Learning improves decoding of odor identity with phase-referenced oscillations in the olfactory bulb

Justin Losacco[1,2,*], Daniel Ramirez-Gordillo[2,*], Jesse Gilmer[1,3] and Diego Restrepo[1,2,#]

[1] Neuroscience Graduate Program

[2] Department of Cell and Developmental Biology

[3] Department of Physiology and Biophysics

University of Colorado Anschutz Medical Campus, Aurora, CO 80045, USA

[*] J.L and D.R.G. contributed equally to this work

[#] Corresponding author: Diego Restrepo, diego.restrepo@cuanschutz.edu

SUMMARY

Local field potential oscillations reflect temporally coordinated neuronal ensembles—coupling distant brain regions, gating processing windows, and providing a reference for spike timing-based codes. In phase amplitude coupling (PAC), the amplitude of the envelope of a faster oscillation is larger within a phase window of a slower carrier wave. Here, we characterized PAC, and the related theta phase-referenced high gamma and beta power (PRP), in the olfactory bulb of mice learning to discriminate odorants. PAC changes throughout learning, and odorant-elicited changes in PRP increase for rewarded and decrease for unrewarded odorants. Odorant identity can be decoded from peak PRP in animals proficient in odorant discrimination, but not in naïve mice. As the animal learns to discriminate the odorants the dimensionality of PRP decreases. Therefore, modulation of phase-referenced chunking of information in the course of learning plays a role in early sensory processing in the high input dimensionality sense of olfaction.

SIGNIFICANCE

Early processing of olfactory information takes place in circuits undergoing slow frequency theta oscillations generated by the interplay of olfactory input modulated by sniffing and centrifugal feedback from downstream brain areas. Studies in the hippocampus and cortex suggest that different information “chunks” are conveyed at different phases of the theta oscillation. Here we show that in the olfactory bulb, the first processing station in the olfactory system, the amplitude of high frequency gamma oscillations encodes for information on odorant identity more accurately when it is observed at the peak phase of the theta oscillation. Furthermore, encoding of information by the theta phase-referenced gamma oscillations becomes more accurate as the animal learns to differentiate two odorants.

INTRODUCTION

Animals must modulate early sensory processing to optimize navigation of their environment (Gire et al., 2013a; Pakan et al., 2018). This experience-dependent shaping involves an interplay between sensory input, behavioral state, arousal and motor activity. Neuronal activity, temporally organized by local field potential (LFP) oscillations, carries distinct information at different phases of each cycle (Jensen, 2001; Kepecs et al., 2006; Lisman and Jensen, 2013). For example, in the hippocampus different phases of the theta LFP (2-12 Hz) preferentially encode present and future spatial locations, illustrating temporal “chunking” of information (Siegle and Wilson, 2014; Amemiya and Redish, 2018). Place cell preferred fields are encoded temporally through phase precession in relation to hippocampal theta oscillations (O'Keefe and Recce, 1993; Skaggs et al., 1996; Buzsáki and Draguhn, 2004; Lisman, 2005). Additionally, LFP oscillations in different frequency bandwidths can be related through cross frequency coupling (CFC), which can further enhance information carrying capacity (Engel et al., 1999; Fries, 2005; Canolty and Knight, 2010; Florin and Baillet, 2015).

One example of CFC is PAC, wherein the amplitude of a faster oscillation (e.g. gamma, 35-95 Hz) is related to the phase of a slower oscillation (e.g. theta, 2-12 Hz) (Soltesz and Deschenes, 1993; Chrobak and Buzsaki, 1998). PAC occurs during task engagement in rodent hippocampus (Soltesz and Deschenes, 1993; Bragin et al., 1995; Lisman and Idiart, 1995), rodent orbitofrontal cortex (van Wingerden et al., 2014), macaque auditory cortex (Marton et al., 2019), human hippocampus (Lega et al., 2016)

and human visual cortex (Daume et al., 2017; Seymour et al., 2017). Interestingly, Van Wingerden et al. (2014) described how theta-gamma PAC strengthens in rat orbitofrontal cortex with learning for trials in which the correct decision was made in the olfactory go no-go task.

The study of PAC in the earliest sensory processing brain areas is virtually non-existent. The olfactory bulb (OB) is an ideal location to study whether PAC is relevant to sensory processing; its intrinsic oscillatory activity is linked to odorant signal processing (Doucette et al., 2011; Smear et al., 2011; Gschwend et al., 2012; Li et al., 2015). The rhythmic activation of olfactory sensory neurons (OSNs) through breathing generates a theta-frequency (2-12 Hz) respiratory-related LFP in the target region of the OSNs, the OB, that is coherent with oscillations in downstream processing areas such as piriform cortex (PC) and hippocampus (Gourevitch et al., 2010; Nguyen Chi et al., 2016; Tort et al., 2018; Heck et al., 2019). Interestingly, PAC is present in human PC, and the respiratory phase for presentation of stimuli modulates the ability to retrieve episodic memory (Zelano et al., 2016). Downstream neural networks receiving olfactory information from the OB could be tuned to respond to sensory input at certain phases of this theta LFP (Jensen, 2001; Buzsaki, 2010) and could send back modulatory signals during other phases (Kepecs et al., 2006; Kay, 2015).

Decades of work have characterized LFP activity in the OB in various states of wakefulness (Adrian, 1950; Li et al., 2012; Chery et al., 2014) and olfactory task engagement (Gourevitch et al., 2010; Martin and Ravel, 2014). Surprisingly, this LFP

signal changes in response to prior experience (Gire et al., 2013b; Courtiol and Wilson, 2017). Work in the go no-go operant conditioning task where the animal learns to respond to the reinforced odorant for a water reward has revealed broadband LFP power changes for the rewarded versus the unrewarded odorant with discrimination learning (Stopfer and Laurent, 1999; Martin et al., 2004; Beshel et al., 2007; Ramirez-Gordillo et al., 2018). Additionally, changes in the balance of sensory and corticofugal inputs to OB interneurons can alter timing of beta and gamma oscillatory states, suggesting that information transfer from the OB is influenced by top-down signaling (David et al., 2015). Also, these investigators posited that slower theta oscillations gate temporal windows for sensory and centrifugal inputs to ascend and descend, respectively, suggesting an important role for PAC in gating information transfer. Since PAC exists in the OB in awake, idle rats (Rojas-Libano et al., 2014), we wondered whether PAC changes with learning, reflecting unsupervised Bayesian learning or supervised learning through top-down modulation sensory information (Hiratani and Latham, 2019), and whether encoding this information by the power of high frequency oscillations differs as a function of the phase of the theta LFP.

Here, thirsty mice implanted with tetrodes in the OB learned to associate an odorant with water reward in the go no-go task, discriminating between two monomolecular odorants (e.g. between isoamyl acetate vs acetophenone, IAAP) or between a monomolecular odorant and an odorant mixture (between ethyl acetate and a mixture of ethyl acetate and propyl acetate, EAPA). LFP was analyzed over task learning for changes in PAC and for the ability of theta phase-referenced high frequency

oscillatory power—termed phase referenced power (PRP) hereafter—to encode for the perceived identity of the odorant.

We demonstrate that the strength of PAC, quantified as the modulation index (Tort et al., 2010), changes for the rewarded odorant (S+) as the animals learn. Interestingly, the variance of the angle of theta with the strongest gamma oscillation (peak angle) increases for the unrewarded odorant (S-) over the course of learning. These changes in PAC raised the question whether a downstream observer focused on high frequency oscillations at different theta phases would find a difference in stimulus prediction with learning. We found that the ability to decode stimulus identity by PRP improves with task proficiency and the dimensionality of PRP decreased as the animal learned to differentiate the odorants.

RESULTS

We report data obtained in two different sets of experiments that we name exp1 and exp2. In exp 1 mice were implanted with optetrodes in the OB and in exp 2 animals were implanted with tetrodes 0.5 mm deeper. Mice learned to discriminate between rewarded (S+) and unrewarded odorants (S-) presented in pseudorandomized order in the go no-go olfactory task (Figure 1A-C). The dataset is comprised of 172 recording sessions in 30 mice. The odorant pairs tested were different volatile compounds (or mixtures) whose nomenclature addresses the odorant names and the experimental set (e.g. IAAPexp1, see Table S1 for the nomenclature).

In the go-no go task mice must decide to either lick a waterspout at least once during each 0.5-second bin in the 2-second response area (green blocks in Figure 1B) to obtain a water reward for the rewarded odorant or refrain from licking for the unrewarded odorant (Figure 1C). We did not punish for licking for the unrewarded odorant, and mice refrain because of the effort it takes to lick. Behavioral performance was termed naïve or proficient when their performance estimated in a 20-trial window was ≤ 65% for naïve and ≥ 80% for proficient. After three 20-trial blocks of proficient performance, the session was ended; valence was reversed the next day (not shown) or another odorant was tested.

Phase amplitude coupling analysis of the LFP recorded in the go no-go behavioral task

Figures 1Di and ii show an example of the extracellular LFP sampled at 20 kHz and filtered between 1-750 Hz. The raw signal (Figure 1Di) was filtered with a 20th order Butterworth filter into different LFP frequency bands (high theta, 6-14 Hz, adapted from (Nguyen Chi et al., 2016) and subsequently referred to as theta; beta, 15-30 Hz; or high gamma, 65-95 Hz). Figures 1Diii and 1Dv show that the filtered theta and high gamma LFP appear to change amplitude in a coordinated manner. We used PAC analysis (Tort et al., 2010) to evaluate the degree of coupling of the amplitude of the envelope of the beta or high gamma LFP on the phase of the theta LFP. Figure 1Div shows the theta phase and the red line in Figure 1Dv shows the envelope for the amplitude of the high gamma LFP, both calculated with the Hilbert transform as detailed in Tort et al. (2010). Figure 1E shows the phase amplitude plot of the normalized distribution of high gamma amplitudes in different theta phase bins. For this example there are two maxima, and the vertical red line corresponds to the phase of theta at which the gamma amplitude is highest (at ~95 degrees); this is called the “peak phase”. Similarly, the phase with the smallest normalized amplitude is called the “trough phase” (vertical blue line at ~5 degrees). The strength of PAC was estimated as the modulation index (MI), a measure computed as the Kullback–Leibler (KL) distance to the uniform phase distribution normalized to range from 0 (uniform distribution) to 1 (sharp gamma bursts at a specific theta phase) (Tort et al., 2010). The MI for the distribution in Figure 1E is 0.018. Figure S1 shows the LFP power spectrogram for simulated data with low gamma oscillations that are either uniformly distributed in the theta phase (MI=0.0004, Figure S1A) or with

gamma bursts at 180 degrees of the theta oscillation (MI=0.01, Figure S1B). Recordings in hippocampus and cortex yield MI values in the range of 0.005-0.03.

Peak angle variance decreases for S+ and/or increases for S- when the animals learn to differentiate odorants in the go-no go task

Figure 2A-F illustrates PAC for all S+ and S- odorant trials in one session for a mouse proficient in differentiating the odorants. Behavioral performance is shown in Figure 2B, the phase amplitude plots for all the trials are shown in pseudocolor in Figure 2A, and the corresponding normalized phase amplitude plots are shown in Figure 2C, while the mean average theta LFP is shown for one cycle in Figure 2D. There is clear (biphasic) theta phase coupling for high gamma amplitude with maxima at ~120° and 270° for S+ (Figure 2A, left and red line in 2C). This strong PAC is reflected by an MI with a mean of 0.02 with small trial to trial variance (Figure 2E, red circles), and small variation in peak phase (Figure 2F, red polar histogram). In contrast, PAC displays variable strength and phase for S- in this session (Figure 2A, right, blue line in 2E and blue polar histogram in 2F). For S- the average theta amplitude is smaller (Figure 2E, blue line) and the MI displays large variance ranging from 0 to 0.06 (mean MI 0.014, Figure 2E, blue circles). The MI and peak angle variance differ significantly between S+ and S- (MI: ranksum p value < 0.001, peak angle variance: ranksum p value < 0.05, n = 50 S- trials, 49 S+ trials).

Browsing through PAC data for different sessions, plotted as in the example shown in Figures 2A-F, raised the question whether PAC changes over the course of

learning in the go-no go olfactory discrimination task—as reported by mean MI and peak angle variance. We noticed that in sessions where the mouse learned to differentiate the odorants the phase of the peak angle became highly variable when the animal became proficient. An example is shown in Figure S2 where a mouse started the session performing close to 50% correct, and became proficient during the last portion of the session. The phase of the peak angle is highly variable for the last 20 S- trials. We proceeded to estimate mean MI and peak angle variance for all mice. The mean MI and peak angle variance was calculated as the mean in all trials within either the naïve or proficient periods for S+ and S-. Figure 2G illustrates cumulative histograms for the mean MI (i,iii) and peak angle variance (ii,iv) for theta/beta (top) and theta/high gamma PAC (bottom) for the APEBexp1 odorant for S+ and S- odorants for trial ranges when the mice were naïve or proficient. The cumulative histograms for mean MI measured per electrode (Figure 2Gi,iii) show small yet statistically significant changes in PAC strength for both the theta/beta and theta/high gamma as the animal learns to discriminate between odorants (generalized linear regression model, GLM analysis, indicates that the differences in mean MI in terms of proficiency and the event type are significant $p<0.001$, 896 degrees of freedom, d.f., for theta/beta and theta/high gamma, 16 electrodes per mouse, 14 mice). In addition, the circles in Figure 2Gi,iii show the per mouse mean MI. A GLM analysis did not yield significant changes for the per mouse average MI as a function of either event type or proficiency for theta/beta PAC ($p > 0.05$, 52 d.f., 14 mice). Per mouse average MI was significant for both event type and proficiency for theta/high gamma PAC ($p < 0.05$, 52 d.f., 14 mice). Finally, for the APEBexp1 odorant pair we found a substantial decrease in peak angle variance for S+

and an increase in peak angle variance for S- for both the theta/beta and theta/high gamma PAC (Figure 2Gii,iv, GLM analysis for changes in peak angle variance for both proficiency and event type yield $p < 0.001$, 892 d.f. for both theta/beta and theta/high gamma, 16 electrodes per mouse, 14 mice). The changes were also statistically significant when a GLM was computed for the average peak angle variance per mouse (circles in Figures 2Gii,iv, GLM p value < 0.01 for theta/beta and $p < 0.001$ for theta/high gamma for both event type and proficiency with 52 d.f.,14 mice).

Similar differences between S+ and S- and changes when the animal became proficient were found for peak angle variance for the other odorant pairs. Figures 2Hi-iv show the summary of changes with learning for S+ and S- odorants for MI (i, iii) and the peak angle variance (ii,iv) for theta/beta (i,ii) and theta/high gamma (iii,iv) for per electrode PAC for all odorant pairs. As mentioned above recordings were performed in two sets of experiments (Table S1). Theta/beta and theta/high gamma PAC strength, quantified as the MI calculated for LFPs recorded in each electrode, was stronger for recordings in exp2 compared to exp1, and there was a statistically significant increase in the MI for S+ with learning (naïve vs. proficient) (GLM analysis p value < 0.001, for event, proficiency and location, with the exception that $p < 0.01$ for proficiency for theta/high gamma, 3255 d.f., number of mice is in Table S1). Interestingly, there were decreases in peak angle variance for S+, and/or increases in peak angle variance for S- for theta/high gamma PAC when the animals became proficient (GLM analysis p values < 0.001 for event type, proficiency and location for both theta/beta and theta/high gamma, 3255 d.f., number of mice is in Table S1). For a downstream observer focused

on the LFP power at the peak of PAC for the rewarded odorant these changes in peak angle variance may reflect a mechanism by which the unrewarded stimulus is devalued while the rewarded odorant is valued. To our knowledge, this is the first demonstration of dynamic cross frequency interactions changing with learning for early processing in sensory systems.

Peak phase referenced power increases for S+ and decreases for S- over learning

The existence and evolution of PAC over learning alone does not demonstrate its potential utility to the animal in its decision-making process. In order to assess whether high frequency LFP power at the peak of theta oscillations changes as the animal learns we performed wavelet analysis to extract the power of high gamma (and beta) LFP referenced to specific PAC phases of theta—peak and trough (as defined in Figure 1E). We term this analysis phase-referenced (high-frequency) power, or PRP. In Figures 3A-C, we show how this PRP analysis was performed. Figure 3A shows Morlet wavelet LFP power analysis (Chery et al., 2014; David et al., 2015) for a single S+ trial characterized by strong, persistent theta power during and after odor delivery (related to respiration and olfactory investigation) accompanied by higher frequency power bursts whose periodicity reflects PAC. Figure 3B shows the high theta (6-14 Hz)-filtered waveform (Figure 3Bi) and the Hilbert-extracted phase (Figure 3Bii, (Tort et al., 2010)) co-registered to the one-second epoch extracted from the wavelet LFP power plot shown in Figure 3Aii. We proceeded to extract the value of the power (in decibels) referenced to the peak or the trough of PAC measured in S+ trials (peak and trough are defined in Figure 1E). The red line in Figure 3Ci shows a clear odor-induced increase in

the peak-referenced wavelet power for high gamma/theta PAC in this single trial example (Figure 3Ci, red line). In contrast, there was little change in power referenced to the trough of PAC (Figure 3Ci, blue line). Figure 3Cii shows the time course for the average peak (red) and trough (blue) PRP calculated for all trials in this go-no go session for a mouse that was performing proficiently in differentiation of the two odorants.

We performed the PRP analysis for all sessions for naïve and proficient periods. Figure 3D shows cumulative histograms for the peak and trough PRP for all mice for the odor period for all go-no go sessions with the IAAPexp2 odorant pair for theta/high gamma PAC. Over the course of learning, peak-related high gamma power increases for S+ and decreases for S- (Figure 3D top). In contrast, the changes elicited by learning in trough PRP are smaller (Figure 3D bottom). The differences were statistically significant between peak and trough, S+ and S- and naïve and proficient (GLM p values < 0.001, 1016 d.f.). The summary of PRP for peak and trough for theta/high gamma PAC for the odor period are shown for all odorant pairs in Figure 3E (peak PRP) and 3F (trough PRP). For all odorant pairs peak PRP increases for S+ over learning and decreases for S-. Trough PRP is weaker yet follows the same trend as the animal learns to discriminate the odorant as peak PRP (GLM p value < 0.001, 6386 d.f. for S+ vs. S-, peak vs. trough, naïve vs. proficient). In addition, the learning-induced changes in PRP are larger for exp2 (GLM p value < 0.001 for location, 6386 d.f.). Similar changes were found for PRP for theta/beta PAC, but there was no difference between peak and trough (not shown).

Linear discriminant analysis classifies stimuli using PRP

The finding that the LFP power referenced to the peak of the theta oscillation increases for S+ and decreases for S- when the animal learns to differentiate the odors (naïve vs. proficient) raises the question whether a downstream reader evaluating peak power could decode which stimulus is presented. To assess decoding of olfactory stimuli using PRP, a linear discriminant analysis (LDA) was used to set a decision boundary hyper-plane between binary stimulus classes (S+ versus S-) (Vizcay et al., 2015) using PRP data referenced to the peak (or the trough) of theta/beta or theta/high gamma PAC. LDA was trained with PRP from each electrode for each mouse (16 electrodes per mouse) for all trials except one (the training dataset) and then the missing trial (test data) was classified as S+ or S- using the PRP data from that trial. This was performed separately for trials where the mouse was naïve or proficient to odorant discrimination, for each time point during the trial and was repeated for all trials. As a control we shuffled the identity of trials in the training set. In addition, we performed a complementary principal component analysis (PCA) of the PRP to visualize the odor-induced divergence in the time course of the first principal component (PC1).

Figures 4 A and B show PC1 for the PCA (Figure 4A) and the performance of the LDA (Figure 4B) calculated with the PRP for 16 electrodes for one example odorant pair (APEBexp1). In the naïve state PC1 diverged after water reinforcement (Figure 4A, left) and the LDA could not classify stimulus identity using peak or trough PRP for either theta/beta or theta/high gamma (Figure 4B, left). However, when the animal became

proficient PC1 diverged between S+ and S- shortly after odorant addition (Figure 4A) and the precision of the LDA classifier increased from chance to values above 70%, significantly higher than the 95% confidence interval of the LDA performed with shuffled trials during the odor delivery epoch for both theta/beta and theta/high gamma (Figure 4B; GLM analysis for learning p value < 0.001, 104 d.f.). In theta/beta, both peak and trough PRP performed significantly better than shuffled data at classifying stimulus class via LDA (GLM p value < 0.001, 104 d.f.). However, in theta/high gamma (Figure 4B, bottom) LDA discriminated stimulus class significantly more effectively using peak PRP than trough PRP (GLM p value < 0.001, 104 d.f.). The LDA performance difference between peak and trough PRP was evaluated by calculating the area under the performance curve (AUC) during the odor application window (0.5-2.5 sec after odorant valve opening). AUC was normalized and it ranged from zero (random decoding at 50%, Figure S3 Ai) to 1 (100% decoding during the odor period, Figure S3 Aii). AUC is plotted for all odorant pairs for peak and trough PRP in Figure 4C for theta/beta and theta/high gamma for naïve and proficient conditions. In the proficient state, AUC for peak PRP LDA for high gamma was significantly higher than AUC for trough PRP LDA (Figure 4C, bottom, GLM for AUC peak vs. trough p value < 0.001, 104 d.f.). AUC for beta PRP LDA did not differ between peak and trough (Figure 4C, top). Finally, Figure 4D shows that the AUC was larger for exp2 (GLM p value < 0.001 for location, 6386 d.f.). To our knowledge, this is the first demonstration of high gamma phase-referenced LFP power developing differential information carrying capacity in early sensory systems with learning.

Decision-making takes place at the same time for peak PRP via LDA and licks

After learning that LDA using peak PRP can decode stimulus identity, we wondered how early in the time course of the trial there is a statistically significant difference between behavioral responses (licks at the waterspout) compared to correct classification of the stimulus by LDA. We estimated the p value for decision-making by computing the ranksum test for the difference between licks (scored 0 for no lick vs. 1 for lick) and for the LDA (scored 1 for correct prediction vs. 0 for incorrect prediction). Figures 5A-B show an example of licks by a mouse in one session (for the EAPAexp2 odorant pair). Figure 5A shows an example of the licks detected at the waterspout for 25 S+ trials when the animal was performing proficiently while Figure 5B shows the lick traces for 25 S- trials in the same session.

Next, we analyzed lick vs. PRP LDA decision making times after pooling data for all mice for each odor pair. An example of the time course of the p values for lick and PRP LDA performing the go-no go task for EAPAexp2 is displayed in Figure 5C for theta/high gamma PAC. At the left, the time-course for performance of decoding stimulus identity for the peak PRP LDA (top) appears to be faster and more accurate compared to the time course of LDA performance evaluated with trough PRP (bottom). Figure 5C (right) shows the p value for LDA performance using the ranksum test. Before odorant presentation the p value fluctuates and sometimes falls below 0.05 transiently, while after odor presentation the p values for peak PRP and lick fall below 0.05 sharply and continuously, while trough PRP decoding performance is significantly worse. We defined decision making time as the time point when the p value falls and stays below

0.05 after odor application. For the pooled data for all mice performing EAPAexp2, peak PRP and lick p values fall below 0.05 earlier than for trough PRP. Both peak PRP and lick perform similarly in LDA decoding of stimulus identity in this example. Figure S3 E and F show that for both lick (E) and PRP peak (F) the decision time is faster when data are pooled for all mice. Lick decision times (Figure S3E) were significantly faster with pooled data (GLM p value < 0.01, 21 d.f.) as were peak PRP values (Figure S3F, GLM p value < 0.01, 21 d.f.). LDA performance as measured with AUC was significantly better with pooled vs. per-mouse data as well (Figure S3B, GLM p value < 0.001, 21 d.f.). Because of this we did the comparison of decision times using pooled animal data.

Analysis of LDA decoding performance via peak PRP and lick for all odorant pairs for pooled mouse data is shown in Figure 5D. Decision times were significantly reduced with task learning in theta/beta (GLM p values < 0.001, 27 d.f.) and in theta/high gamma (GLM p values < 0.001, 27 d.f.). However, lick vs. PRP LDA (including peak and trough) did not differ between naïve and proficient states (GLM p values > 0.05, 20 d.f.).

Dimensionality of PRP LFP space decreases when the animal learns to discriminate the odorants

Figure 6A shows the change in the scatter pattern of the first two principal components (PCs) of the theta/high gamma peak PRP calculated for pooled data from all mice as the animals learn to differentiate the two odorants for the EAPAexp1 odorant pair. Figure 6Ai shows twenty trials of the 2D PCA spatial pattern before and during

odorant application (S+, red and S-, blue) when the mice are naïve, while Figure 6Aii shows that the PCA patterns become clearly separable after the animals become proficient. There appears to be a change in the geometrical arrangement of the points in principal component space suggesting that there is a change in dimensionality as the animal learns to differentiate the odorants (Litwin-Kumar et al., 2017).

In order to provide a more nuanced analysis of whether the spatial arrangement of the PRP LFP data measured by the 16 electrodes changes as the animals learn to discriminate between odorants we calculated a quantitative measure that characterizes the dimensionality in the PRP LFP space (see methods). The 16 electrodes are closely spaced (12.5 μm diameter electrodes in the same tetrode are directly abutted against each other and tetrodes are separated by 100-200 μm in both preparations: exp1 with optetrodes and exp2 with tetrodes). Because of this close proximity the extracellular field potentials would not be expected to differ greatly (Gold et al., 2006) and therefore the dimensionality of the PRP LFP space would be expected to be close to one. Interestingly, when the dimensionality was calculated on a per mouse basis where the maximum dimensionality (M) is 16 we found that the dimensionality for electrodes in exp2 was indeed close to one, but it was significantly higher (2-4) for theta/high gamma PRP LFP in exp1, with the highest dimensionality for PRP LFP trough (Figure S4Aii, GLM p value < 0.001 for peak vs. trough and location, 379 d.f.). We found similar differences in dimensionality between exp1 and exp2 for theta/beta PRP LFP, but in this case there was not a difference between peak and trough dimensionality (Figure S4Ai, GLM p value < 0.001 for location and >0.05 for peak vs. trough, 379 d.f.). Furthermore,

the dimensionality decreased after addition of the odorant (Figure S4A, GLM p value < 0.001 for pre-odor vs odor, for theta/beta, GLM p value < 0.05 for pre-odor vs odor, for theta/high gamma, 379 d.f.). However, statistical analysis of dimensionality when assayed on a per-mouse basis (Figure S4B) does not yield a significant change in dimensionality of PRP LFP as the mouse learned to differentiate the odorant for either theta/beta or theta/high gamma PRP LFP (GLM p value > 0.05 for proficiency, 379 d.f.).

We proceeded to analyze the dimensionality of PRP LFP space when the data were pooled for all mice. Figure 6B shows the time course for dimensionality of the PRP LFP space for two odorant pairs (Figure 6Bi is APEBexp1 and Figure 6Bii is EAPAexp2). The red trace is the dimensionality for the peak PRP LFP while the blue trace is the dimensionality for the trough PRP LFP. Because the measure of dimensionality shown in Figure 6B was calculated for the pooled mouse data the maximum dimensionality was the product of the number of animals x 16 electrodes (the number of animals per odorant pair is shown in Table S1). Thus, for the pooled animal measure of PRP LFPs the dimensionality could range from 1 to M where M is 224 for Figure 6Bi (14 mice x 16 electrodes) and 144 for Figure 6Bii (9 mice x 16 electrodes). In these examples, for the proficient mice there is a clear decrease in dimensionality for peak PRP after the odorant is applied (as indicated by the horizontal black bar), and the dimensionality in the presence of the odorant appears to change when the animal becomes proficient. Figure 6C summarizes theta/beta (Figure 6Ci) and theta/high gamma (Figure 6Cii) peak/trough pooled mouse PRP LFP dimensionality estimates for all odorant pairs. Both exp1 and exp2 were included in these plots where we normalized

the dimensionality to the dimensionality before addition of the odorant. There is a clear decrease in dimensionality after odorant addition for proficient mice for theta/beta PRP (both peak and trough), and for high gamma/theta peak PRP. For theta/beta PRP dimensionality the GLM yielded significant changes for time course and for the interaction of time course x learning ($p < 0.001$, 1696 d.f.), but there were no differences for peak vs. trough ($p > 0.05$). For theta/high gamma the GLM yielded significant changes for peak vs trough ($p<0.001$), learning ($p < 0.01$) and for the interaction of peak vs. trough x learning x time course ($p < 0.001$, 1696 d.f.). Therefore, for the pooled animal PRP LFP space there is a decrease in dimensionality as the animal learns to discriminate between odorants. This is evidence for a learned tightening of dimensionality for the odor representation manifold.

DISCUSSION

We asked whether information carried by high frequency oscillations in brain regions involved in early sensory processing changes when observed at different phases of theta frequency oscillations. Specifically, we asked whether information on odorant identity conveyed by LFP power of high gamma and beta frequency oscillations of the OB differs when the downstream observer filters the input through different theta LFP phase windows. We focused on the OB where it has been suggested that amplitude modulation may be used by higher-order processing centers such as PC to decode stimulus identity (Freeman and Schneider, 1982; Heck et al., 2019). We found that identity could not be decoded from beta or high gamma OB LFP power when the mouse was naïve to the identity of the reinforced odorant (left panels in Figure 4B). However, the identity of the odorant could be decoded from the PRP after the animal learned to discriminate between the odorants (right panels in Figure 4B). Furthermore, decoding was more accurate when the high gamma LFP power was observed at the peak of the theta oscillation, but it was similar at both the peak and trough phases for beta LFP oscillations. LDA correctly identified the stimulus with a time course for decision making similar to the behavioral readout (licking on the water-delivery spout, Figure 5). Thus, stimulus representation by theta phase-referenced beta and high gamma oscillations of the OB evolves over the course of learning. To our knowledge this is the first demonstration that decoding of stimulus identity from theta phase-referenced high gamma oscillations, reflecting synchronized neuronal firing, changes with learning in brain areas involved in early sensory processing.

The importance of PAC in information processing is based on the notion that different information is carried at different phases of the slow carrier oscillation. This discretization of information is a potential mechanism for encoding short-term memories (Lisman and Idiart, 1995) and more broadly as a neural communication protocol, the theta-gamma neural code (Hopfield, 1995; Lisman and Jensen, 2013). Circuit activation at theta frequency was discovered to be essential for long-term potentiation in the rabbit hippocampus (Bliss and Lomo, 1973), arguing that neural information transfer may go beyond mean rate codes and involve spike timing with respect to oscillatory timekeepers. Subsequent studies indicated that temporal windows set by theta cycles allow for local circuit interactions and a considerable degree of computational independence in subdivisions of the entorhinal cortex-hippocampal loop (Mizuseki et al., 2009). For CA1, it was hypothesized that theta partitions the encoding of new information versus the retrieval of stored information (Hasselmo et al., 2002). Indeed, optogenetic recruitment of fast spiking inhibition in CA1 delivered at specific theta phases altered murine performance in a spatial navigation task (Siegle and Wilson, 2014). When the mice were encoding a location, triggering of inhibition of CA1 at theta peaks improved performance, whereas inhibition at theta troughs improved performance during the retrieval phase. Recently, age related cognitive decline in working memory in humans was reversed temporarily via transcranial alternating-current stimulation frequency-tuned to the endogenous theta peak frequency of each individual’s frontotemporal network (Reinhart and Nguyen, 2019). It appears that working memory is undergirded by theta-gamma PAC and that inter-regional communication is facilitated by theta phase coherence (Daume et al., 2017; Reinhart

and Nguyen, 2019). Thus, the theta LFP which has been cited as one of the most global oscillations in the brain acts as a timekeeper (Siegle and Wilson, 2014) that is coherent across numerous cortical and subcortical structures arguing for its role in transfer of discrete chunks of information (Buzsaki, 2002).

Here we asked whether odor identity information was carried at different phases of theta. Our study of PAC in the OB involved learning of odor discrimination in the go no-go task. In orbitofrontal cortex, Van Wingerden et al. (2014) showed that this task elicited strongest theta/gamma PAC during odor sampling preceding correct decision-making, which they attributed to temporal coordination between cell assemblies responsible for stimulus representation and reward association. The OB circuit is characterized by beta and gamma oscillations nested in a slower theta carrier wave pertaining to respiration (Kay, 2005; Kepecs et al., 2006; Rebello et al., 2014; Rojas-Libano et al., 2014; Heck et al., 2019). However, Granger directionality analysis shows that the theta oscillation in the OB is also influenced by hippocampal slow oscillations indicating that the origin of OB slow carrier oscillations is complex, involving both respiration and centrifugal input from downstream brain regions whose relative contribution likely changes under different behavioral states (Nguyen Chi et al., 2016). Cyclical activation of OSNs through respiration (Adrian, 1950; Bressler, 1988; Rosero and Aylwin, 2011; Zhuang et al., 2019) provides input to mitral/tufted (M/T) cells, and coordinated activity by M/T cells in the OB caused by mitral cell-granule cell reciprocal dendrodendritic interactions establishes gamma oscillations (Stopfer et al., 1997; Schoppa, 2006; Osinski and Kay, 2016; Arnson and Strowbridge, 2017; Pouille and

Schoppa, 2018)—see review by Heck et al., 2019. Gamma LFP is thought to represent the stimulus identity (Bathellier et al., 2006; Beshel et al., 2007; Rojas-Libano and Kay, 2008; Li and Cleland, 2017) and beta LFP oscillations (15-35 Hz) have been linked to learned reward (Martin et al., 2006). However, the classification of these two bandwidths as carrying learning vs. odor identity information is not always appropriate, and Frederick et al. (Frederick et al., 2016) proposed that they represent different cognitive states that vary depending on particular behavioral demands. PAC is present in the OB (Rojas-Libano et al., 2014) and it may be of great albeit relatively unexplored utility. We contributed to the understanding of which information is carried by different frequency oscillations by asking whether the high gamma and beta oscillations carry different information when they are observed within particular phase windows of the theta carrier wave (Figure 4). We hypothesized that PAC may be used to convey olfactory stimulus information to theta phase-synchronized downstream regions like PC (Kepecs et al., 2006; Lisman and Buzsaki, 2008).

We demonstrate that PAC changes in the OB during associative learning (Figure 2). Coupling strength (modulation index) increased for the rewarded stimuli and peak angle variance increased for the unrewarded odorants with learning. The theta phase-referenced PRP of the high gamma and beta LFPs also follows this trend (Figure 3), increasing for S+ and decreasing with S- over learning. The changes in PRP correlating with learning may reflect contrast enhancement that facilitates the discrimination between similar stimuli by a downstream observer. Neuronal adaptation to facilitate contrast enhancement has been shown at the level of the OSNs (Haney et al., 2018),

the glomerular layer (Cavarretta et al., 2016), and the granule cell layer (Adams et al., 2019) of the OB. We speculate that the increase in PRP for S+ (concurrent with decrease in PRP for S-) reflects adaptational mechanisms in the OB to enhance contrast between similar stimuli with different associated values.

To quantify whether these changes in PAC increased discriminability between odorants using PRP we asked whether the stimulus that was delivered could be decoded from PRP data. Previous investigations have decoded olfactory stimuli from spike counts aligned to sniffing in the mouse (Cury and Uchida, 2010; Gschwend et al., 2012), in the antennal lobe of the locust (Stopfer and Laurent, 1999), and M/T neural activity assayed through calcium transients in mice (Chu et al., 2016), but trial classification via PAC remains unexplored. Here, we performed LDA on PRP data to determine whether a single hyperplane could differentiate between stimuli as the animals learned the task. Figure 4B indicates that the identity of the stimuli cannot be decoded from PRP recorded in the naïve mice, but that the stimuli can be decoded when the mice learn to discriminate the odorants. Furthermore, decoding the stimuli from peak high gamma PRP is significantly better than trough PRP, and decoding does not differ between peak and trough for beta PRP. Thus, decoding of the odorant identity by downstream observer using a phase window for the slow carrier oscillation could detect the identity of the odorants after the animal learns to differentiate between odors in the go no-go task. Finally, this study contributes to the mounting evidence that the OB is an important locus for learning-related activity changes (Doucette and Restrepo,

2008; Doucette et al., 2011; Abraham et al., 2014; Li et al., 2015; Chu et al., 2016; Koldaeva et al., 2019).

We wondered how timing for decision-making would relate comparing stimulus decoding with PRP with decision-making estimated from the behavioral readout of the olfactory-driven decision in the go no-go task (differential licking between the two odorants). The time course for the behavioral decision, reflected in the graph of p value for differential licks, is overlaid by the p value time course for peak PRP decoding (Figure 5C, right) with trough PRP lagging behind. Neither beta nor high gamma peak PRP p values are discriminable from lick p values in terms of decision time in this task. As shown in Figure 5D, for proficient animals licking differs between odorants within about 0.5 seconds from stimulus onset, at the same time that PRP decoding becomes significant. This is slower than discrimination times for odorant responses recorded in mice performing two alternative forced choice odorant discrimination or receiving odors passively where discrimination takes place within the first respiratory cycle (~250 ms) after stimulus onset (Cury and Uchida, 2010; Shusterman et al, 2011, Short and Wachowiak, 2019). However, the 0.5 sec decision-making time is consistent with studies with the go no-go task where there is not a strong motivation for the animal to stop licking for the unrewarded odorant (Abraham et al., 2004; Rinberg et al., 2006; Doucette and Restrepo, 2008; Jordan et al., 2018)—with the exception of rapid detection of optogenetic activation of glomerular input in the go no-go task (Smear et al., 2011). Thus, we conclude that PRP is an accurate reflection of decision-making in olfactory discrimination.

Finally, we turn to dimensionality. The measurements in the experiments presented here are inherently low dimensional because: 1) The animal is discriminating between two odors, and the behavior is constrained by the go no-go task (Gao et al., 2017) and 2) While the number of independent degrees of freedom for inputs to the OB is potentially large (~2000 glomeruli), the four tetrodes are close together and they should sample a small subset of these input dimensions. Indeed, because the tetrodes are within 200 μm of each other (and electrodes are adjacent within tetrodes), the LFP recorded from each electrode is likely to be similar, assuming that the cellular milieu around the electrodes does not separate the inputs (Gold et al., 2006). Interestingly, the dimensionality measured from the 16 electrodes (four tetrodes) varies between 1 and 6 suggesting that subsets of electrodes are detecting different signals (Figure S6 B and C). However, as expected, when the mice learn to discriminate the dimensionality decreases in the presence of the odorant (Figures 6 and S6). Furthermore, the dimensionality does not decrease for high gamma trough PRP. This likely reflects experience-dependent, selective attention (Cañas and Jones, 2010; Niv et al., 2015). In the naïve state, relevant dimensions are unknown to the mouse. After learning to lick the waterspout in the presence of an odorant stimulus during training, they lick indiscriminately for both S+ and S-. However, reward prediction error accumulates when they are repeatedly unrewarded for licking for S-. Attention, guided by this reward prediction error, biases/constrains the learning to the relevant dimensions (Leong et al., 2017)—in this task, odor identity.

The changes in signal processing elicited by learning, habituation, and passive odorant exposure in the OB, the earliest processing center in the olfactory system, are remarkable. Why is processing in this circuit so plastic and influenced by centrifugal input from downstream targets, such as PC, and from a variety of neuromodulatory centers (Gire et al., 2013b; Linster and Cleland, 2016)? Likely the reason is that this processing center receives massive parallel glomerular input and projects M/T output of high dimensionality to downstream brain processing areas (Chae et al., 2019). Furthermore, unsupervised learning modifying cellular interactions in the OB would increase dimensionality (Hiratani and Latham, 2019). Therefore, odors must be identified in the midst of a high dimensionality cocktail party (Rokni et al., 2014). Based on the present data showing a large increase in the accuracy for decoding of odor identity from the PRP as the animal learns to differentiate the odorants we speculate that when a subset of the degrees of freedom of the sensory input become behaviorally relevant, centrifugal feedback elicits supervised learning that results in a decrease in dimensionality of the output from the OB to downstream brain areas. This would result in more accurate discrimination of the input channels that are relevant to the current task. Whether (and how) supervised learning-induced OB circuit plasticity results in a decrease in dimensionality of the output remains to be determined.

Acknowledgments: We thank Nicole Arevalo for animal husbandry, Dnate' Baxter for laboratory support, and Arianna Gentile-Polese for immunohistochemical assays. This research was supported by NIDCD 1F31DC016483-01A1 to J.L. and NIDCD 5R01DC000566-30 to D.R.

Author contributions: J.L., D.R-G., and D.R. designed all experiments, while J.L. and D.R-G. performed them. J.L. and D.R. set up the Intan recording system for LFP data. D.R. wrote all programs for data collection and analysis. J.G. shared code with D.R. regarding dimensionality and influenced the development of the dimensionality work. J.L., D.R-G., and D.R. analyzed the data. J.L. and D.R. generated the figures. All authors discussed the results. J.L. and D.R. wrote the manuscript, and all authors participated in commentary and revision of the manuscript.

Declaration of interests: The authors declare no competing interests.

METHOD DETAILS

Animals

The study was performed in two sets of experiments. For the first set of experiments (exp1), DBH-eNpHR3.0 mice were used as per Ramirez-Gordillo et al., 2018. These mice express halorhodopsin under the dopamine beta hydrolase promoter (DBH-Cre eNpHR3.0) in LC-NA neurons and were produced by crossing DBH-Cre mice (032081-UCD, Mutant Mouse Resource and Research Center) with mice expressing halorhodopsin in a Cre-dependent fashion [eNpHR3.0, 129S Gt(ROSA)26Sortm39(CAG-hop/EYFP) Hze/J, Jackson labs]. C57BL/6 mice (Jackson stock number: 000664) were used for the second set of experiments (exp2). All animal procedures were performed under approval from the Institutional Animal Care and Use Committee (IACUC) at the University of Colorado Anschutz Medical Campus under guidelines from the National Institutes of Health. The number of mice used in each experiment is shown in Table S1.

(Op)tetrode implantation

Surgery was performed under approval from the Institutional Animal Care and Use Committee (IACUC) at the University of Colorado Anschutz Medical Campus, using aseptic technique. As per Li *et al.*, 2015, tetrode boards (EIB-16, Neuralynx) with optional fiber optic cannula for photostimulation (optetrodes) were populated with four tetrodes consisting of four 12.5 µm nichrome wires coated with polyimide (Sandvik RO800). For optetrode fabrication the tetrode boards received the addition of a fiber optic cannula 105 µm in diameter (Thorlabs FG105UCA) with 1.25 mm OD ceramic

ferrule (Precision Fiber Products) for coupling to a DPSS laser (473 nm, Shanghai Laser and Optics Century). The optogenetic experiments of Ramirez-Gordillo et. al. 2018 are not included in this study. In the rest of the text we refer to the inserts as (op)tetrode with the understanding that optetrodes were implanted for exp1 while tetrodes were implanted for exp2. Electrode tips were electroplated to 0.2-0.4 MΩ impedance.

Two-month-old male mice were anesthetized with 5% isoflurane in oxygen. Intraperitoneal ketamine/xylazine (100 mg/kg and 10 mg/kg respectively) was then administered along with 100 µl of 2% lidocaine injected subcutaneously over the skull. After the mouse was found to be unresponsive to a toe pinch, the animal's head was then secured in the stereotaxic apparatus (Narishige SR-5M-HT) and the skull was leveled (≤ 50 µm difference DV between bregma and lambda). Gentamycin ophthalmic ointment was applied to the eyes to maintain hydration. After incising the skin overlaying the skull, the periosteum was cleared with 15% H_2O_2. A manipulator (Sutter MP-285) was zeroed at bregma and midline and the target location for OB implantation was marked with respect to bregma (exp1: AP +4.28 mm, ML +0.4 mm, exp2: AP +4.28 mm, ML +0.5 mm).

A craniotomy performed at this site (Marathon III drill) exposed dura mater which was removed prior to implantation. Another craniotomy was performed more caudally for implantation of one ground screw (Plastics1 00-96 x 1/16). The (op)tetrode was positioned above the craniotomy over the OB while the ground wire was wrapped

around the ground screw with the connection coated in silver paint (SPI Flash-Dry silver conductive paint). After securing the ground screw to the skull, the (op)tetrodes were lowered into position at the rate of 1 mm/minute (exp1: AP +4.28 mm, ML +0.4 mm, DV 0.53 mm, exp2: AP +4.28 mm, ML +0.5 mm, DV 1.0 mm). After reaching the target depth, the optetrode was adhered to the skull with C&B Metabond, followed by Teets 'Cold Cure' dental cement. After curing (10 minutes), the (op)tetrode was detached from the manipulator, the animal was removed from the stereotax and received subcutaneous injections of carprofen (10 mg/kg) and buprenorphine (0.05 mg/kg) and recovered on a heating pad kept at 37°C. The mice were monitored daily and received additional carprofen injections daily for the first two days postoperatively.

Go-no go behavioral task

Mice were water deprived until they reached 80% normal body mass. Then they were placed into the Slotnick olfactometer (Bodyak and Slotnick, 1999; Li et al., 2015) chamber where they could move freely. All mice were first trained to lick the water spout to obtain water in the presence of odor (1% isoamyl acetate in mineral oil, v/v) in the "begin" task (Slotnick and Restrepo, 2005), then to lick for 1% isoamyl acetate (S+) versus mineral oil (S-) in the "go no-go" task (Doucette et al., 2011; Li et al., 2015). Mice self-initiated trials by poking their head into the odor delivery port, breaking a photodiode beam (Figure 1A). During odor delivery (lasting 2.5 seconds) they must lick a water delivery spout at least once during each of four 0.5 second-long response areas in order to register the decision as a lick (Figure 1B), detected as electrical connectivity between the water spout and the ground plate on which they stand (Slotnick and

Restrepo, 2005). If the mice licked during a rewarded odorant trial, they received ~10 µl water reinforcement (Figure 1C). The mice learn to refrain from licking for the unrewarded odorant due to the unrewarded effort of sustained licking. Performance was evaluated in blocks of 20 trials, with 10 S+ and 10 S- trials presented at random. Animals performed as many as 10 blocks per session. Sessions were terminated when animals demonstrated satiety/disengagement from the task or when they performed at or above 80% correct discrimination in three or more blocks in a session. Once criterion was reached, the next session had reversed valence, meaning that the previous S+ was S- and the previous S- was S+. This reversal disambiguates the identity of the odor from its valence (reversed session data are not shown). Data were analyzed within two performance windows: when the animal was performing below 65% (naïve) or above 80% (proficient). The data were also analyzed in the 65-80% performing window, and the results fell between the naïve and proficient windows (data not shown).

Odor stimulus delivery time was measured with a photoionization detector (miniPID, Aurora Scientific). The time difference between valve opening and detection of odor at the odor port was between 66 – 133 ms, depending on which olfactometer was used.

Neural recording

Extracellular potentials from the four tetrodes were captured and digitized at 20 kHz on the RHD2216 amplifier of the Intan RHD2000 Evaluation System with a 1-750 Hz bandpass filter. Information on behavioral events (valve times, mouse presence at

the odor port) was sent through a digital line from the Slotnick olfactometer to the Intan board. Licks detected by the Slotnick olfactometer were recorded as an analog signal by the Intan board.

PAC

PAC data were processed using the Hilbert transform method described by Tort *et al.*, 2010. Briefly, data were bandpass filtered with a 20th order Butterworth filter using Matlab's filtfilt function with zero phase shift to extract LFP in the low frequency oscillation used for phase estimation (theta, 2-14 Hz, Figure 1Diii) and the high frequency oscillation used for estimation of the amplitude of the envelope (either beta, 15-30 Hz, or high gamma, 65-95 Hz, Figure 1v). Hilbert transform established the theta phase (Figure 1Div) and, separately, the envelope for beta or high gamma (red line in Figure 1Dv). A plot of the distribution of beta/high gamma amplitude in 51 bins of the phase of theta (PAC) was plotted to gauge the strength of PAC (Figure 1E). To quantify the strength of PAC we calculated the MI estimating the KL distance to quantify the difference between the observed beta/high gamma amplitude distribution along the phase of theta from a uniform distribution. If PAC is non-existent, MI = 0, meaning the mean amplitude is distributed uniformly over theta phases, and if PAC is a delta function MI = 1. MI for signals measured in brain areas such as the hippocampus typically fall between 0-0.03 (Tort et al., 2010).

Phase-referenced LFP power

Beyond the mere existence of PAC that has been documented in numerous brain regions, we sought to evaluate whether the information carried by beta or high gamma power at different theta phases could be used to discriminate between olfactory stimuli. To accomplish this analysis, we developed the PRP approach using custom Matlab code. Briefly, PAC was calculated using the approach documented by Tort *et al.*, 2010, as described above and summarized in Figure 1. Peak and trough theta phases are defined as the phase for maxima and minima of the PAC distribution measured for the S+ trials (red plot in Figure 2C). A continuous Morlet wavelet transform was used to estimate the power for the high frequency oscillations. PRP was estimated as the power of the high frequency oscillations (beta or high gamma) measured at the peak or trough of PAC (Figure 4). The Matlab code used for data analysis has been deposited to https://github.com/restrepd.

Statistical analysis

Statistical analysis was performed in Matlab. PAC parameters and PRP were calculated separately per electrode (16 electrodes per mouse) for all electrodes per mouse. Statistical significance for changes in measured parameters for factors such as learning and odorant identity was estimated using a GLM, with post-hoc tests for all data pairs corrected for multiple comparisons using false discovery rate (Curran-Everett, 2000). The post hoc comparisons between pairs of data were performed either with a t-test, or a ranksum test, depending on the result of an Anderson-Darling test of normality. 95% confidence intervals (CIs) shown in the figures as vertical black lines

were estimated by bootstrap analysis of the mean by sampling with replacement 1000 times.

PAC Simulation

To validate our analysis of PAC, we performed simulations with uncoupled oscillations (no PAC) with a low frequency cosine oscillation at 8 Hz and a high frequency cosine oscillation at 40 Hz (Figure S1A) or with coupled oscillations where high frequency bursts took place at 180 degrees for the low frequency cosine oscillation (Figure S1B). The bursts were generated by a cosine oscillation at the high frequency multiplied by gaussians with a full width half maximum of (1/8 Hz)/2 centered at 180 degrees of the low frequency cosine wave. Top panels in both S1A and S1B are pseudocolor Morlet wavelet power time courses, while the bottom panels show the PRP of the 40 Hz oscillation. For the case with no PAC (Figure S1A), peak and trough PRP are equal. However, for the PAC example (Figure S1B) the peak PRP is significantly higher than trough PRP.

Linear discriminant analysis

Classification of trials using PRP was accomplished via LDA in Matlab whereby PRP for every trial except one were used to train the LDA, and the missing trial was classified by its fit into the pre-existing dataset. This was repeated for all trials and was performed separately for peak and trough PRP, and for analysis where the identity of the odorants was shuffled.

Dimensionality

Following Litwin-Kumar et al. (2017) we defined the dimension of the system (dim) with *M* inputs as the square of the sum of the eigenvalues of the covariance matrix of the measured PRP LFP (*pLFP*) divided by the sum of each eigenvalue squared:

$$\dim\,(pLFP) = \left(\sum_{i=1}^{M} \lambda_{\mathrm{i}}\right)^{2} / \left(\sum_{i=1}^{M} {\lambda_{\mathrm{i}}}^{2}\right)$$

where λ_i are the eigenvalues of the covariance matrix of *pLFP* computed over the distribution of PRP LFP signals measured in the OB. If the components of *pLFP* are independent and have the same variance, all the eigenvalues are equal and dim(*pLFP*) = *M*. Conversely, if the *pLFP* components are correlated so that the data points are distributed equally in each dimension of an *m*-dimensional subspace of the full *M*-dimensional space, only *m* eigenvalues will be nonzero and dim(*pLFP*) = *m*.

REFERENCESBodyak, N., and Slotnick, B. (1999). Performance of mice in an automated olfactometer: odor detection, discrimination and odor memory. Chem Senses 24, 637-645.

Chery, R., Gurden, H., and Martin, C. (2014). Anesthetic regimes modulate the temporal dynamics of local field potential in the mouse olfactory bulb. J Neurophysiol 111, 908-917.

Curran-Everett, D. (2000). Multiple comparisons: philosophies and illustrations. Am J Physiol Regul Integr Comp Physiol 279, R1-8.

David, F., Courtiol, E., Buonviso, N., and Fourcaud-Trocme, N. (2015). Competing Mechanisms of Gamma and Beta Oscillations in the Olfactory Bulb Based on Multimodal Inhibition of Mitral Cells Over a Respiratory Cycle. eNeuro 2.

Doucette, W., Gire, D.H., Whitesell, J., Carmean, V., Lucero, M.T., and Restrepo, D. (2011). Associative cortex features in the first olfactory brain relay station. Neuron 69, 1176-1187.

Gold, C., Henze, D.A., Koch, C., and Buzsaki, G. (2006). On the origin of the extracellular action potential waveform: A modeling study. J Neurophysiol 95, 3113-3128.

Li, A., Gire, D.H., and Restrepo, D. (2015). Upsilon spike-field coherence in a population of olfactory bulb neurons differentiates between odors irrespective of associated outcome. J Neurosci 35, 5808-5822.

Litwin-Kumar, A., Harris, K.D., Axel, R., Sompolinsky, H., and Abbott, L.F. (2017). Optimal Degrees of Synaptic Connectivity. Neuron 93, 1153-1164.e1157.

Nguyen Chi, V., Muller, C., Wolfenstetter, T., Yanovsky, Y., Draguhn, A., Tort, A.B., and Brankack, J. (2016). Hippocampal Respiration-Driven Rhythm Distinct from Theta Oscillations in Awake Mice. J Neurosci 36, 162-177.

Slotnick, B., and Restrepo, D. (2005). Olfactometry with mice. Curr Protoc Neurosci Chapter 8, Unit 8 20.

Tort, A.B., Komorowski, R., Eichenbaum, H., and Kopell, N. (2010). Measuring phase-amplitude coupling between neuronal oscillations of different frequencies. J Neurophysiol 104, 1195-1210.

Vizcay, M.A., Duarte-Mermoud, M.A., and Aylwin Mde, L. (2015). Odorant recognition using biological responses recorded in olfactory bulb of rats. Comput Biol Med 56, 192-199.

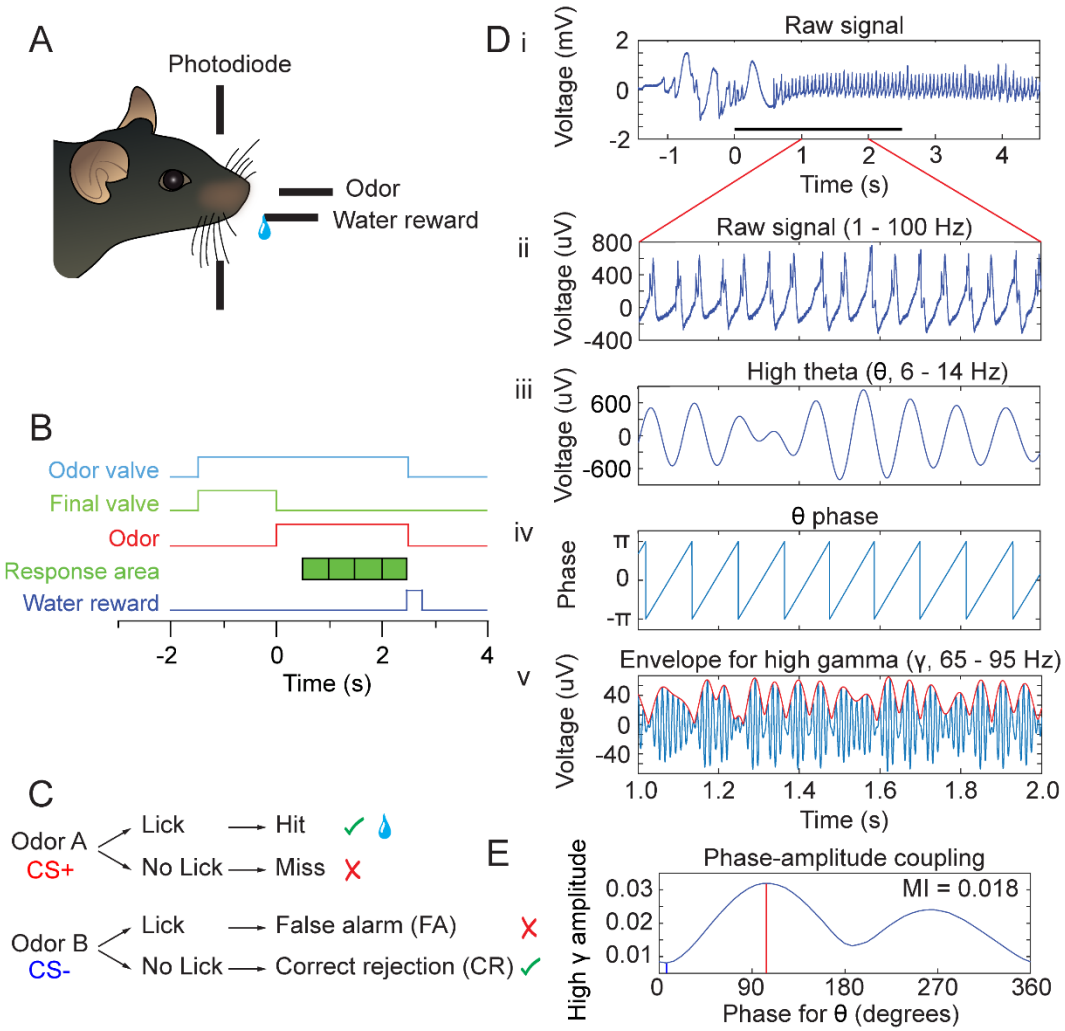

Figure 1. Behavioral task and PAC analysis.

(A) Behavioral apparatus. Mouse self-initiates trials by breaking a photodiode beam. Odorants and water are delivered according to the timeline in B and the decision tree in C.

(B) Timeline for a single trial. When the animal enters the port the air flow is diverted by turning the final valve output to the exhaust and the odorant valve is turned on. At this time the odor builds up to steady-state concentration flowing away from the animal for

1-1.5 seconds. At time 0 the final valve air flow is turned towards the odor port for 2.5 s resulting in odor onset 100 ms after the valve is actuated. In order to trigger water reward for the rewarded odorant the mouse must lick at least once during each 0.5 s block for four blocks in the response area. If the stimulus is rewarded and the animal licked appropriately, a water reward is delivered.

(C) Decision tree for each stimulus. Green check mark means the correct decision was made. Red “X” mark indicates an incorrect decision. Water reward is represented by the water droplet symbol in the case of a hit.

(D) PAC data analysis of the LFP. For each electrode, raw signal collected at 20 kHz (i) is bandpass filtered to 1-100 Hz for broadband LFP (ii) or filtered into different frequency bands as needed (e.g. theta 6-14 Hz in iii or high gamma 65-95 Hz in v, blue line). Hilbert transform is used to calculate the theta phase (iv) and the amplitude envelope of higher oscillations such as high gamma (65-95 Hz) (red line in v). Theta phase and the envelope of the amplitude of high gamma are then used to calculate the probability for high gamma amplitude at specific phases of theta (PAC, shown in E).

(E) Probability for high gamma amplitude at specific phases of theta. In this example, PAC is strongest at ~100° and 270°. The coupling strength is quantified by the modulation index, MI. Peak phase (the phase of high theta with the highest amplitude in high gamma) is indicated with a red line and trough phase is indicated with a blue line.

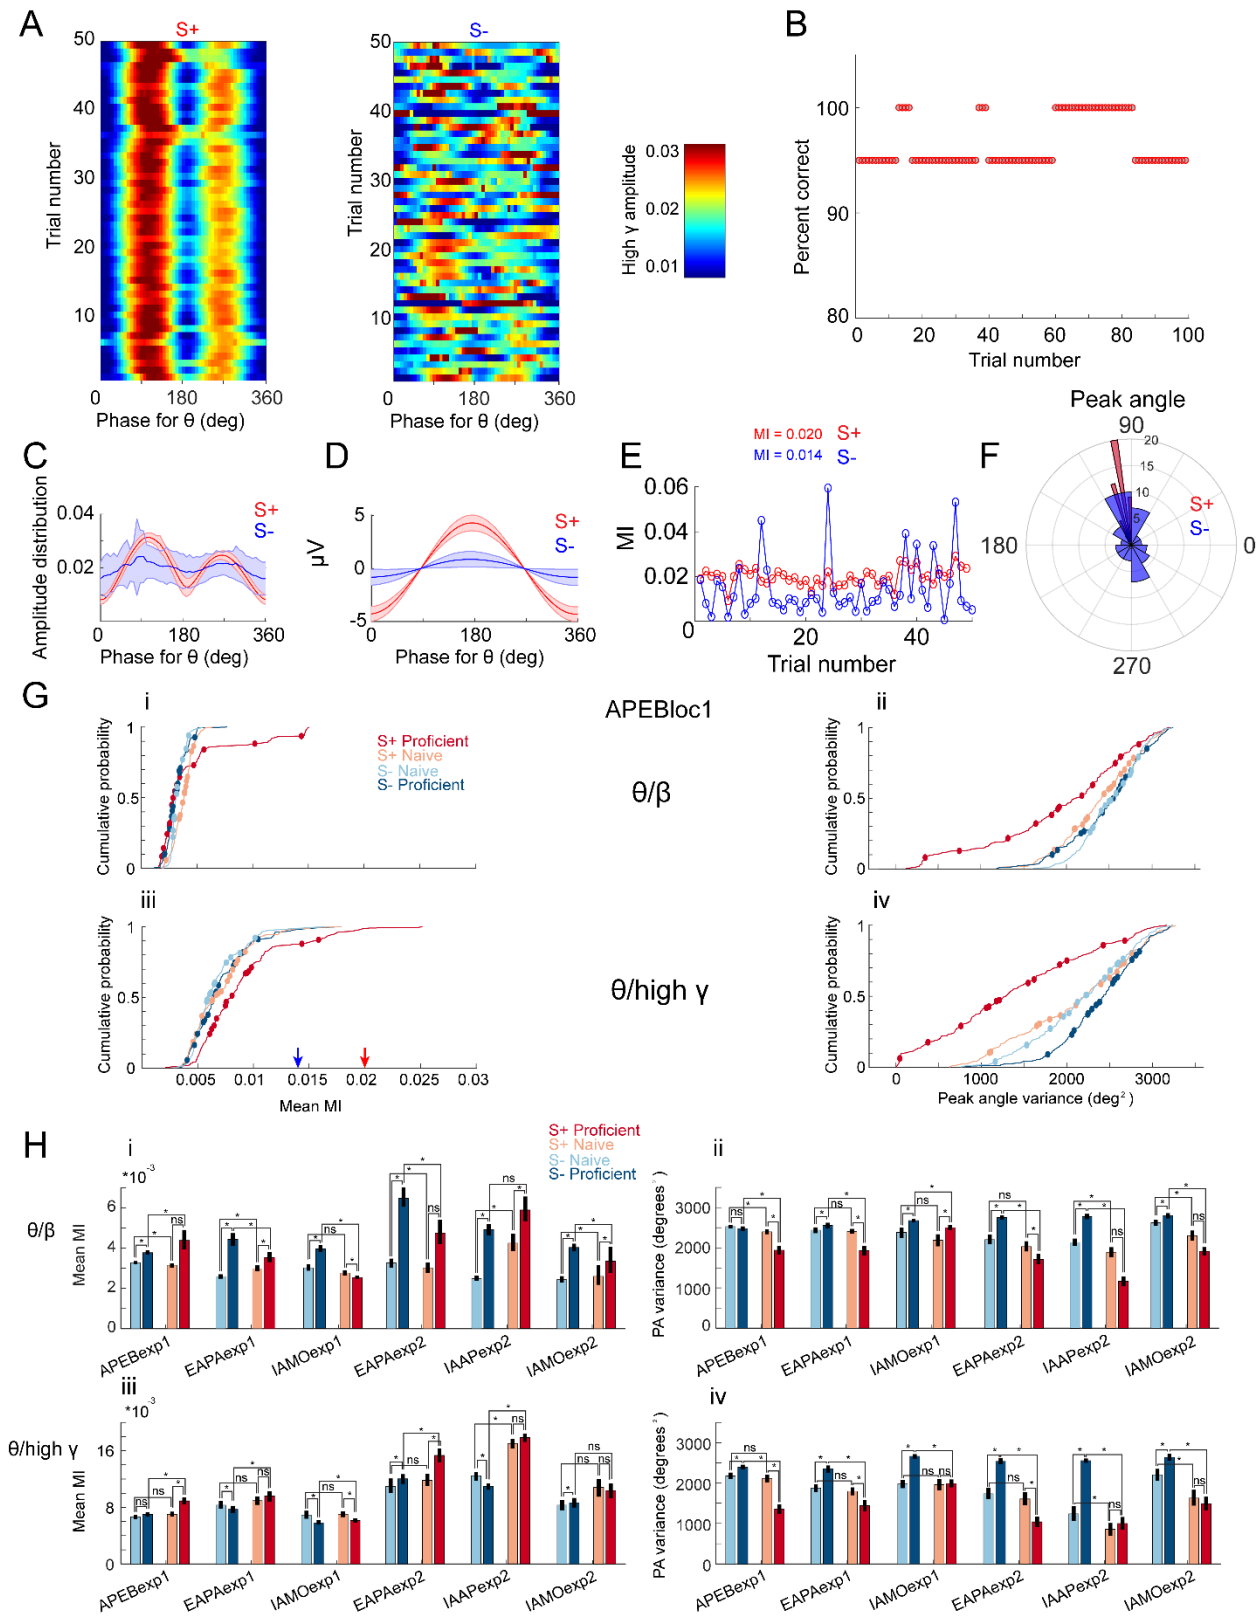

Figure 2. The peak angle variance of PAC changes as the animal learns to

discriminate odorants.

(A) Pseudocolor plots showing the phase amplitude relationship for S+ (left) and S- (right) for an example go no-go session. PAC between theta (6-14 Hz) phase and high gamma (65-95 Hz) amplitude was more variable for S- compared to S+.

(B) Behavioral performance for this session. Over 100 trials, this trained animal discriminated between isoamyl acetate and acetophenone with 95-100% accuracy.

(C) Theta phase distribution for high gamma amplitude (mean $\pm$ CI) for S+ (red) and S- (blue).

(D) Mean theta LFP (mean $\pm$ CI) for one cycle for S+ (red) and S- (blue).

(E) MI per trial for S+ (red) and S- (blue) as a function of trial number. The difference in MI between S+ and S- is statistically significant (ranksum p value < 0.001, n=50 S- trials, 49 S+ trials).

(F) Polar histograms for high gamma peak theta angles for S+ (red) and S- (blue).

(G) Cumulative histograms showing the differences in mean MI (i,iii) and peak angle variance (ii,iv) for theta/beta PAC (i,ii) and theta/high gamma PAC (iii,iv) for the APEB odorant in location 1 for the S+ and S- odorants for naïve and proficient trial ranges. Peak angle variance for S+ displays a substantial decrease with learning, while mean MI shows small changes. Changes with learning and between events for mean MI and peak angle variance are statistically significant (GLM p value < 0.001, 896 d.f. for theta/beta and theta/high gamma, 16 electrodes per mouse, 14 mice). Blue and red arrows in Gi are MIs derived from single-session example in C.

(H) Bar graph displaying differences in mean MI (i,iii) and peak angle variance (ii,iv) between naïve and proficient trial ranges and S+ and S- events for all sessions for all the odorants. For all odorants GLM yielded significant p values (< 0.001) for both mean MI and peak angle variance. These parameters are shown for theta/beta (i,ii) and theta/high gamma PAC (iii,iv). GLM analysis p value was <0.001, for event, proficiency and location, with the exception that p < 0.01 for proficiency for theta/high gamma, 3255 d.f., number of mice is in Table S1. Asterisks denote a significant post-hoc difference evaluated with either a t-test or a ranksum test (p < pFDR, for the number of mice per odorant, see Table S1). Changes in mean MI were variable for different odorants. MI displays differences between events and learning states that differ in magnitude between odorant pairs. Peak angle variance (at right) increases with task proficiency for S- and decreases for S+ for all odorants.

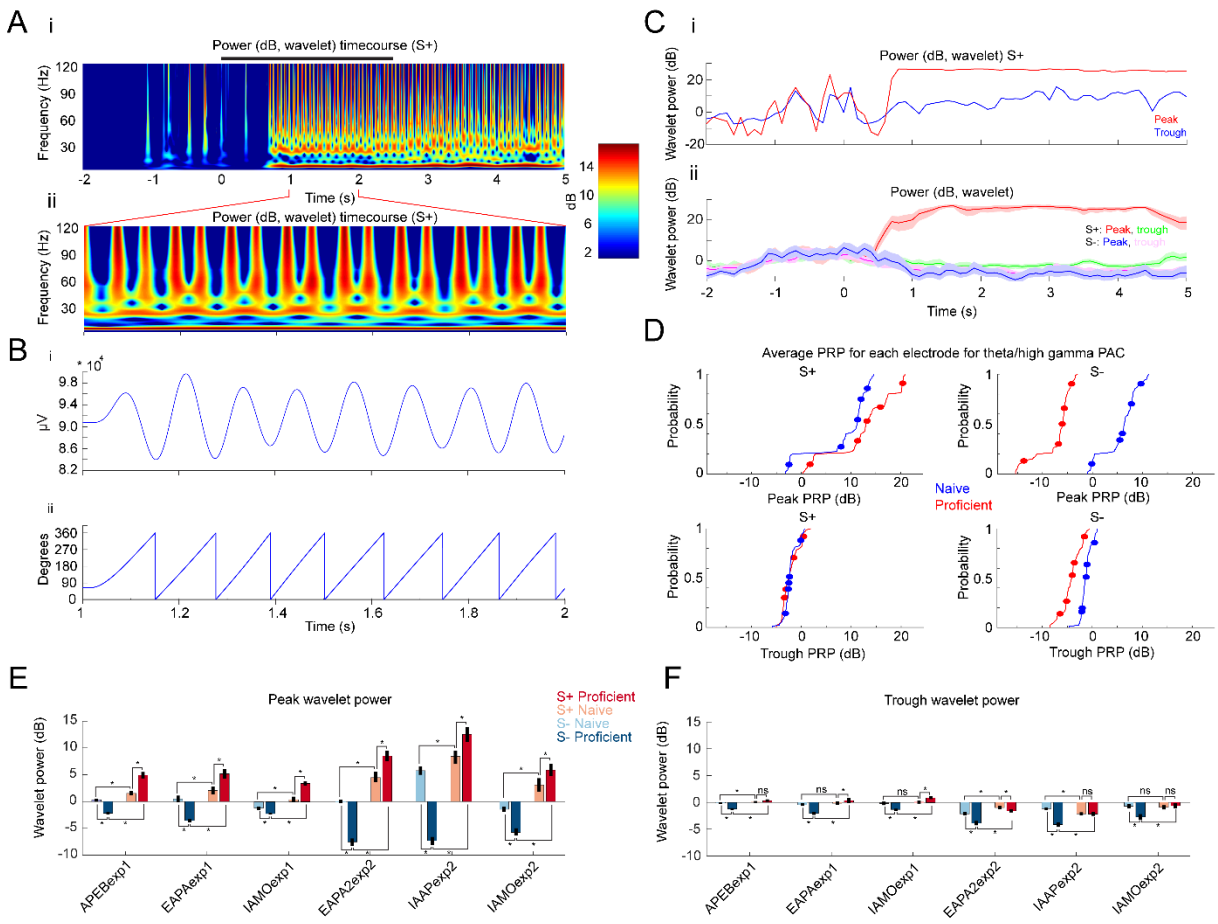

Figure 3. PRP increases for S+ and decreases for S- over learning.

(A) Example wavelet broadband LFP spectrogram for one S+ trial. Right: Pseudocolor scale for LFP power (in dB). Black bar represents the odor stimulation epoch (2.5 s). Ai: Full trial length demonstrating pre-odor baseline LFP, odor driven LFP, and reward epoch. Aii: One-second epoch during odor stimulation.

(B) Bi: Theta waveform (6-14 Hz) during odor stimulation. Bii: High theta phase, extracted with Hilbert transform.

(C) Ci: Single trial wavelet power referenced to either peak (108°, red) or trough (0°, blue) of the high gamma amplitude distribution along the theta phase. Cii: PRP (± CI) averaged across the entire session for S+ and S-. Peak PRP for S+ displays the largest odor-induced increase in power that is sustained after the end of odor stimulation.

(D) The thin lines are cumulative histograms for PRP for the LFP recorded in each electrode averaged for theta/high gamma for all proficient or naïve trials for mice performing to differentiate odorants in the IAAPexp2 odorant pair. As the animals learn the task, peak PRP (top panels) increases for S+ and decreases for S-, whereas trough PRP (bottom panels) is relatively stable. The differences were statistically significant between peak and trough, S+ and S- and naïve and proficient (GLM p values < 0.001, 1016 d.f.). Circles are per mouse averages for PRP.

(E-F) Odor-induced changes in high gamma PRP averages for all odorant pairs evaluated per electrode during the odor delivery epoch (0.5 – 2.5 s). Peak PRP (E) increases for S+ and decreases for S- when the mice become proficient. Trough PRP changes (F) are markedly weaker, yet they follow a similar trend (GLM p value < 0.001, 6386 d.f. for S+ vs. S-, peak vs. trough, naïve vs. proficient). In addition, the learning-induced changes in PRP are larger in location 2 (GLM p value < 0.001 for location, 6386 d.f.).

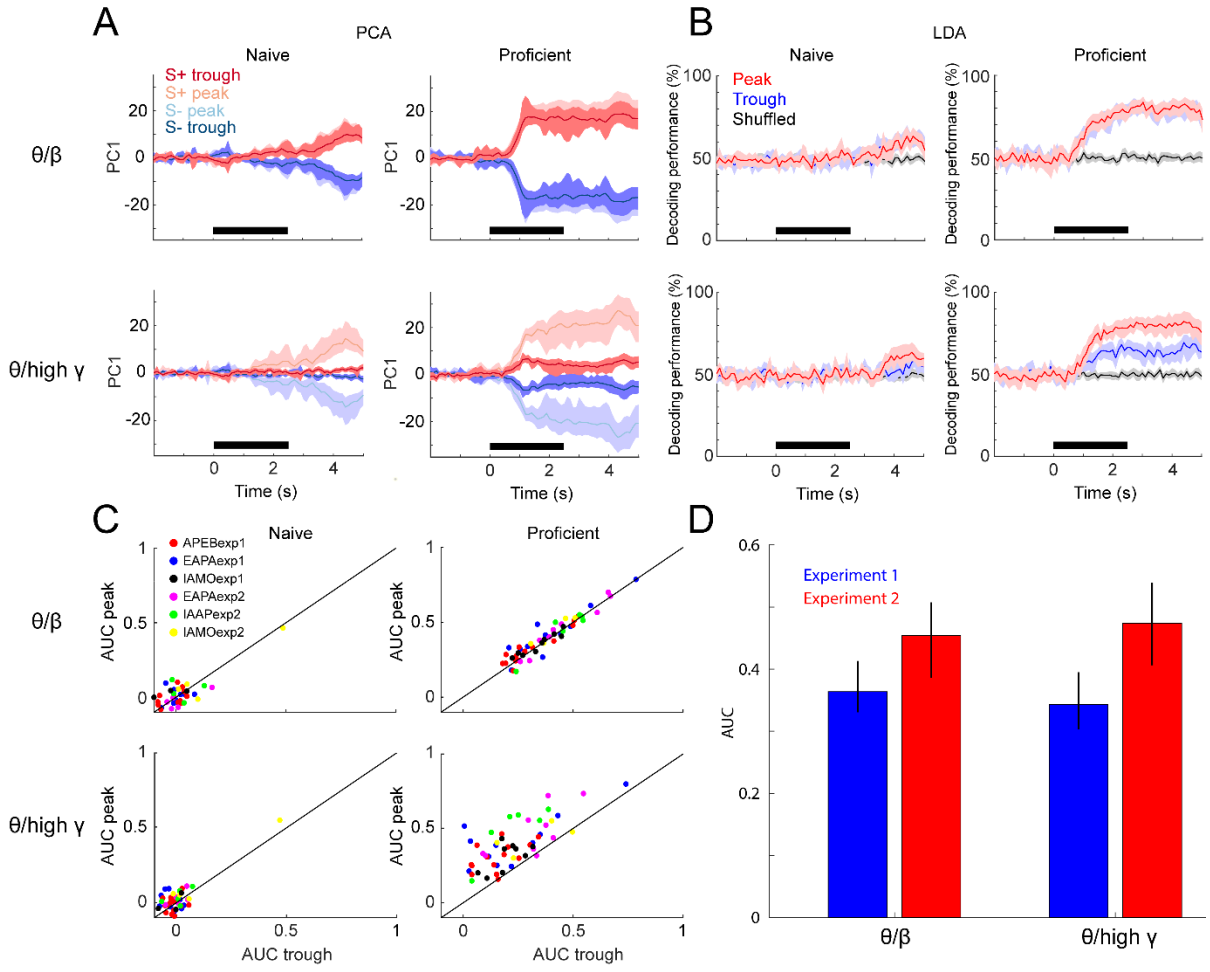

Figure 4. Linear discriminant analysis classifies stimuli using peak PRP.

(A-B) PCA (in panel A) and linear discriminant analysis (LDA; in panel B) of averaged PRP data for one example odorant pair, APEBexp1.

(A) PCA of PRP data from naïve to proficient—theta/beta on top, theta/high gamma on bottom. Both PRPs for S+ and S- are discriminable with learning.

(B) LDA of PRP data over learning, with theta/beta on top, theta/high gamma on bottom. In the proficient state, LDA can decode stimulus identity using both PRPs (theta/beta PRP GLM naïve-proficient p value < 0.001, 104 d.f.; theta/high gamma PRP GLM naïve-proficient p value < 0.001, 104 d.f.). Peak and trough carry similar information for theta/beta; both peak- and trough-referenced beta power can be used to decode

stimulus identity with learning. In theta/high gamma, peak PRP is significantly better at decoding the stimulus identity than trough PRP (GLM p value < 0.001, 104 d.f.), which is still better than shuffled data.

(C) Area under the curve (AUC) for peak versus trough PRP over task learning for all odorant pairs and locations. Top: Theta/beta PRP. Bottom: Theta/high gamma PRP. Generalized linear model statistics show that peak AUC is significantly higher than trough (p value < 0.001 for 5/6 odorant pairs) and that there are significant changes over learning (p value < 0.001 for 5/6 odorant pairs). These data suggest that peak PRP decodes stimulus identity better than trough PRP.

(D) Average peak AUC (±CI) for all odorant pairs for theta/beta and theta/high gamma PRP for the two experimental settings. The learning-induced changes in PRP are larger in exp2 (GLM p value < 0.001 for location, 6386 d.f.).

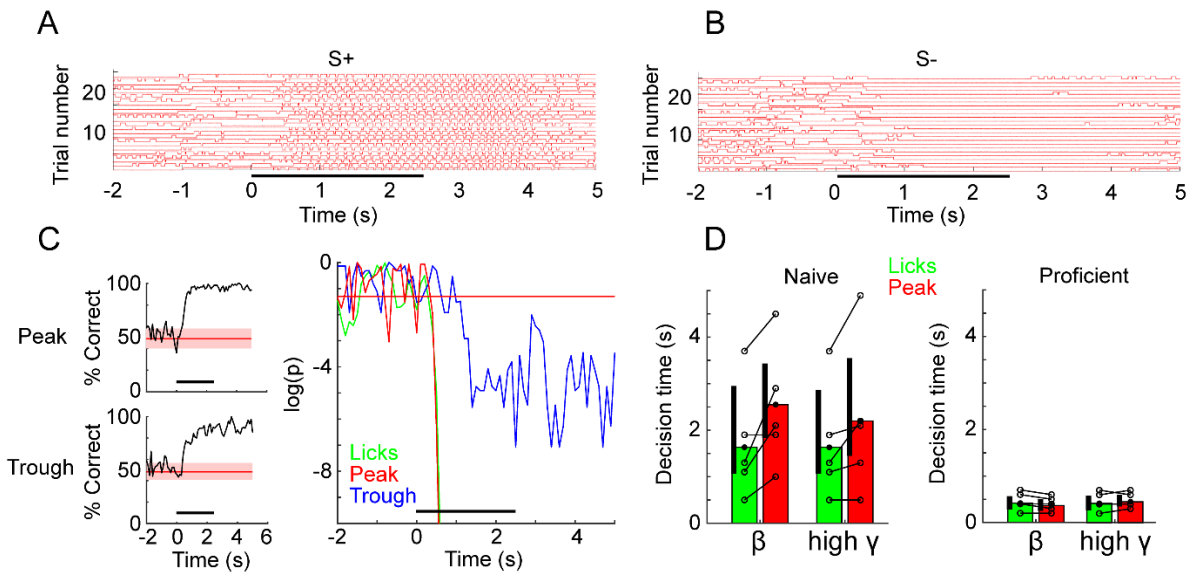

Figure 5. Decision making time does not differ between licks and peak PRP LDA decoding of odorant identity

(A) Lick time course for 25 S+ trials in one session of EAPAexp2 for one proficient mouse. Horizontal black bar indicates odor delivery epoch for all time course subpanels**.**

(B) Lick time course for 25 S- trials in the same session as in A. Mouse refrains from licking in the presence of the unrewarded odorant.

(C) Decoding performance and decision making time for licks and for PRP LDA using data pooled for all mice performing odorant discrimination with the EAPAexp2 odorant pair. Left: LDA decoding performance appears to be faster and more accurate using peak PRP (top) vs. trough PRP (bottom). Right: p value for the ranksum test compares the time-course for decision-making using licks for S+ vs. S- (green) or LDA prediction of odorant identity with peak (red) or trough (blue) PRP. This example illustrates that that peak PRP LDA prediction and lick divergence perform similarly, while LDA prediction with trough PRP is significantly slower in stimulus decoding. Horizontal red line indicates p = 0.05.

(D) Decision times for licks and for pooled animal peak theta/beta and theta/high gamma PRP LDA in naïve (left) and proficient (right) mice. Decision times were significantly reduced with task learning for theta/beta PRP LDA, theta/high gamma PRP LDA and licks (GLM p values < 0.001, 27 d.f.). Lick vs. PRP LDA (including peak and trough) did not differ between naïve and proficient states (GLM p values > 0.05, 20 d.f.).

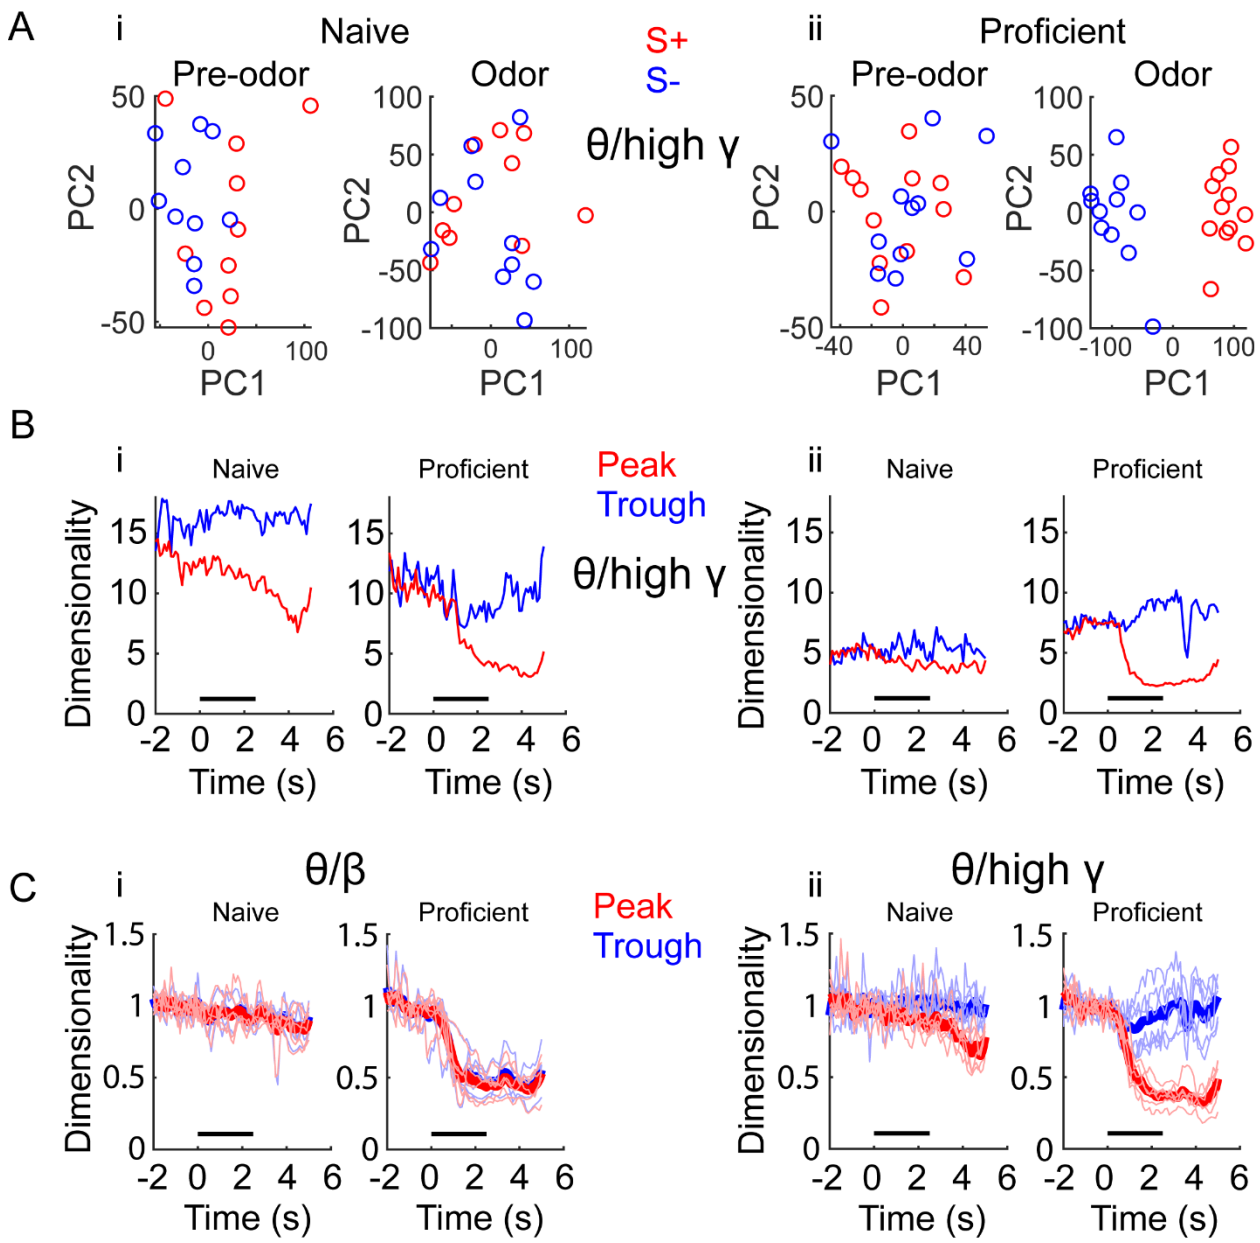

Figure 6. Dimensionality changes with learning and during the time course of the trial.

(A) Scatter plot showing the first two principal components for twenty trials of a PCA performed with peak theta/high gamma PRP recorded in mice learning to differentiate the odorants in the EAPAexp1 odorant pair calculated for pooled animal data. The input to the PCA was the peak PRP recorded in 16 electrodes. Principal components are

intermingled between S+ (red) and S- (blue) trials in the naïve state (left) in pre-odor and odor-delivery epochs, yet they separate during odor delivery upon task acquisition (right). i. Naïve animals. ii. Proficient animals.

(B) Examples of the within-trial time course for theta/high gamma dimensionality calculated for pooled animal peak (red) and trough (blue) PRP LFP for mice learning to differentiate the odorants in the APEBexp1 (i) and EAPAexp2 (ii) odorant pairs. The dimensionality of peak PRP in the presence of the odorant decreases when the animal becomes proficient.

(C) Within-trial peak (red) and trough (blue) PRP dimensionality time course for data pooled across mice and experiments 1 and 2, normalized to pre-trial dimensionality. (i) Theta/beta PRP dimensionality did not differ between peak and trough yet decreased significantly during odorant delivery with learning. For theta/beta PRP dimensionality the GLM yielded significant changes for time course and for the interaction of time course x learning ($p < 0.001$, 1696 d.f.), but there were no differences for peak vs. trough ($p>0.05$). (ii) Peak theta/high gamma PRP dimensionality decreased significantly during the odorant delivery epoch in proficient animals in both experiments. Peak PRP dimensionality was significantly smaller than trough. For theta/high gamma the GLM yielded significant changes for peak vs trough ($p < 0.001$), learning ($p < 0.01$) and for the interaction of peak vs. trough x learning x time course ($p < 0.001$, 1696 d.f.).

SUPPLEMENTAL INFORMATION

Experiment	Device implanted	Recording coordinates	Odor pair	Odor name	Concentration (v/v)	Number of mice
Exp1	Optetrodes	+4.25 mm AP, +0.4 mm ML, -0.53 mm DV	APEBexp1	Acetophenone	1%	14
				Ethyl benzoate	1%	
			EAPAexp1	Ethyl acetate	0.1% EA	11
				Propyl acetate	0.05% EA + 0.05% PA	
			IAMOexp1	Isoamyl acetate	1%	9
				Mineral oil	100%	
Exp2	Tetrodes	+4.28 mm AP, +0.5 mm ML, -1.0 mm DV	EAPAexp2	Ethyl acetate	1%	9
				Propyl acetate	1%	
			IAAPexp2	Isoamyl acetate	1%	6
				Acetophenone	1%	
			IAMOexp2	Isoamyl acetate	1%	4
				Mineral oil	100%	

Table S1. Odorant pairs and electrode locations

Odorant identities for each pair are listed along with concentrations represented as volume/volume dilutions in mineral oil. In the case of EAPAexp1, odorant A is 0.1% ethyl acetate and odorant B is a mixture of 0.05% ethyl acetate + 0.05% propyl acetate. The two recording locations (exp1 and exp2) are listed in millimeters with respect to bregma. The number of mice per odorant pair experiment (e.g. EAPAexp2) is also listed at the right.

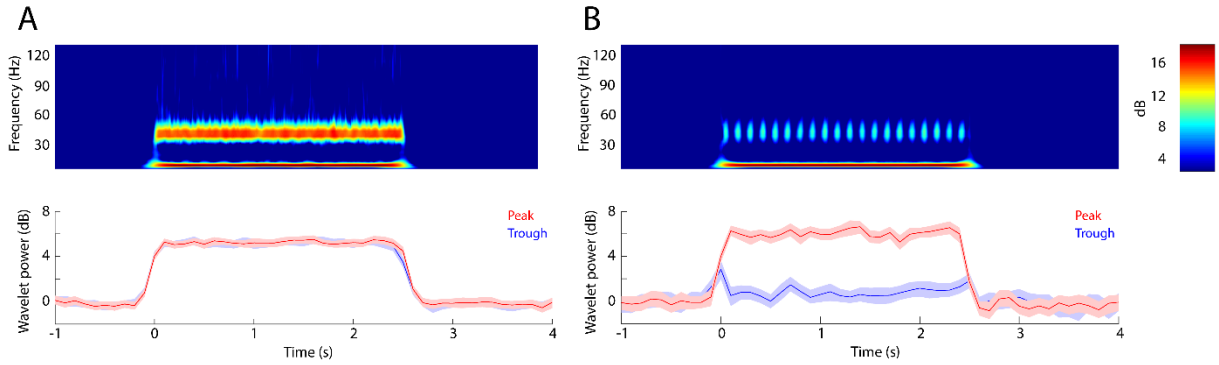

Figure S1. Simulation: PRP wavelet analysis for theta-gamma waves with/without PAC.

LFP data were generated by sinusoids supplemented with gaussian random noise. The data was analyzed for PRP as described in the text.

(A) Example with no theta/gamma PAC generated by input of two continuous cosine waves of 8 Hz and 40 Hz. Top: Morlet wavelet spectrogram shows a continuous increase in power at the two frequencies. Bottom: peak (red) and trough (blue) PRP do not differ. MI = 0.0004.

(B) Example of theta/gamma PAC generated with input of 40 Hz gamma bursts placed at the 180° phase of an 8 Hz theta sine wave (the bursts take place at the peak of PAC). Top: Morlet wavelet spectrogram shows discrete increases of gamma power and a continuous increase in power at 8 Hz. Bottom: peak (red) PRP shows a large increase in power while trough PRP (blue) does not increase. MI = 0.01.

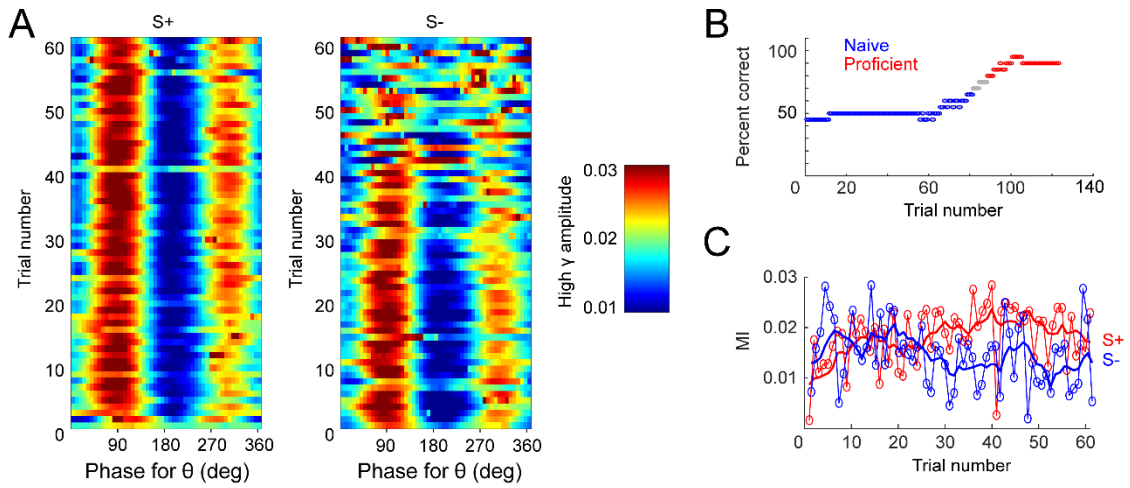

Figure S2. Example of PAC changes over learning in one session.

(A) PAC for all S+ trials (left) and S- trials (middle), with high γ amplitude color bar quantified at right.

(B) Behavioral performance. Mouse starts the session performing at chance, then learns the task, acquiring proficiency by the end. Blue: naïve, performance ≤ 65%. Red: proficient, performance ≥ 80%.

(C) PAC quantified as mean MI for S+ (red) and S- (blue). With learning, PAC increases for S+ and decreases for S-.

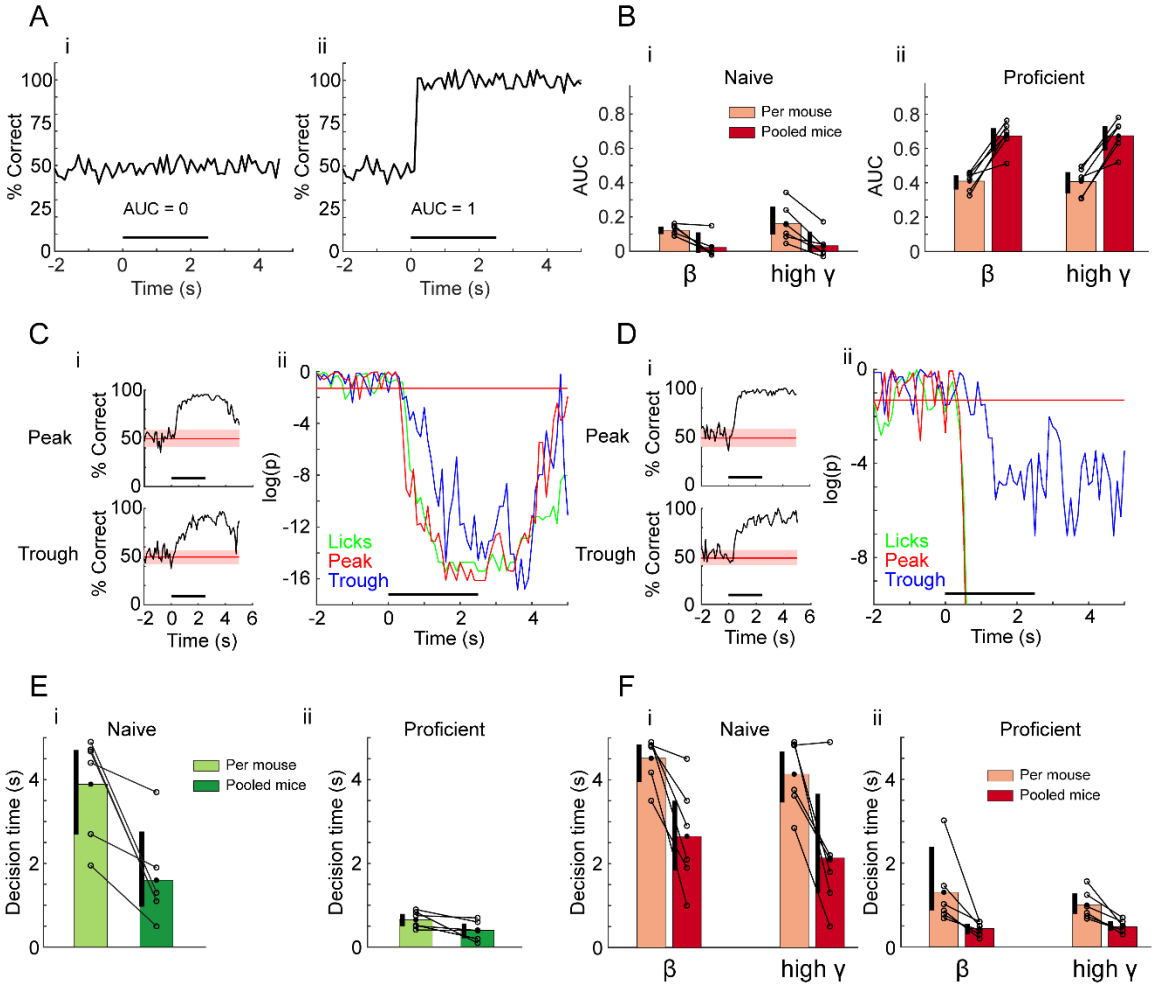

Figure S3. Decision times are faster for peak PRP LFP LDA computed when data are pooled across mice.

(A) Simulated data illustrating the time course for decoding performance for LDA for discrimination of odor identity with: i, The area under the curve (AUC) = 0, LDA was not able to identify the stimuli, and ii, AUC = 1, perfect identification of the stimulus after odor application. AUC was calculated as the area under the curve during odorant stimulation normalized to 1.

(B) AUC estimates of PRP LDA proficiency for theta/beta and theta/high gamma. The data compare AUC values obtained for a peak PRP LDA estimated separately for each mouse using LFP recordings from all electrodes (per mouse, light red) vs. AUC

estimates for peak PRP LDA obtained using LFP recordings from all electrodes from all mice (pooled mice, red). Data are presented for all odorant pairs. i: Naïve mice, ii: proficient mice. There is a significant difference between per-mouse and pooled mouse estimates (GLM p values < 0.001, 27 d.f.), and there is no statistically significant difference between theta/beta and theta/high gamma (GLM p values > 0.05, 27 d.f.).

(C) i. Decoding performance time course for LDA of PRP recorded with all 16 electrodes for one mouse discriminating odorants in the EAPAexp2 odorant pair (per-mouse PRP LFP LDA). LDA decoding precision is shown for peak (i, top) and trough (i, bottom) PRP. ii. p value for the ranksum test yielding information on the time-course for decision-making using licks for S+ vs. S- (green) or LDA prediction of odorant identity with peak (red) or trough (blue) PRP.

(D) i. Decoding performance time course for LDA of PRP estimated with all electrodes recording the LFP in nine mice discriminating odorants in the EAPAexp2 odorant pair (pooled animal PRP LFP LDA). ii. Decision making, measured as a sustained drop of the p value below 0.05, occurs at the same time for peak PRP and licks, while trough PRP is significantly slower.

(E) Comparison of lick decision times between per mouse (light green) vs. pooled mouse (dark green) p value estimates. GLM p value for the difference in lick decision times between per-mouse vs. pooled GLM < 0.01, 20 d.f.

(F) Decision times for peak PRP decoding of stimulus identity for the per mouse (light red) vs. pooled mouse (red) LDA for naïve (i) vs. proficient (ii) mice. Vertical black bars represent the 95% bootstrap confidence interval of the data in each condition; they are shifted to the left to better display the mean and individual odorant pair values. For both

theta/beta and theta/high gamma decoding of stimulus identity via peak PRP is faster using data pooled across mice than the per mouse estimate (GLM p values < 0.001, 20 d.f.). There is no difference in decision making times between theta/beta and theta/high gamma (GLM p values > 0.05, 20 d.f.).

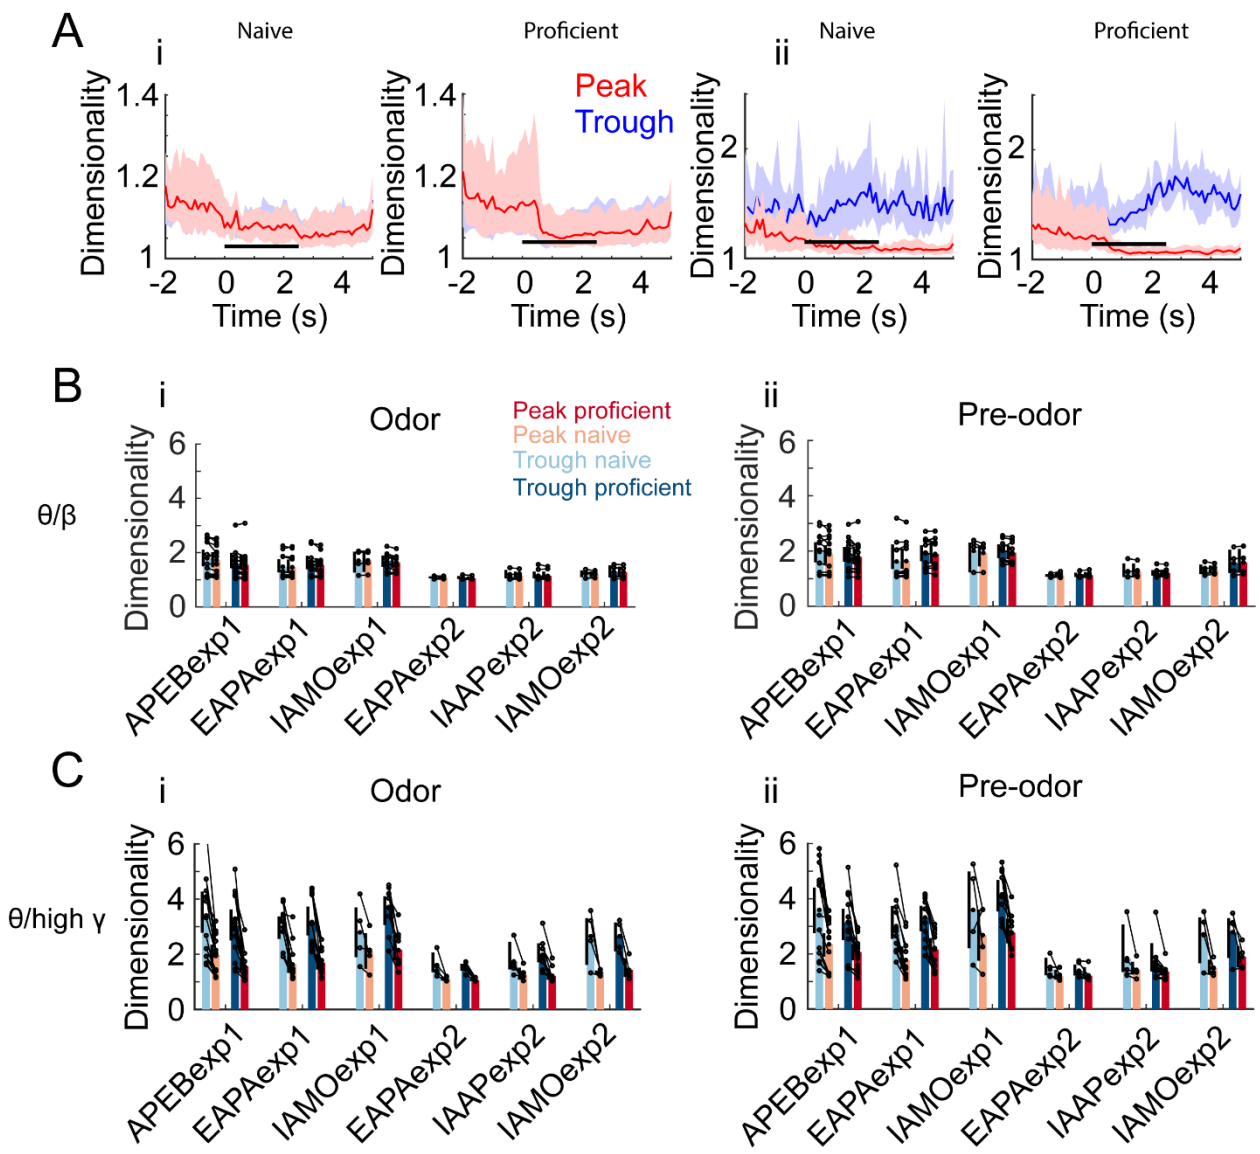

Figure S4. PRP dimensionality estimated per mouse.

(A) Example of the per mouse dimensionality time course for mice learning to discriminate odorants for the odorant pair IAAPexp2. i: Theta/beta PRP in naïve (left) and proficient (right) mice. ii. Theta/high gamma PRP in naïve (left) and proficient (right) mice.

(B) Per mouse theta/beta PRP dimensionality summary for all six odorant pairs. i: During the odor delivery epoch, ii: before odorant delivery. GLM p value < 0.001 for location and > 0.05 for peak vs. trough, 379 d.f.

(C) Per mouse theta/high gamma PRP dimensionality summary for all six odorant pairs. i: During the odor delivery epoch, ii: before odorant delivery. GLM p value < 0.001 for peak vs. trough and location, 379 d.f.