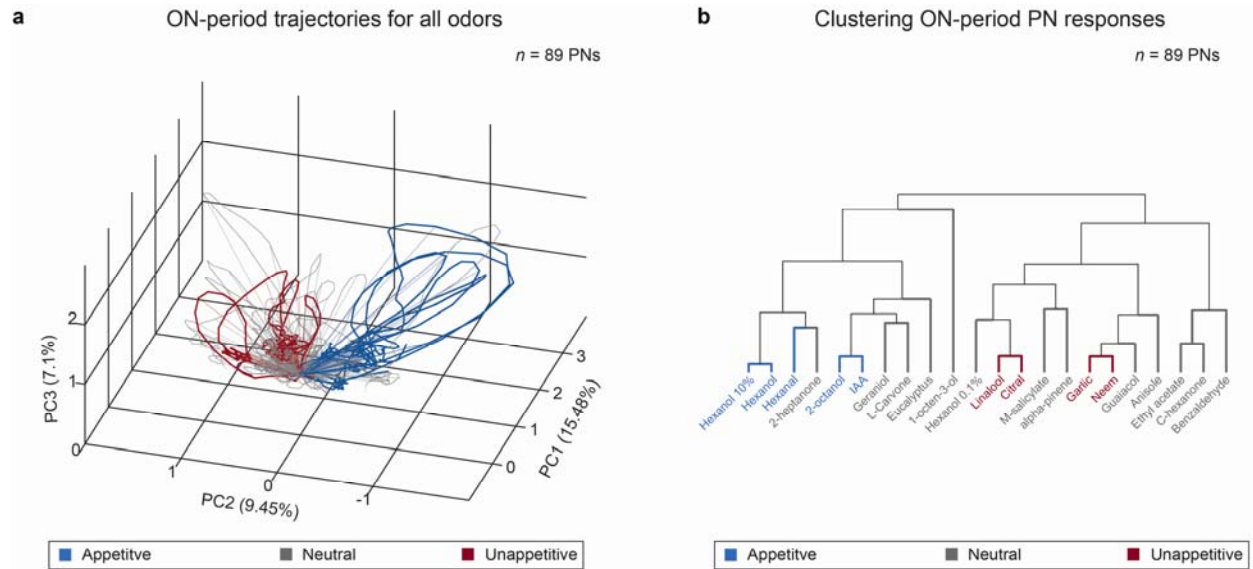


Figure 3

Figure 3: Ensemble PN responses for appetitive and non-appetitive odors

a) Visualization of the ensemble ($n = 89$) PN responses to the odor panel after Principal Component Analysis (PCA) dimensionality reduction are shown (see **Methods**). 4 s of ON-responses for all twenty-two odorants were used for this analysis, and the data were projected on to the first 3 principal components that captured the highest variance ($\sim 30\%$ captured along the three axes shown). Neural response trajectories evoked by innately appetitive odors are colored in blue, neutral odors response trajectories are indicated in gray, and unappetitive odors responses are shown as red trajectories. Note that the ensemble neural response trajectories cluster based on overall appetitive valence.

b) Dendrogram showing the overall hierarchical organization of 89-dimensional PN ON-responses. Odorants are again colored based on the corresponding behavioral preferences (blue indicates appetitive odors, gray indicates neutral odors, red indicates unappetitive odors). Appetitive odors cluster along the left branch, while unappetitive odors cluster on the right branch. It is worth noting that these results are similar to the overall arrangement of responses shown after dimensionality reduction in **panel a**.

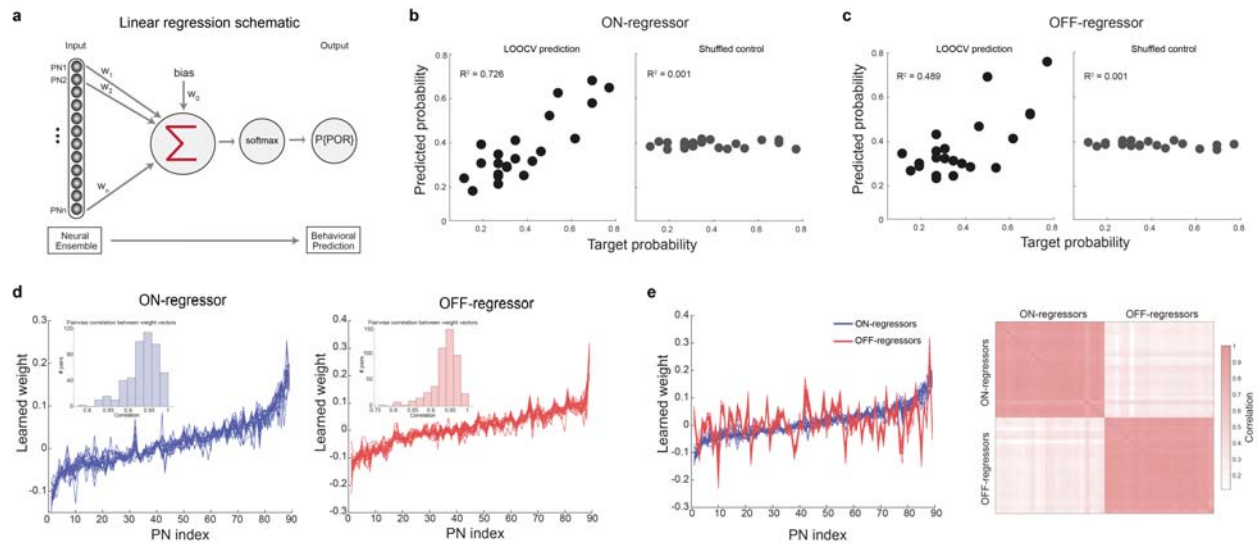


Figure 4

Figure 4: Neural response patterns robustly predict innate behavioral preferences for odors

a) Schematic of the linear regression approach is shown. The input data was the mean PN spiking activity during odor onset for the ON-regressor, or the mean PN responses in a 4 s window post odor termination for the OFF-regressor (i.e., 89 dimensional ON or OFF response vectors). The output to be predicted was the normalized preference score (interpreted as a probability of POR; see **Methods**) for each odorant. The regressors were trained using a gradient descent approach and validated using a leave-one-odorant-out-cross-validation (LOOCV) approach. Therefore, the POR probability for each odorant (that was left out of the training), was based on a regression model learned using the data for the remaining twenty-one odorants. This resulted in twenty-two ON-regressors (one for each odorant), and twenty-two OFF-regressors (again one for each odorant).

b) Left: Predictions from the ON-regressor versus the actual probabilities obtained from the behavioral assay for all odorants in the panel are shown. Overall, the R^2 value between the predicted value and the actual behavioral response was high ($R^2 = 0.726$). **Right:** Similar plot but for the shuffled control is shown. Here, the behavioral POR probabilities were randomized, and a regression model was fit similar to learning the unshuffled case. Note that the predictions are centered around the mean valence of ~ 0.4 ($R^2 = 0.001$).

c) Similar plots as panel **b**, but using models trained on the OFF-period responses are shown. The OFF-regressors ($R^2 = 0.489$) performed poorer than the ON-regression models but were still well above shuffled control performance levels ($R^2 = 0.001$).

d) Left: The ON-period linear regression model was validated by training 22 different models, leaving 1 of the 22 odors out each time for validation. The weights obtained for each PN are shown for all 22 models trained using this leave-one-odor-out cross-validation approach. The weights assigned to eighty-nine PNs were sorted (i.e., lowest to highest) based on the model used to predict POR responses to hexanol. Inset shows the distribution of pairwise correlations between each weight vector obtained for predicting POR for different odorants. **Right:** Similar plot as left panel, but for the twenty-two OFF-regressors are shown.

e) Left: Blue curves indicate weight vectors obtained from the ON-period regressors as shown in panel **d**. Red traces show weights learned by the OFF-period regressors but sorted using the same indices as the

ON-period vectors. As can be seen, the blue and red curves are uncorrelated. **Right:** Correlation analysis quantifying the similarities in weights assigned to PNs by the ON- and the OFF- regressors. As can be expected from panel **d**, weights learned by the PNs are highly correlated within the ON-period and OFF-periods (darker colors along the diagonal blocks). However, as shown in the left panel, the weights assigned to each PN are different between the ON- and OFF-regressors, and hence the off-diagonal blocks have lower correlations (lighter colors).

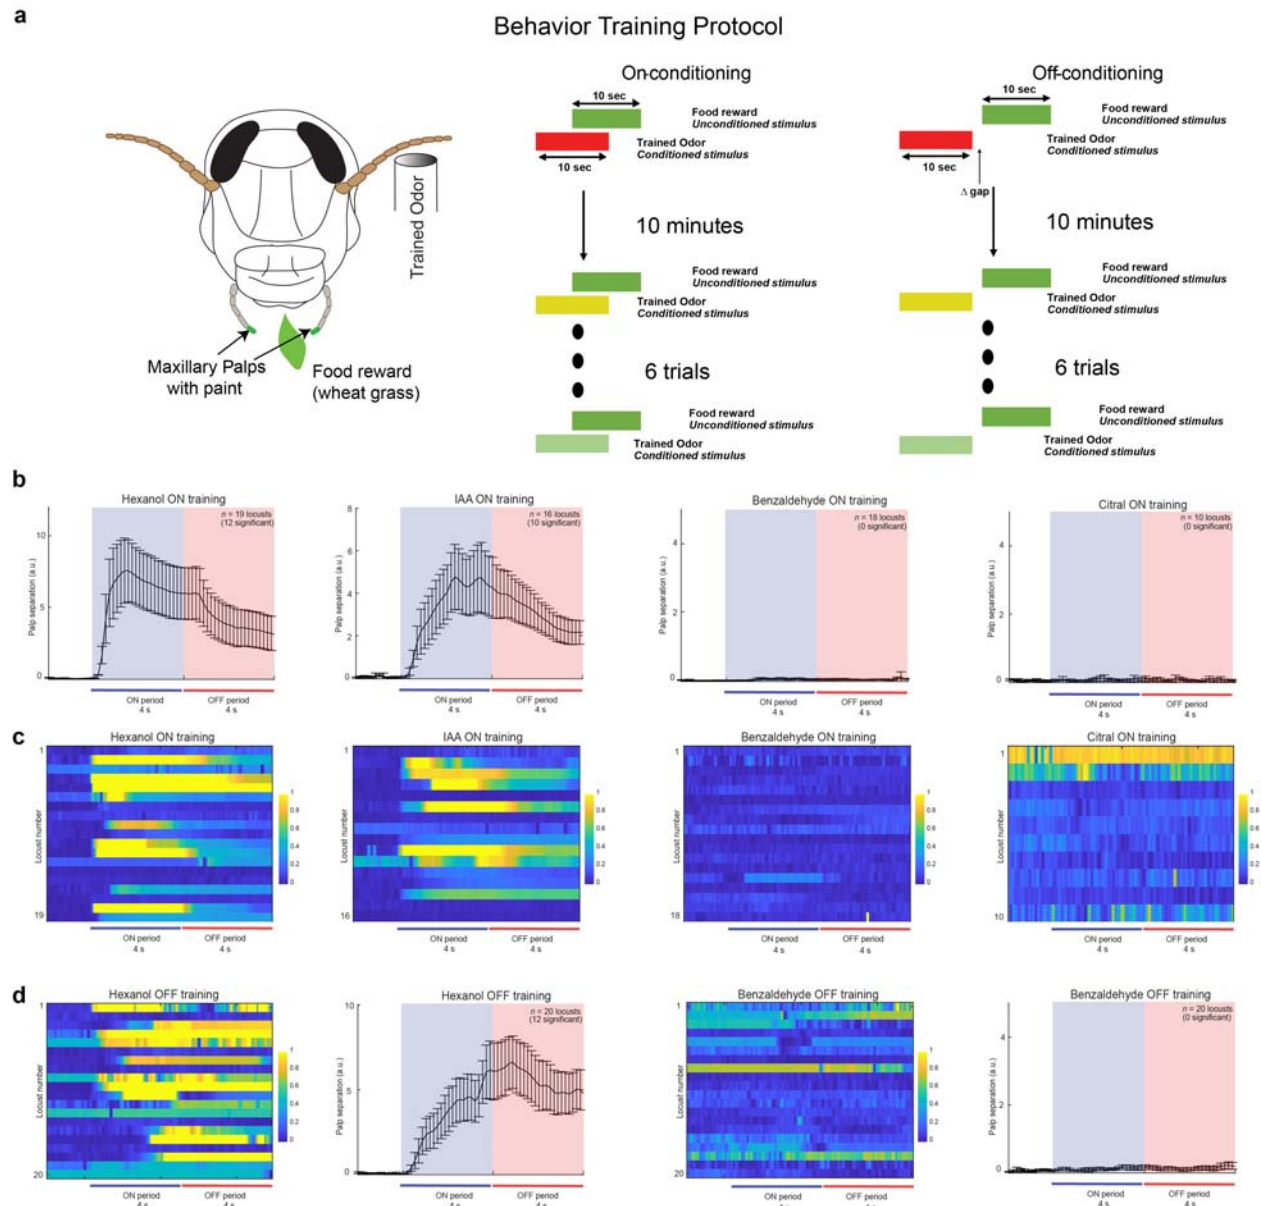


Figure 5

Figure 5: Only innately appetitive odorants can be reinforced using classical conditioning

a) A schematic showing the training protocol followed for both ON- and OFF- classical appetitive conditioning assays (see **Methods** for details). Following the training phase, locusts were then tested for palp-opening responses (PORs) in an unrewarded phase.

b) Results from ON-conditioning using 4 different odors are shown. The mean POR response of locusts during the unrewarded testing phase is shown in each plot. The testing odor was the same as the training odor, as indicated on each plot. Colored bars indicate 4 s of odor presentation and 4 s immediately following odor termination. Error bars indicate s.e.m., and the number of locusts that had significant PORs for each conditioning odorant are shown in parentheses. As can be seen, locusts trained with hexanol and isoamyl acetate were able to produce POR responses in the test phase, while benzaldehyde

and citral training yielded no responses. Note that different sets of locusts were trained/tested for each odorant.

c) POR traces for the four sets of locusts trained with hexanol, isoamyl acetate, benzaldehyde, or citral are shown. The PORs shown were recorded during the testing phase. Each row corresponds to the response observed in one locust. The responses were normalized to range between [0, 1] for each locust (see **Methods**; blue – 0 and yellow -1). Note that for a small fraction of locusts (such as citral, first row) that only had minimal palp movement during the entire trial, the normalization protocol followed produced spurious shades of yellow, but these locusts still did not have a significant response to that odorant.

d) Similar traces as shown in panels **b** and **c** but for OFF-conditioning using hexanol or benzaldehyde are shown. Hexanol-OFF training produced significant PORs in 12/20 locusts, whereas benzaldehyde-OFF training yielded no significant responses. Note that the PORs for hexanol-OFF training were delayed and persisted well into the OFF period (compared to hexanol-ON trained responses shown above).

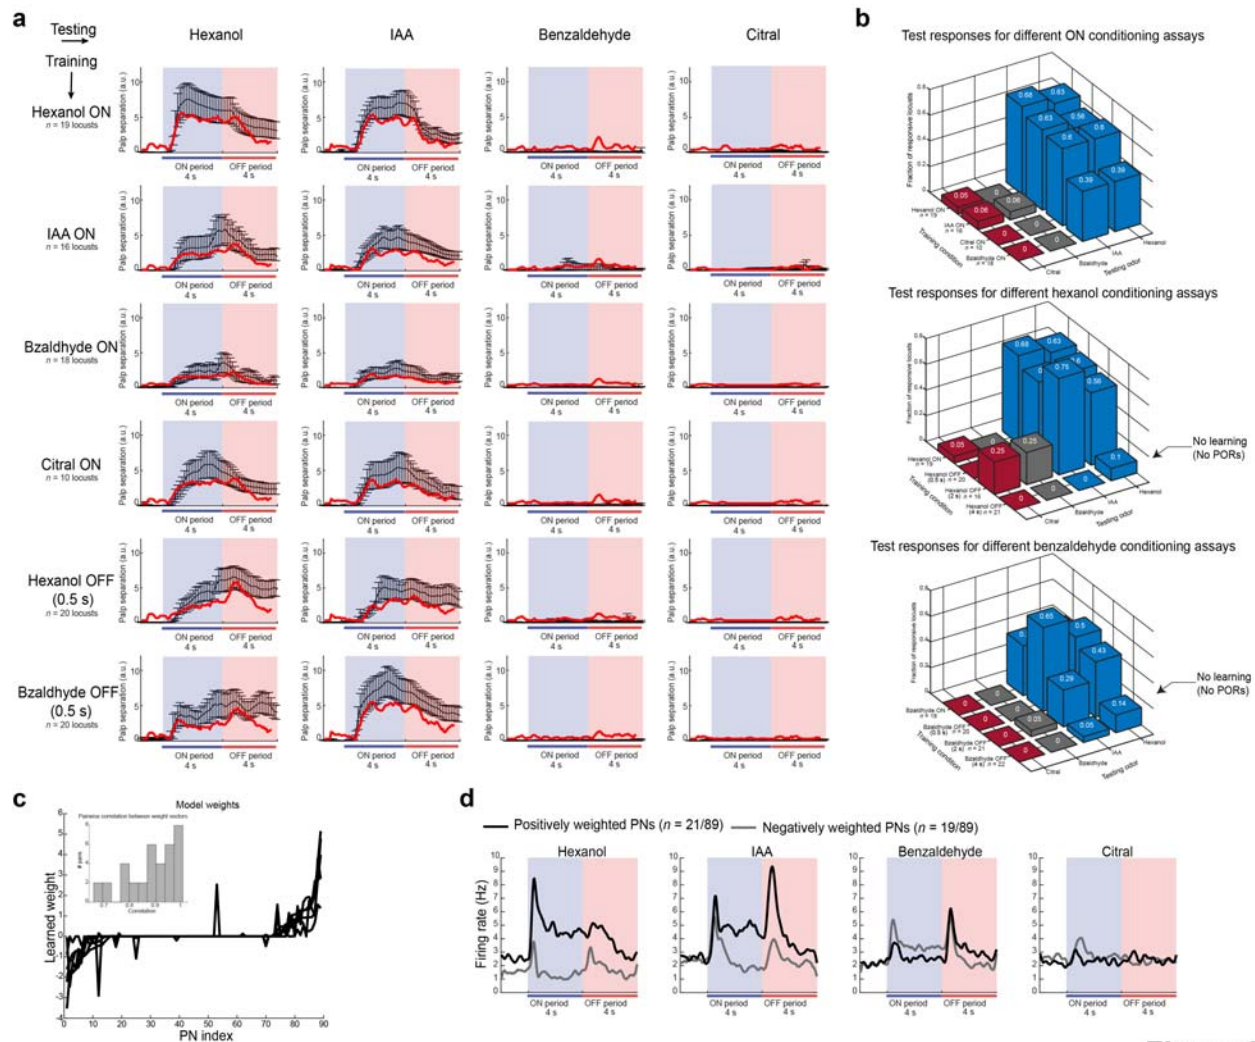


Figure 6

Figure 6: Predictable behavioral response dynamics, cross-learning, and generalization between trained odors

a) Summary of observed and predicted POR responses for six different training conditions are shown: row 1 – ON-trained with hexanol, row 2 – ON-trained with isoamyl acetate, row 3 – ON-trained with benzaldehyde, row 4 – ON-trained with citral, row 5 – OFF-trained (0.5 s gap) for hexanol and row 6 – OFF-trained with benzaldehyde. The number of locusts tested in each training paradigm is shown on the left. Responses of the trained locusts were examined for all four odorants during the unrewarded testing phase. The mean PORs to each odorant are shown in black and error bars indicate s.e.m. Colored bars indicate odor ON and OFF periods. Red traces on each plot show PORs produced by linear regression model that used ensemble PN activity for the four different odorants as inputs (see **Methods**).

b) The fraction of locusts that had statistically significant PORs to each of the four odorants used in the testing phase is shown. Top panel summarizes the results from ON-training using all 4 odors. Middle panel summarizes results from training with hexanol during the ON-period and during the OFF-period with 0.5 s, 2 s, and 4 s latencies. Bottom panel shows similar ON- and OFF- conditioning results but using benzaldehyde. The number of locusts tested in each case is indicated along the axis labels. Blue – appetitive odorants, gray – neutral and red – unappetitive odorants.

c) Weights learned by linear regression models used to predict mean POR responses using ensemble PN activity are shown. The weights were sorted based on the values assigned to each PN for the hexanol-ON training model. Inset shows the distribution of pairwise correlations between different pairs of weight vectors (i.e., across different regression models).

d) Summed spiking activities of all PNS that were assigned positive (black) or negative (gray) weights are shown. In total twenty-one PNS were assigned positive weights, nineteen PNS received negative weights < 0 , and the remaining 49 PNS were assigned a weight of 0.

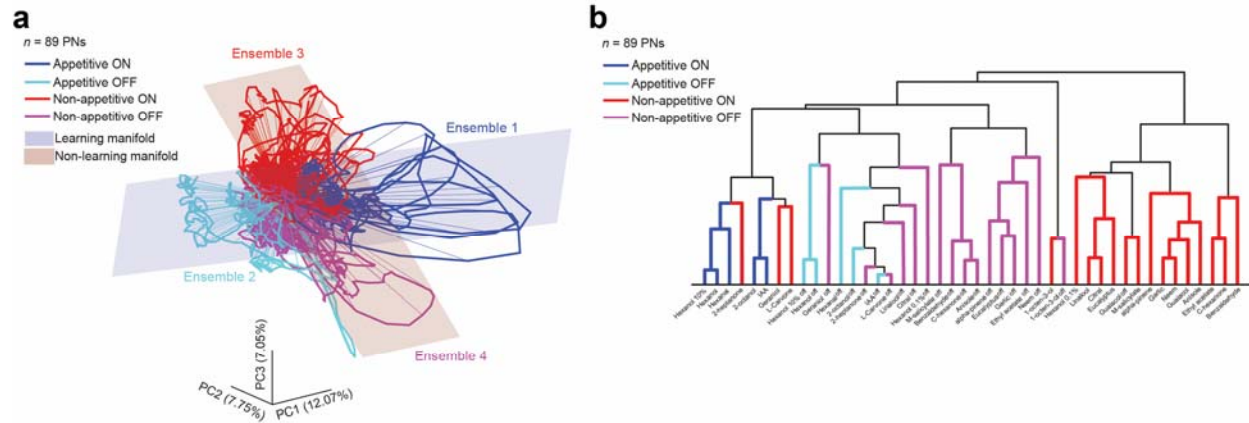


Figure 7

Figure 7: Neural manifolds can explain innate and acquired behaviors

a) PCA trajectories showing ensemble neural responses during both the ON- and the OFF- periods for all 22 odors are shown along the top 3 principal components ($n = 89$ PNs; see **Methods**). The trajectories were colored as follows: blue – appetitive odorants ON responses, cyan – appetitive odorants OFF responses, red – non-appetitive odorants ON responses, and magenta – non-appetitive odorants OFF responses. Variances in odor-evoked responses of appetitive odorants were not uniformly distributed but confined a subspace and are schematically shown as using a linear plane (plane colored in blue that encompasses appetitive ON and appetitive OFF neural ensembles). Similarly, non-appetitive odorants ensemble responses are confined to a distinct neural manifold schematically shown in red.

b) Dendrogram showing the categorization of odor-evoked ON and OFF responses of all twenty-two odors in the panel are shown. A correlation distance metric was used to assess the similarity between 89-dimensional PN response vectors. Coloring convention similar to panel **a**. Note that the appetitive and non-appetitive odorants form supra-clusters, each containing ON and OFF responses sub-clusters.