nature portfolio

Peer Review File

Fast updating feedback from piriform cortex to the olfactory bulb relays multimodal identity and reward contingency signals during rule-reversal

Corresponding Author: Dr Dinu Albeanu

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
This study describes a unique and important finding: the feedback fibers from the anterior piriform cortex (aPCx) to the olfactory bulb carry reward contingency signals, where the modulation can be multimodal. They show this with an elegant behavioral paradigm where mice discriminate between odor vs. sound cues and where the reward contingency switches between blocks so that sometimes an odor predicts a reward, and in other blocks, a sound predicts a reward. By imaging from the aPCx terminals in the olfactory bulb, they show that rapid contingency changes modulate the cue-driven signals in these structures. Surprisingly, for feedback from an olfactory cortex, the modulation also occurred for the sound cues, which the authors confirm using an additional, non-olfactory (purely auditory) discrimination task. The results are rigorously tested and analyzed, and the data are high quality. The authors may further strengthen the study by addressing the following points.

Regarding the nature of odor vs. sound-driven signals conveyed: For the odor, as the authors note, there is an increase in responsivity as mice learn the task; For the sound, there seems to be a decrease – an opposite effect. This point is in addition to the observation that the modulation by olfactory cues seems generally much more robust. So, while the modulation is multimodal, there is a difference. A discussion clarifying the nature of multimodal modulation would be beneficial.

The sound discrimination task is a very nice control related to the above point. A more direct comparison of sound-driven responses across the two tasks (e.g., the magnitude of modulation, etc.) would be helpful. This is especially the case because the sound trials in the odor-sound discrimination can be interpreted as a trial without an odor.

Given the discussion about regional differences within the piriform cortex (e.g., lines 513 – 515), it would be good if the authors could present an image showing the AAV injection site if they have it. Such an image would allow the readers to judge the location and extent of infections.

Related to the above point, Supplementary Figure 3 shows a range of effect sizes (proportions of ROIs that are modulated). Out of curiosity, do the variabilities relate to the injection sites or the imaged depths? Related to this, the authors write that they imaged at 200-300 um depths from the brain surface. This depth range sounds too superficial for GCL imaging. Could the authors please clarify?

Optogenetic silencing of the presynaptic terminals is interesting and worth expanding on. It is surprising but exciting that behavioral performance could be impaired this way, especially with only 2.5 mW of irradiation. However, this result section needs to be better described. The method for this experiment was missing, too. How exactly was the light presented? Did it overlap with the cues or the delay period? Please describe the conditions used.

Reviewer #2

(Remarks to the Author)
Unfortunately, we cannot recommend this manuscript for publication at this time. While we commend the authors for their admirable efforts to train mice on this difficult rule-reversal task and acknowledge the valuable insights into feedback processing within the olfactory bulb presented in the data, we have major reservations about the primary conclusions of the paper. The central claim that the piriform cortex conveys multimodal identity and reward-contingency feedback to the olfactory bulb lacks robust support in the provided data. Our reservations are grounded in three key reasons, which we will outline in the following paragraphs:

First, previously published work suggests that the piriform cortex does not strongly encode reward information. For example, Miura et al. (2012) found that very few neurons in anterior PCx (<5%) exhibit choice or outcome selectivity within the first sniff of an odor stimulus. Similarly, Millman and Murthy (2020) reported that posterior PCx only weakly encodes reward contingency in an odor reward categorization task (<5% of responsive neurons reward selective). Moreover, Wang et al. (2020) demonstrated that PCx odor representations are largely unchanged throughout learning in an appetitive odor discrimination task, while odor-value representations emerge in OFC and mPFC. The results of these studies are in stark contrast to the strong contingency representations observed in the manuscript's axonal imaging data. Notably, a few classical studies have reported a degree of reward representation in piriform, and neurons providing feedback to the bulb may receive more reward-related information than the general population. However, without direct evidence, it is difficult to conceive of how PCx could provide a rapid and strong update of reward-contingency to the bulb. Similarly, almost all studies of odor responses in the piriform cortex report piriform neurons are relatively odor-selective and that only ~10% of piriform neurons respond to a given odor. In this study, the authors report that >50% of axon terminals were odor-responsive.
Second, given how drastically different the results reported here are from the extant literature, we are concerned that the observed in the present study, which is interpreted as the encoding of reward contingency observed in feedback from the piriform cortex may actually arise due to neuromodulatory effects on presynaptic terminals of the piriform cortex axons within the bulb, as opposed to signaling piriform cortex activity itself. The OB receives neuromodulatory inputs from many different brain regions involved in reward processing including cholinergic input from the horizontal limb of the Diagonal Band of Broca and noradrenergic input from the locus coeruleus. The experiments within this manuscript do not address the possibility that neuromodulation may be required to gate the observed reward signaling. While the experiments investigating the effect of suppressing feedback axons in figure 4 f-g suggest that intact axonal activity is required for accurate task performance, they do not exclude the possibility that reward information within the axons does not originate in piriform.
Third, the authors make the claim that piriform feedback encodes reward-contingency signals across multiple modalities. In support of this claim, they demonstrate that in animals trained to perform a sound-sound discrimination task, feedback boutons in the OB exhibit different responses to the rewarded and unrewarded stimulus. However, there is no analysis to compare hit and false alarm trials for each stimulus, making it possible that the observed responses are simply the result of motor preparatory activity.
Given these reservations, we believe that major revisions to the current manuscript would be necessary to support the central claims. These include either 1) recordings in piriform, or specifically in bulb-projecting neurons within piriform during task performance to determine whether reward-contingency signals are present in this region or 2) pharmacological manipulations to rule out the influence of presynaptic neuromodulatory signaling on reward-contingency responses in the OB. Without these important and difficult controls, we recommend that in future submissions, the authors reframe their results to specifically focus on feedback information available in the olfactory bulb without specifically attributing these computations to the piriform cortex.

Reviewer #3

(Remarks to the Author)
Summary

In this manuscript the authors present evidence that piriform cortical feedback to the olfactory bulb carry multimodal information. Interestingly, they show that response properties of individual boutons can change across blocks within a session depending on reward contingencies. This is novel and significant. However, while the main observation and data showing cortical feedback to OB is multi-modular is convincing, there are some specific claims about neural responses were not fully supported by experimental design and analysis. In particular, the behavioral task design makes many interpretation of the data confounded, and this is not acknowledged. The authors should address these concerns about experimental design, framing, and interpretation of the data before the manuscript is accepted.

Major concerns:
1. Block structure in behavioral task is predictable. What is the exact structure of the block lengths and how was this decided? If the hazard rate function of block length is not flat, (that is, if mice can predict block length), mice would be able to learn the trial structure of block switches, and solve the task uni-modally. For example, task mice learn can be: odor-> go for 45 trial; and then rule switch to odor -> non-go for 45 trials. If mice only used odor and not sound at all, they can still solve each block at 75% chance level (100% correct for all odor trials, and 50% correct for all sound trials). Similarly, mice could also only use the sound to solve the task. This is a particularly likely strategy given that trial structure is also temporally predictable: ITI intervals and cue lengths (ITI hazard is not flat) makes it such that mice can know the temporal likelihood of when sensory cue will occur. This means that if mice ignored sounds completely, they can successfully use the strategy of "lick at time ~X if I don't smell an odor" for the tone trials of the "sound-go" blocks and achieve high performance. In fact, given the variety of behavioral and neural responses that authors report (for example, results in Sup Fig 8), it is likely that different mice are using different strategies to solve the task. The authors should present evidence to support their claim that mice are using different sensory modalities to solve this task. In lieu of behavioral evidence for this, the most parsimonious explanation for how mice would be solving the task. In this case, there are a few issues:

a. Authors do not know which sensory modality mice are using to solve the task for any given trial. Therefore, interpretation of behavioral relevance of neural signal is inherently limited and all claims in manuscript about sensory-domain specific learning should reflect this. For example, Page 20: “As learning of the sound-reward associations progressed, …” sound be rephrased.
b. Even if mice are solving the task using sound only in the “sound go” blocks, and odors only in the “odor-go” blocks, because the block length is predictable, expectation (sensory prediction, reward prediction) is systematically manipulated relative to block switch. It is known that OB activity is modulated by association, memory, and attention, behavior or neural responses that depend on analysis aligned to block is confounded by this. Therefore, the claim that neural and behavioral differences across block switches are due to the sensory domain differences is not supported. This means that the main conclusion/claim of this study needs to be carefully framed, and perhaps reinterpreted completely.

2. Given the confounds presented in 1), the results in Figure 3 would be a lot more completing if they were compared to neural recordings from awake passive mice over days of exposure to predictable events (same trial structure and block structure without rewards). Which part of the reported results here would be different? Figure 3f-h shows that experience does matter, but given that we do not know what the mice are learning, these panels must be interpreted carefully. Given that Varga and Wesson (2013) showed that piriform cortex responds tones in anesthetized conditions. In my view, the authors should make clear the expansion of their results beyond this study (aka tones evoke responses from piriform cortex neurons). Can the authors compare results between awake passive naïve mice vs awake passive experienced mice, and contrast these differences with results in their current manuscript (even if statements about domain-specific learning cannot be made). For example, highlighting the difference between naïve and experienced mice.

3. Go/NoGo task introduces asymmetries and confounds. There is not a clear axis in the behavioral paradigm that helps to discern activity from reward/motor/sensory cues clearly. All the distinction come from the quantification of the imaging data. In this regard, authors should include a more careful comparison between EGFP signals and GCaMP signals to clearly state what is the expected noise level, and what is clear signal.

4. It was not clear to me if the authors are able to track the same boutons over days/sessions. The manuscript would be significantly stronger if analyzes included properties of a subset of individual boutons over sessions and days.

Minor concerns:
1. In the last sentence of the abstract, it is not clear what “identity” refers to.
2. Fig 1d is very surprising. It seems very unlikely that there would be no drop in session performance at the first day of rule reversal. How do the authors explain this? There are only 6 reversals in average in one session. Perhaps a session example of a rule reversal training day might give a bit more clarity on what happens? Looking at the Methods section that details the training, it looks like there are a lot of details that are overlooked in the main body of the text. For instance, the first day of reversals animals were given automatic water and other tricks to help them achieve performance. The authors should either make this transparent or tone down the claim regarding fast learning made in Figure 1 and beyond.
3. In general, the paper is missing statistical comparisons for several of the claims they make. Information such as number of trials per session, reversals, accuracy, etc. should also be including the SD of the values.
4. Is the sampling of the odor locked in with the breathing? Does it take the same time to process both information?
5. Comparison between the task with delay and the no delay. It is unclear in which way the authors expected to see differences between delay and no delay task. They mention how the delay can eliminate possible confounds but that is not directly observed in any analysis or metric that they access.
6. Page 6. “Reaction time” for delay version of the task in not meaningful, given the cue length and enforced delay. Its misleading to use that term.
7. Page 7. Licking rate really overlapping in the figure but then substantially different in the text. Include statistics and/or change the axis size to be able to perceive the subtle differences.
8. Supp 1c. Don’t understand the switch frequency plot.
9. Page 7. Where is the analysis that Supp. Fig. 3e points at?
10. Page 9. Is this the best way to look at responsiveness?
11. In line 193, Pg. 9 authors state that "across the cue-modulated feedback boutons … 10.1 ± 2.7 % were selectively tuned to the tone". In line 207, Pg. 10 they say: "only a small fraction of EGFP boutons passed our responsiveness criterion … 15.5 ± 7.5%; ". If we take the fraction of responsive EGFP boutons as chance level, this means that tone-modulated boutons are below chance level? Are these two statistically different?
12. Page 11. Isn’t the fact that the boutons respond only in reward trials for sound indicates that it’s mostly either reward/movement related?
13. Page 11, line 238. Where can these observation be found?

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors have addressed all my concerns.

One thing is that some of the preprints that are cited have since been published in peer-reviewed journals. I have listed these citations below.

Congratulations to the authors on the beautiful study.

31. Young, H., Belbut, B., Baeta, M. & Petreanu, L. Laminar-specific cortico-cortical loops in mouse visual cortex. bioRxiv 773085 (2019) doi:10.1101/773085.
-> Now published at eLife (https://doi.org/10.7554/eLife.59551)

107. Cole, N., Harvey, M., Myers-Joseph, D., Gilra, A. & Khan, A. G. Prediction error signals in anterior cingulate cortex drive task-switching. 2022.11.27.518096 Preprint at https://doi.org/10.1101/2022.11.27.518096 (2022).
-> Now published at Nature Communications (https://doi.org/10.1038/s41467-024-51368-9)

117. Lindeman, S., Fu, X., Reinert, J. K. & Fukunaga, I. Reward contingency modulates olfactory bulb output via pathway-dependent peri-somatic inhibition. 2023.08.17.553686 Preprint at https://doi.org/10.1101/2023.08.17.553686 (2023).
-> Now published at PLOS Biology (https://doi.org/10.1371/journal.pbio.3002536)

Reviewer #2

(Remarks to the Author)
My major concern was that the responses the authors reported were inconsistent with previous studies and that this might be attributable to neuromodulatory signals on the axons of piriform cortex cells that project back to the bulb. The authors argue against this based on spatial selectivity. Specifically, they argue that if the signals observed in the axon terminals are neuromodulatory, the signals recorded in nearby terminals should be correlated. However, if the signals they record are due to the activity of the neurons in the piriform cortex, then responses on boutons from the same axon will be highly correlated while responses for similarly proximal boutons on different axons should be uncorrelated.

I agree with this argument; the data the authors provide to support it are compelling. Most of my other major concerns (e.g., why responses are so different from many previous studies) hinged on the concern that the authors weren't really recording cortical responses, and so I now consider my previous concern in this regard to be resolved. I think the authors have done a commendable job addressing the concerns of all three reviewers and I have no major reservations about this manuscript. Indeed, the Federman et al. paper that was published in the meantime, and the recent paper from Kehl et al. just published in Nature, show that the piriform cortex is much more than a simple primary sensory cortex (of course it is, that is a straw man), but that it exhibits both contingency and multimodal signals. I am still a little surprised by how dense their odor responses are, but I think the authors have done everything correctly. As such, this is an important paper that should be of interest to a large audience.

Reviewer #3

(Remarks to the Author)
The authors have addressed the issues I raised during my initial review. In my opinion the manuscript is significantly strengthened. Particularly considering the recent publication from Federman et al., Nature Communications, 2024, I believe this manuscript is a timely and significant advance for the field. In particular, the additional explanation about the ITI durations is indeed helpful for ruling out a timing-based strategy that can explain behavior. Similarly, I am happy with the explanation provided about how block length estimation is also not a strong confound given the behavioral observations about failures.

We appreciate the authors for conducting extra experiments to answer this question. When comparing figure Reviewer 9c and Figure 3h, it is striking that the peak relative classification performance is already around 0.3 for the naïve session in the habituation only experiments but between 0 and 0.1 for the Auditory only experiments. Why is this the case? I would expect a performance around 0 (chance level). The authors may consider clarifying this in their final manuscript.

Response to reviewers

Reviewer #1 (Remarks to the Author):

This study describes a unique and important finding: the feedback fibers from the anterior piriform cortex (aPCx) to the olfactory bulb carry reward contingency signals, where the modulation can be multimodal. They show this with an elegant behavioral paradigm where mice discriminate between odor vs. sound cues and where the reward contingency switches between blocks so that sometimes an odor predicts a reward, and in other blocks, a sound predicts a reward. By imaging from the aPCx terminals in the olfactory bulb, they show that rapid contingency changes modulate the cue-driven signals in these structures. Surprisingly, for feedback from an olfactory cortex, the modulation also occurred for the sound cues, which the authors confirm using an additional, non-olfactory (purely auditory) discrimination task. The results are rigorously tested and analyzed, and the data are high quality. The authors may further strengthen the study by addressing the following points.

We thank the Reviewer for their assessment of our findings and would like to address the issues raised.

Regarding the nature of odor vs. sound-driven signals conveyed: For the odor, as the authors note, there is an increase in responsivity as mice learn the task; For the sound, there seems to be a decrease – an opposite effect. This point is in addition to the observation that the modulation by olfactory cues seems generally much more robust. So, while the modulation is multimodal, there is a difference. A discussion clarifying the nature of multimodal modulation would be beneficial.
The sound discrimination task is a very nice control related to the above point. A more direct comparison of sound-driven responses across the two tasks (e.g., the magnitude of modulation, etc.) would be helpful. This is especially the case because the sound trials in the odor-sound discrimination can be interpreted as a trial without an odor.

We agree. While the modulation we observe in cortical bulbar feedback axons is multimodal, the exact dynamics differ across odor and sound stimuli. In expert mice performing the reversal task, there is robust modulation of both odor and sound responses as a function of reward contingency, albeit the sound responses are in general weaker in amplitude than odor responses. On average however, odor responses were stronger when the odor was rewarded (**Fig. 2f**). Similarly, sound responses were larger in sound-rewarded vs. sound-not rewarded blocks of trials (**Fig.2f**; **Supplementary Fig. 9f**). We did not explicitly compare the odor responses during learning of the reversal task. We note that compared to previous work, including reports from our lab (Otazu et al. 2015; Boyd et al., 2015), the odor responses in expert mice performing the rule-reversal task were denser and stronger in amplitude compared to average responses in naïve mice.

To address the point raised by the Reviewer, for the *sound responses*, we include additional analysis (**Supplementary Figs. 15a**,**b**,**e-j**, also shown below as **Reviewer Fig. 1**) to compare the response strength across days in the auditory-only Go/No-Go task and expert mice engaged in the odor/sound reversal task. Across the population, the magnitude of the sound responses was similar in mice performing the odor/sound rule-reversal task than the auditory-only Go/No-Go expert mice (6 days vs. several weeks, Wilcoxon rank-sum test, enhanced and suppressed: $p>0.9$, **Reviewer Fig. 1a**, **b**, **d**). In the auditory-only Go/No-Go task, in expert mice, rewarded sound responses were stronger in amplitude than responses to the non-rewarded sound (Wilcoxon rank-sum test, enhanced: $p<0.001$; suppressed: $p<0.001$, **Reviewer Fig. 1b**, **e**). In contrast, in naïve sessions, responses to the Go and No-Go sound cues were more similar (Wilcoxon rank-sum test, enhanced and suppressed: $p>0.1$).

In response to R3's comments, we performed additional experiments (n=3 mice) in which we imaged bouton responses to the same two sound cues across 6 consecutive days, without providing water rewards (habituation controls, **Reviewer Fig. 1c**, **f**; also included as **Supplementary Figs. 15d-f**). Compared to the task-engaged mice, in these habituation experiments, the feedback bouton activity was sparser and their response strength substantially lower (Wilcoxon rank-sum test, enhanced and suppressed: $p<0.0001$). In these habituation experiments, we did not observe differential modulation of the cortical feedback responses to the two sound cues (A vs. B) in a systematic manner across sessions (days 1-to-6). The strength of responses to each cue did not follow a systematic trend across sessions. The performance of classifiers was stable across sessions, in contrast to the steady increase observed in the decoding performance during learning in mice engaged in the auditory-

only Go/No-Go task. Interestingly, across fields of view, we observed differences in feedback bouton response amplitude as a function of sound stimulus identity (A vs. B) which were present throughout the imaging sessions (**Supplementary Fig. 15 d-f**).

We conclude that sound cues trigger sparse responses in the cortical bulbar feedback axons in naïve mice, whose strength and specificity are shaped (augmented) by reward contingency, but also potentially depend on other stimulus features. We discuss these new results and analyses in the revised manuscript (rows **464-493**).

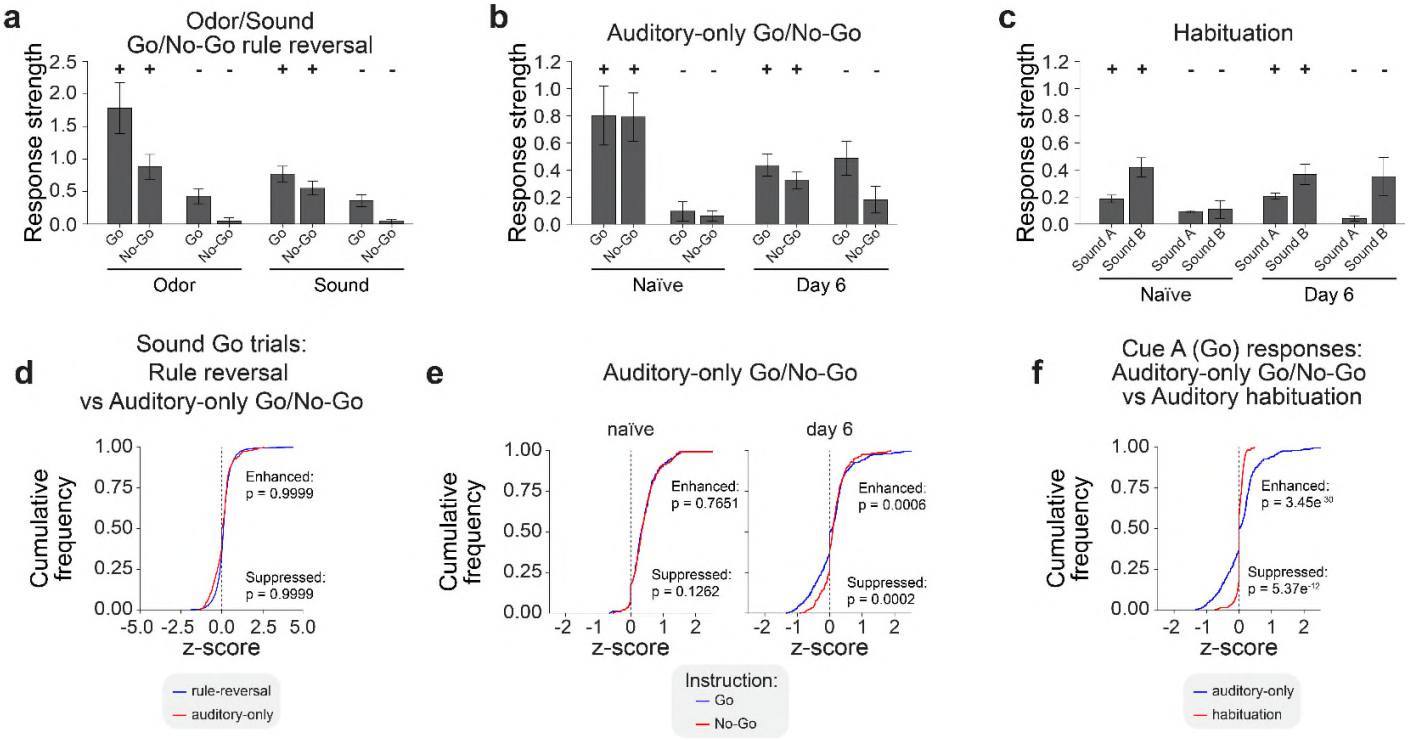

Reviewer Figure 1. *Comparison of response strength differences between odor and sound responses in the reversal task (**a**), Two-sounds Go/No-Go task (**b**), and Two-sounds habituation sessions (**c**). We quantified response strength using the average z-score response during the delay period for **a** and **b**, or an equivalent period for **c**, sampling in each field of view the top 20 response amplitude boutons. +/- mark enhanced and respectively suppressed bouton responses. **d**. Cumulative distributions of response strength for expert mice engaged in the odor/sound rule-reversal task (blue) and auditory-only Go/No-Go task (day 6, red). **e**. Cumulative distributions of response strength (peak amplitude during delay period) for naïve (Left) and task-engaged (day 6, Right) mice performing the auditory-only Go/No-Go task. Go responses are shown in blue; No-Go responses in red. **f**. Cumulative distributions of response strength to Sound A for day 6 in mice performing the auditory-only Go/No-Go task (Go cue, blue) or passively exposed to the same sound cue (habituation, red).*

In terms of interpreting the sound trials in the odor-sound discrimination task as trials without odor, our behavioral data suggests otherwise. Expert mice precisely time their licks in a short window after the cue onset for both the odor and the sound trials (**Fig. 1e**; **Supplementary Fig. 1b**). As we use a flat hazard rate for the inter-trial intervals (ITI), expert mice cannot accurately predict when the sensory cues will occur. In conjunction with variable cue-to-cue intervals for consecutive trials based on each trial's outcome (hits vs. correct rejections vs. false alarms vs. misses), and the lack of overt signals to mark the start of a trial, this strategy ensures that mice could not predict the onset of the sensory cues (rows **825-830**; **891-902**), and suggests that mice detect and use both cues (**Supplementary. Figs. 1f, g**). Also, on average, expert mice had high behavioral performance for both the odor and sound trials, independent of the block type (Odor-Go, as well as Sound-Go blocks), with slight biases as a function of cue and mouse identity (**Supplementary Figs. 1d**,**e**; **10d**,**e**). This rules out the potential strategy of solving the task (receiving rewards) by relying only on one of the sensory cues, and engaging in random responses to the other cue. In the revised manuscript, we clarify these points (rows **127-134**; **825-830**; **891-902**).

Given the discussion about regional differences within the piriform cortex (e.g., lines 513 – 515), it would be good if the authors could present an image showing the AAV injection site if they have it. Such an image would allow the readers to judge the location and extent of infections.

We agree and thank the Reviewer for the suggestion. In the revised manuscript, we included a new supplementary figure showing the injection sites (**Supplementary Fig. 2**).

Related to the above point, Supplementary Figure 3 shows a range of effect sizes (proportions of ROIs that are modulated). Out of curiosity, do the variabilities relate to the injection sites or the imaged depths? Related to this, the authors write that they imaged at 200-300 um depths from the brain surface. This depth range sounds too superficial for GCL imaging. Could the authors please clarify?

We apologize for the vague statement. We have imaged between 200-300 µm from the bulb surface, just below the mitral cell layer. Unfortunately, we do not have systematic records of the exact depth of the imaging optical plane for each session. As such, the variability in response across fields of view may indeed be a function of multiple variables. These include the exact location of the injection site (in each animal, we aim to inject in the same 3 sets of AP-ML-DV coordinates, as mentioned in the Methods, rows **795-798**), as well as differences in the statistics of odor responses in the input glomeruli and local bulbar circuits around the imaged fields of view, across different locations in the bulb. In the revised manuscript, we clarified these points in the Methods accordingly (rows **190-192**; **949-953**).

Optogenetic silencing of the presynaptic terminals is interesting and worth expanding on. It is surprising but exciting that behavioral performance could be impaired this way, especially with only 2.5 mW of irradiation. However, this result section needs to be better described. The method for this experiment was missing, too. How exactly was the light presented? Did it overlap with the cues or the delay period? Please describe the conditions used.

We apologize for this oversight. Light stimulation started 500 ms before the onset of the cue and lasted until the end of the report period. This and other related info were present in the Methods (Analysis section, see below). We inserted a clearer description in the revised Methods, as a separate section on optogenetic suppression of the cortical feedback. We include timing of the stimulation protocols, light intensity, and comparisons to the EGFP controls (rows **955-971**).

‘ … Mice were bilaterally injected at the same aPCx coordinates as for GCaMP expression with AAV to express the inhibitory opsin Jaws. On the same day, mice were bilaterally implanted with cannulas loaded with 200 µm diameter optic fibers (Doric: MFC_200/230-0.48_2.0mm_ZF1.25(G)_FLT) on top of the olfactory bulb (coordinates: AP: +1.2 from inferior cerebral vein; ML: ± 1.2 mm from inter-frontal suture; DV: 0.0 mm OB surface). The space between the optic fiber and the edges of the skull craniotomy was filled with white petrolatum (Dynarex) and the optic fiber cannulas and a metallic head-bar were attached to the skull using a combination of dental cement (Metabond®), black dental acrylic resin (Ortho-Jet™) and cyanoacrylate glue (Krazy-Glue®). After training and reaching > 80% performance, mice were ready for optogenetic suppression experiments (0.25 probability of experiencing a trial with light stimulation). Optic fiber cannulas were connected to a branched dual patch cable (Doric, Cat # SBP(2)_200/230/900-0.48_1m_SMA-2xZF1.25) using ceramic sleeves (ThorLabs, Cat # ADAL1). Light stimulation was performed using a 590 nm LED (ThorLabs Cat # M590F3) calibrated to deliver 2.4 mW at the tip of the patch cable. Stimulation (5 ms, 30 Hz light pulses) was triggered 500 ms before the start of the cue period and continued until the end of the reporting period (2.85 s). Equivalent experiments were performed in control animals expressing EGFP instead of Jaws.’

Reviewer #2 (Remarks to the Author):

Unfortunately, we cannot recommend this manuscript for publication at this time. While we commend the authors for their admirable efforts to train mice on this difficult rule-reversal task and acknowledge the valuable insights into feedback processing within the olfactory bulb presented in the data, we have major reservations about the primary conclusions of the paper. The central claim that the piriform cortex conveys multimodal identity and reward-contingency feedback to the olfactory bulb lacks robust support in the provided data. Our reservations are grounded in three key reasons, which we will outline in the following paragraphs:

We thank the Reviewer for their comments, and would like to address the concerns raised.

First, previously published work suggests that the piriform cortex does not strongly encode reward information. For example, Miura et al. (2012) found that very few neurons in anterior PCx (<5%) exhibit choice or outcome selectivity within the first sniff of an odor stimulus. Similarly, Millman and Murthy (2020) reported that posterior PCx only weakly encodes reward contingency in an odor reward categorization task (<5% of responsive neurons reward selective). Moreover, Wang et al. (2020) demonstrated that PCx odor representations are largely unchanged throughout learning in an appetitive odor discrimination task, while odor-value representations emerge in OFC and mPFC. The results of these studies are in stark contrast to the strong contingency representations observed in the manuscript's axonal imaging data. Notably, a few classical studies have reported a degree of reward representation in piriform, and neurons providing feedback to the bulb may receive more reward-related information than the general population. However, without direct evidence, it is difficult to conceive of how PCx could provide a rapid and strong update of reward-contingency to the bulb. Similarly, almost all studies of odor responses in the piriform cortex report piriform neurons are relatively odor-selective and that only ~10% of piriform neurons respond to a given odor. In this study, the authors report that >50% of axon terminals were odor-responsive.

Indeed, previous studies provide contradicting pictures of the existence of reward contingency signals in the anterior piriform cortex. Some studies report a lack of reward contingency signals (*Miura, 2012; Millman and Murthy, 2020, Wang et al., 2020*), while others do report some degree of reward representations in the anterior piriform (*Roesch et al., 2007; Schoonover et al., 2021*). In addition, a recent study reports stimulus reward information within the anterior piriform cortex in mice engaged in a virtual reality task (*Federman et al., Nature Communications, 2024*). In our study, we find strong reward contingency signals in the cortical feedback fibers. As the Reviewer points out, this may be because bulb projecting piriform neurons are enriched in representing reward-related information compared to the general population of aPCx neurons. Other possible explanations include specific differences in the behavioral tasks used to assess the reward signals. In particular, our task differs from previous work in that it specifically requires that expert animals repeatedly switch fast back and forth between different rules of engagement within the same behavioral session.

With regard to the overall responsiveness of aPCx to odors, indeed previous studies have reported sparse responses (~10%) of piriform neurons to a given odor in naïve animals. In naïve animals, work from several groups including ours (*Boyd et al, 2015; Otazu et al., 2015*) also found sparse (< 10 %) odor responses in the cortical bulbar feedback. However, studies involving task-engaged animals reported denser responses in the anterior piriform cortex (*Schoonover et al. 2021*). In addition, *Wu & Komiyama, 2020* showed that responses of cortical bulbar feedback fibers become stronger and denser with task learning. Our observation of high responsiveness in expert mice is in agreement with these studies. In the revised manuscript, we further expand on these points, and highlight the potential causes for differences between our study and previous work (rows **620-622**).

The authors make the claim that piriform feedback encodes reward-contingency signals across multiple modalities. In support of this claim, they demonstrate that in animals trained to perform a sound-sound discrimination task, feedback boutons in the OB exhibit different responses to the rewarded and unrewarded stimulus. However, there is no analysis to compare hit and false alarm trials for each stimulus, making it possible that the observed responses are simply the result of motor preparatory activity.

We thank the Reviewer for the suggestion to compare directly the hit and false alarm trials in the auditory-only Go/No-Go task. We share the Reviewer's concern that reward-contingency signals may be hard to untangle from motor preparatory activity. To address this very issue, in the odor/sound rule reversal task, we compared Hit and False Alarm trials for both odors and sounds independently, and found that the ensemble feedback responses during both the cue and delay periods were different for Hits vs. False alarms, despite the animal licking in both cases (**Fig. 3a-c** in the submitted manuscript).

"… Across fields of view, cortical feedback response amplitude during the hit and false alarm trials was generally higher than for correct rejection trials (Supplementary Figs. 7e,f). However, we also observed differences in the responses of individual boutons when performing pairwise comparisons between trials of different contingencies in which mice licked the reward port (Fig. 3b; Odor hits vs. false alarms, 65.0 ± 7.5 %; Odor vs. Sound hits, 60.1 ± 11.9 %, Odor vs. Sound false alarms, 42.7 ± 12.1%; Sound Hits vs. False Alarms, Fig. 3b). Since in all these

cases, mice subsequently licked the reward port, changes in the response of individual boutons to same stimulus across contingencies cannot be solely attributed to motion artifacts and/or preparatory motor activity…"

To further address the Reviewer's concerns, we also performed additional analysis and compared the Hit and False Alarm trials specifically in the auditory-only Go/No-Go task. Similarly, we found that the ensemble feedback responses during both the cue and delay periods were different for Hit vs. False alarm trials, despite the animal licking in both cases (**Reviewer Fig. 2**, also included in the revised manuscript, **Supplementary Fig. 15c,** rows **478-482**). As such, the observed responses cannot be explained as simply the results of motor preparatory activity.

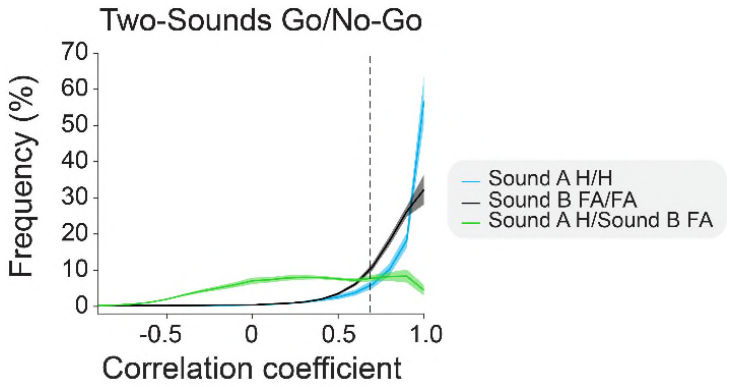

Reviewer Figure 2. *Histogram of individual bouton response correlation values across trials as a function of trial behavioral contingency (Hits, H vs. false alarms, FA) in the two-sounds Go/No-Go task. Bouton responses were sampled between cue onset and the end of delay period (before licking) for each of the two sound cues (A and B). Shaded areas correspond to SEM; the stability analysis was performed focusing on trials where mice subsequently licked the reward spout (hits and false alarms). Note the differences in trial-to-trial response correlation distributions when comparing H/H and FA/FA trials and respectively H/FA trials.*

Given how drastically different the results reported here are from the extant literature, we are concerned that the observed in the present study, which is interpreted as the encoding of reward contingency observed in feedback from the piriform cortex may actually arise due to neuromodulatory effects on presynaptic terminals of the piriform cortex axons within the bulb, as opposed to signaling piriform cortex activity itself. The OB receives neuromodulatory inputs from many different brain regions involved in reward processing including cholinergic input from the horizontal limb of the Diagonal Band of Broca and noradrenergic input from the locus coeruleus. The experiments within this manuscript do not address the possibility that neuromodulation may be required to gate the observed reward signaling. While the experiments investigating the effect of suppressing feedback axons in figure 4 f-g suggest that intact axonal activity is required for accurate task performance, they do not exclude the possibility that reward information within the axons does not originate in piriform.

Given these reservations, we believe that major revisions to the current manuscript would be necessary to support the central claims. These include either 1) recordings in piriform, or specifically in bulb-projecting neurons within piriform during task performance to determine whether reward-contingency signals are present in this region or 2) pharmacological manipulations to rule out the influence of presynaptic neuromodulatory signaling on reward-contingency responses in the OB. Without these important and difficult controls, we recommend that in future submissions, the authors reframe their results to specifically focus on feedback information available in the olfactory bulb without specifically attributing these computations to the piriform cortex.

The Reviewer raises the possibility that while feedback axonal activity is required for accurate task performance (**Fig. 4g**), the reward information that we observe within the feedback axons may not originate in the piriform cortex. Rather, it may reflect neuromodulatory input acting on these cortical bulbar feedback axons locally, within the olfactory bulb. This is a valid concern given the significant neuromodulatory inputs to the olfactory bulb. A multitude of scenarios can be considered in terms of neuromodulation given its diffuse action and the diversity of neuromodulators, types of receptors, as well as the diversity of cell types neuromodulators act on within the olfactory bulb (*Kapoor et al., 2016; Petzold et al., 2009; Smith et al., 2015; Rothemel et al., 2014; Pressler et al., 2007; Castillo et al.,1999; Collins et al., 2023; Linster & Fontanini, 2014*). However, we think that the spatial statistics of our data are not consistent with known mechanisms of neuromodulatory action in the olfactory bulb.

We outline below our reasoning supported by new analyses (see **Reviewer Figs. 3-6,** also shown in the revised manuscript as new **Supplementary Fig. 12**).

If neuromodulatory mechanisms were at play, we expect that boutons within the same spatial vicinity would be similarly modulated, irrespective of whether they belong to the same axon or not. New analyses show that, while responses of bouton pairs identified as belonging to the same axon are highly correlated, the responses of bouton pairs from different axons in the same neighborhood (ranging from 3 -to-15 µm apart) are very diverse in terms of amplitude, shape, and polarity (enhanced and suppressed). This is the case both when considering responses to a given stimulus (odor or sound), and/or reward contingency (**Reviewer Figs. 3-6**). Further, upon rule reversals, responses of nearby boutons from different axons (<15 µm apart) change in an uncoordinated manner, while responses of boutons on the same axon switch in a correlated fashion across blocks (**Reviewer Figs. 3-6**).

In other words, we find that for pairs of equidistant boutons, response similarity is tightly correlated with the axonal identity of the boutons (whether they belong to the same feedback axon or not). To the best of our knowledge, there are no known neuromodulatory mechanisms that can reformat responses of boutons in an axon-specific manner over such small (micrometric) spatial scales, i.e., change the responses of boutons on the same axon in the same manner, but affect the responses of adjacent boutons belonging to different axons in a hugely diverse fashion. A more parsimonious explanation for this exquisite coordination only among the boutons belonging to a given axon is that they reflect the activity of the parent piriform neuron.

We think that pharmacological manipulations of neuromodulation on the cortical bulbar feedback would be hard to interpret due to the large diversity of neuromodulators and receptor types present on virtually all the cell types of the olfactory bulb, in addition to the cortical feedback neurons. As such, pharmacological manipulation of neuromodulatory action in the bulb would alter not only the cortical feedback, but also the feedforward inputs to the cortex. This, in turn, could potentially impact the dynamics of feedback responses, rendering the results hard to interpret.

We acknowledge that the cortical recordings suggested by the Reviewer would be very interesting. However, in the light of these new analyses, and of recent work reporting reward related signals in the anterior piriform cortex in a different task (*Federman et al., Nat. Communications, 2024*), we believe they belong to follow-up projects. Importantly, pursuing these experiments involves a new set of experimental settings e.g. (opto-tagging and chronic electrophysiology recordings from deep structures like the piriform cortex and/or using miniscopes/two photon approaches for deep tissue imaging to monitor activity of specifically tagged sets of neurons). We estimate these will take 1-2 years to implement and perform so as to reach robustly interpretable results. Indeed, we would like to focus on addressing these questions, record activity in the piriform cortex, and further analyze the logic of top-down inputs the cortex receives over the coming few years. We aim to assess how the activity of specific neurons in the anterior piriform cortex is shaped by rule reversals and to further understand to what degree such changes depend on: 1) cell type identity, and 2) top-down signals to the piriform from other cortical areas (OFC, mPFC, etc.). We discuss these directions in the revised manuscript (rows **637-668**).

Results of new analyses: We systematically compared the responses of boutons within a small neighborhood (<15 µm) to various features including stimuli (odor, sound), as well as reward contingency. Specifically, we identified boutons that visibly belonged to the same axon (*yellow arrows* - **Reviewer Fig. 3a**), as well as equidistant boutons in the vicinity that appeared to lie on other axons (*white arrow heads* - **Reviewer Fig. 3a**). Consistently, we found that boutons belonging to the same axon responded substantially more similarly to a given stimulus (*black traces,* **Reviewer Fig. 3b***)* than boutons on different axons within the same neighborhood (<15 µm, color traces, Odor response correlation, Avg±SEM: 0.62±0.12 vs. 0.26±0.15, **Reviewer Fig. 4a**; Sound response correlation, Avg±SEM: 0.50±0.11 vs. 0.22±0.15, $p<0.0001$ Wilcoxon rank-sum test, **Reviewer Fig. 4b**).

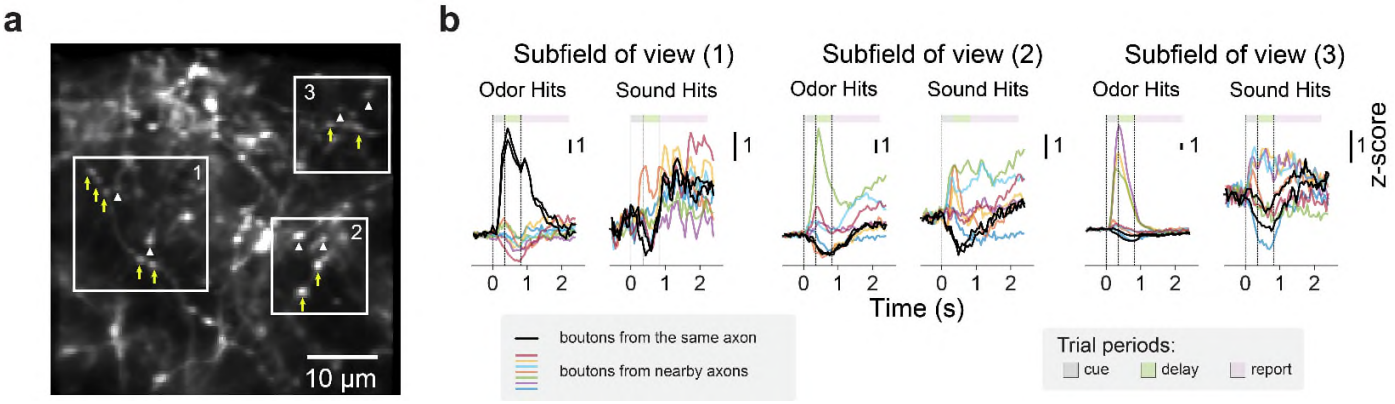

***Reviewer Figure 3**. **a**. Example field of view (48 x 48 µm, ~300 µm from bulb surface); cortical bulbar feedback axons expressing GCaMP5: several axonal branches can be visually identified (squares); yellow arrows mark example boutons on the same axon; white arrow heads mark neighboring boutons from potentially different feedback axons. **b**. Example trial-averaged responses of boutons in **a** for Odor Hit trials (Left) vs. Sound Hit trials (Right) across three example subfields (#1, # 2, # 3). Horizontal grey lines mark the duration of sensory cue (odor or sound) delivery. Responses of boutons on the same axon are shown in black (two example boutons per panel are included). Other color traces correspond to the responses of neighboring boutons. In a given subfield, the response of a given bouton across Odor vs. Sound Hit trials is shown in the same color. Note the diversity of responses in a small spatial vicinity in terms of amplitude, shape and polarity. Boutons on the same axon appear to respond in a similar fashion to the same stimulus, while the responses of nearby boutons from different axons can be widely different. Across trial types (Odor Hits vs. Sound Hits), a given bouton may respond differently.*

Indeed, example feedback boutons, as close as <5 µm apart, but putatively belonging to different axons, showed widely different responses to the same stimulus compared to equally spaced boutons on the same axon (*multicolor traces,* **Reviewer Fig. 3b and Fig. 4**). This is consistent with many reports including work from our lab (*Petreanu et al., 2012; Glickfeld et al., 2013; Otazu et al., 2015; Tanaka et al., 2017; Yoshida et al, 2018*) that boutons on the same axon are more similar in their responses to sensory stimuli, as well as in their spontaneous activity than boutons on different axons.

This was even more apparent when considering the pairwise correlation of responses across boutons for trials of the same outcome (Odor Hits response correlation, Avg±SEM: 0.91±0.04 vs. 0.50±0.15; Sound Hits response correlation, Avg±SEM: 0.74±0.08 vs. 0.40±0.15; $p<0.0001$, Wilcoxon rank-sum test, **Reviewer Fig. 4b**). Importantly, the nature of the response changes across stimuli and respectively reward contingency conditions was correlated for boutons on the same axon. In contrast, adjacent boutons on different axons changed their responses in an uncorrelated manner (**Reviewer Fig. 4b**).

Specifically, we assessed whether the responses of bouton pairs on the same axon varied in a coordinated manner *across different contingencies for the same odor stimulus* (rewarded vs. non-rewarded, **Reviewer Fig. 5**), and *across different stimuli for the same reward contingency* (odor hits vs. sound hits, **Reviewer Fig. 6**). Further, we asked whether stimulus and contingency dependent changes in the responses of these boutons were different for similarly spaced-apart nearby boutons presumably belonging to different axons (within ~15 µm). We found that, on average, responses of pairs of boutons from the same axon were substantially more similar than for pairs of boutons from different axons, within equally small vicinities (<15 µm). This was true both when parsing responses by a given stimulus (same odor across different contingencies, **Reviewer Fig. 5**) or by a given contingency (same contingency across different stimuli, **Reviewer Fig. 6**). We discuss these results in the revised manuscript (rows **357-383**).

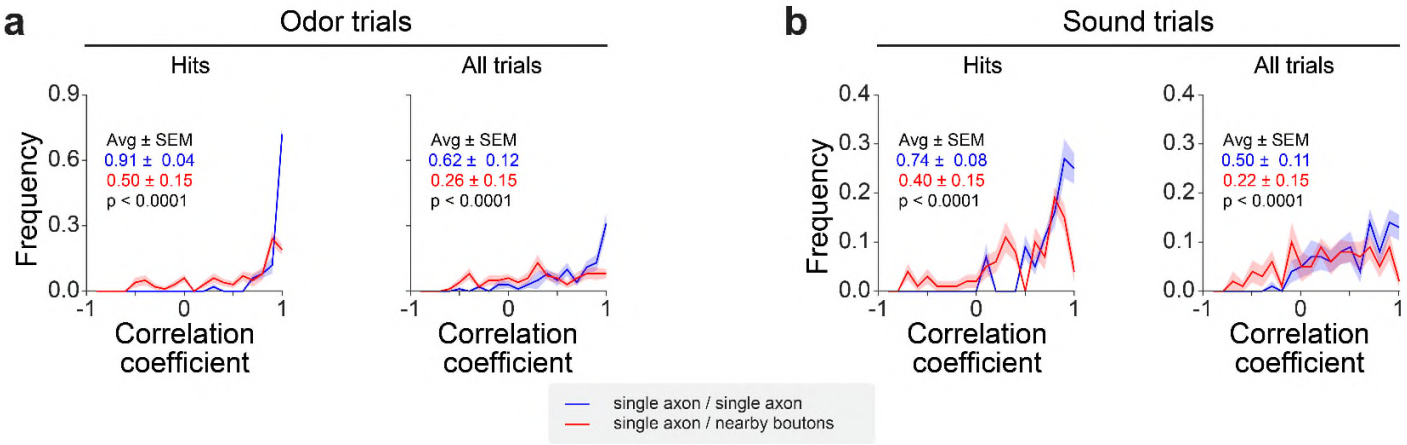

Reviewer Figure 4. ***a***. *Histograms of average response pairwise correlations for boutons on same (blue) or different, nearby (red) axons in odor hit trials (Left), and across all types of odor trials (hits, false alarms, correct rejection, misses, Right). Bouton pairs were selected to be at most 15 µm apart. Note that boutons on the same axon were substantially more similar in response than boutons on different axons;* ***b***. *Same as* ***a*** *for sound trials: hits (Left); all sound trials (Right). Boutons on same axon were substantially more similar in response than boutons on different axons. Pairwise bouton analysis included 46 pairs of boutons on the same axon and 400 bouton pairs across axons; we were able to identify boutons on the same axon in 7 out of 20 FOVs. Bootstrapping: random samples were picked from each group to perform numbers-matched comparisons, iterating 10 times. Traces and shaded areas correspond to average and SEM.Statistical comparissons: Wilcoxon rank-sum test.*

Note that our results outlined above represent a lower bound estimate of the local diversity between boutons from different axons as some of these boutons may, in fact, belong to the same axon even though they may not visually appear so in the low resting baseline fluorescence images.

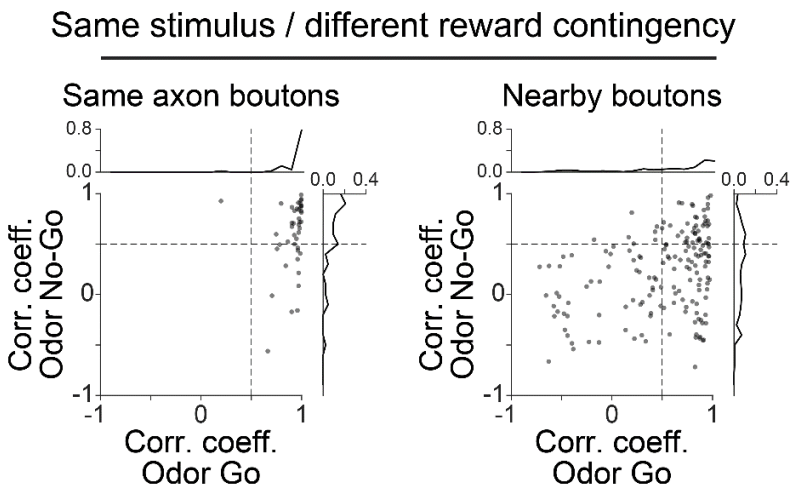

Reviewer Figure 5. ***Same stimulus, different reward contingencies (Odor Go vs. Odor No-Go blocks)***. ***a***. *Scatter plot and histograms of the correlation of average responses of pairs of boutons on the same axon (<15 µm apart) to the same odor stimulus in Odor Go blocks (odor cue was rewarded) vs. Odor No-Go blocks (sound cue was rewarded). Each dot represents the correlation value of a bouton pair on the same axon. Odor Go vs. Odor No-Go same axon bouton pairs: 0.91±0.14 vs. 0.52±0.04* ***b***. *Same when comparing pairs of boutons across nearby axons; in particular, we compared the responses of boutons from* ***a***. *to those of boutons from neighboring axons (<15 µm apart); Each dot represents the correlation value of a bouton pair across nearby axons. Odor Go vs. Odor No-Go nearby axon bouton pairs: 0.59±0.09 vs. 0.23±0.02. Note that pairs of boutons on the same axon were on average substantially more similar in their responses than pairs of boutons on different axons. Higher dispersion of correlation values between boutons on the same axons in the odor No-go blocks (**a**) reflects lower responsivity compared to the Odor Go blocks.*

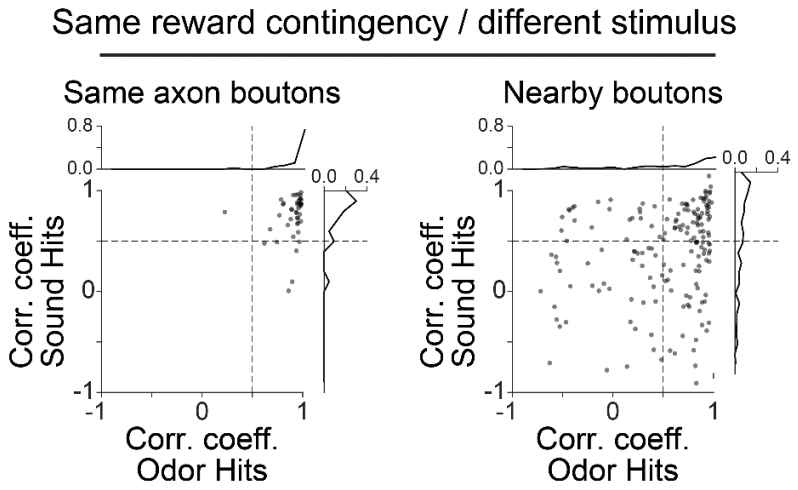

***Reviewer Figure 6. Same reward contingency, different stimuli (Odor Hit vs. Sound Hit trials)**. **a**. Scatter plot and histograms of the correlation of average responses of pairs of boutons on the same axon (<15 µm apart) to the same reward contingency across stimuli in Odor Hit vs. Sound Hit trials. Each dot represents the correlation value of a bouton pair on the same axon. Odor Hits vs. Sound Hits same axon bouton pairs: 0.91±0.04 vs. 0.74±0.08. **b**. Same when comparing pairs of boutons across nearby axons; in particular, we compared the responses of boutons from **a**. to those of boutons from neighboring axons (<15 µm apart); Odor Hits vs. Sound Hits nearby axon bouton pairs: 0.50±0.15 vs. 0.40±0.15. Note that pairs of boutons on the same axon were on average substantially more similar in their responses than pairs of boutons on different axons. Higher dispersion of correlation values between boutons on the same axons in the Sound Hit trials (**a**.) reflect the lower responsivity compared to the Odor-Hit trials. 46 pairs of boutons on the same axon; 24 axons; 400 bouton pairs across axons; we could identify boutons on the same axon in 7 out of 20 FOVs.*

One potential alternative scenario that our analyses cannot discard is that neuromodulation acts within the bulb in a cortical feedback axon specific manner (i.e. boutons on the same axon are modulated in the same manner) due to unique combinations of receptors and downstream signaling cascades in individual feedback axons. In this scenario, the feedback responses may indeed not reflect the spiking activity of the cortical neuron *per se*. To account for this possibility, we adjusted the wording throughout the text, and modified the title of the revised manuscript accordingly (rows **637-655**), while noting that, to the best of our knowledge, this has not been documented by other studies to date. Our manuscript now acknowledges this possibility. Importantly, even if this unlikely scenario would be at play, this would not infringe on our results that show rapid reformatting of cortical feedback axon activation in response to reward contingency. On the contrary, our manuscript motivates new exploration avenues to understand the nature and source of this rapid reformatting of cortical feedback axon bouton dynamics.

We thank the Reviewer for their thoughtful and constructive critique. We think the added analyses provide a clearer picture and significantly strengthen our claims, which we state now in a more cautious manner, as suggested by the Reviewer.

Reviewer #3 (Remarks to the Author):

In this manuscript the authors present evidence that piriform cortical feedback to the olfactory bulb carry multimodal information. Interestingly, they show that response properties of individual boutons can change across blocks within a session depending on reward contingencies. This is novel and significant. However, while the main observation and data showing cortical feedback to OB is multi-modular is convincing, there are some specific claims about neural responses were not fully supported by experimental design and analysis. In particular, the behavioral task design makes many interpretation of the data confounded, and this is not acknowledged. The authors should address these concerns about experimental design, framing, and interpretation of the data before the manuscript is accepted.

We thank the Reviewer for her/his assessment of our work and would like to address the concerns raised.

Major concerns:
1. Block structure in behavioral task is predictable. What is the exact structure of the block lengths and how was this decided? If the hazard rate function of block length is not flat, (that is, if mice can predict block length), mice would be able to learn the trial structure of block switches, and solve the task uni-modally. For example, task mice learn can be: odor-> go for 45 trial; and then rule switch to odor -> non-go for 45 trials. If mice only used odor and not sound at all, they can still solve each block at 75% chance level (100% correct for all odor trials, and 50% correct for all sound trials). Similarly, mice could also only use the sound to solve the task. This is a particularly likely strategy given that trial structure is also temporally predictable: ITI intervals and cue lengths (ITI hazard is not flat) makes it such that mice can know the temporal likelihood of when sensory cue will occur. This means that if mice ignored sounds completely, they can successfully use the strategy of "lick at time ~X if I don't smell an odor" for the tone trials of the "sound-go" blocks and achieve high performance. In fact, given the *variety of behavioral and neural responses* that authors report (for example, results in Sup Fig 8), it is likely that different mice are using different strategies to solve the task. The authors should present evidence to support their claim that mice are using different sensory modalities to solve this task.

The Reviewer raises two concerns regarding the behavioral strategy of expert mice in our task: 1) whether mice solve the task unimodally, or use both odor and sound cues; 2) whether mice can predict the block transitions given the fixed block size. These are indeed important issues to consider in interpreting our results, and we address them below.

The concern regarding the possible reliance on only one of the sensory cues stems from the assumption that mice may be able to predict the cue onsets. If so, they could respond to one cue in an informed manner (~100% performance), and respond randomly to the other (~50% performance), but in a timely fashion given the predictable cue onset timings. This is indeed a potential strategy mice could take. However, within the settings of our task, this is not possible since we actually used a flat inter-trial interval (ITI) hazard rate. The ITI is drawn from an exponential decay distribution within a 0.3 s to 1.2 s range on top of a bias value (9 s): $f(x|\mu) = \frac{1}{\mu} e^{\frac{-x}{\mu}}$ with: $\mu = 0.2\ s$; $0.3\ s \leq f(x) \leq 1.2\ s$ (**Reviewer Fig. 7a**, also shown as **Supplementary Fig. 1g**). We apologize for not stating this clearly in the initial version of the submitted manuscript. In the revised manuscript, we explain these important technical aspects in detail (rows **128-134**; **823-830**; **891-902**), and also include a distribution of the inter-cue periods, i.e. the time intervals between the offset of the sensory cue in a given trial and the onset of the sensory cue in the next trial. For each pair of trials, this includes a fixed 500 ms *delay* period, followed by the *report* period (up to 1.5 s, when the mouse can lick the reward spout). The animal's behavior adds additional variability to the effective inter-cue duration: a) in a Hit trial, a fixed water reward is delivered, and the behavior control software moves to the next state (ITI) once the animal ceases to lick for 100 ms; b) in a false alarm trial, as the animal licks, no water reward is delivered from the spout, and a 10 s additional time-out is imposed before moving to the ITI state; c) in a correct rejection or a miss trial, the control software awaits until the end of the 1.5s report period before moving to the ITI state. As the inter-cue period and the ITI vary from trial-to-trial (**Supplementary Fig. 1f**), mice cannot lock their licking (or lack thereof) to a strict time window following the cue onset (as we observed in our behavioral data, **Fig. 1e, Supplementary Fig. 1b**), unless they actually detect and act on both types of sensory cues.

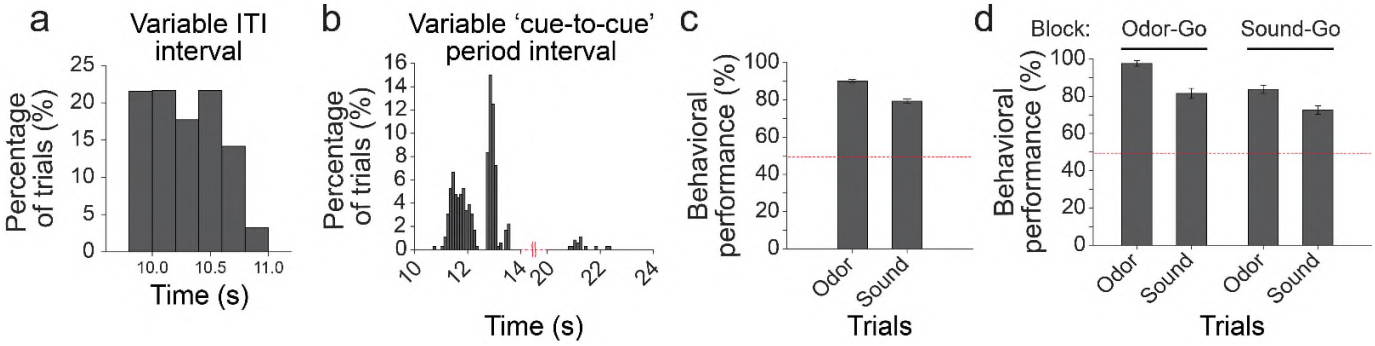

Reviewer Figure 7. a. *Distribution of variable inter-trial intervals (flat hazard rate);* ***b****. Distribution of cue-to-cue intervals for the delay task (20 fields of view, 3 mice).* ***c****. Bar plots showing average session behavioral performance for odor and sound*

trials (20 fields of view, 3 mice, delay version of the odor/sound reversal task, Avg±SEM. ***d.*** *Bar plots showing the average session behavioral performance from the same dataset presented in* ***c****, parsed by cue identity (odor vs. sound trials) for each block type ('Odor-Go' and 'Sound-Go' blocks).*

Furthermore, in **Fig. 4g** (optogenetic suppression and control experiments), we show that expert mice perform at high success rate (> 80%) in both the odor, as well the sound trials, thus ruling out the strategy of random responses to one of the sensory cues. To further address the Reviewer's concern, in the revised manuscript we also included an analysis of the performance of expert mice specifically for the odor versus the sound trials in the imaging sessions (delay version of the reversal task). Indeed, we find that expert mice perform at > 80% success rate in both the odor as well as sound trials (**Reviewer Fig. 7b**, also shown in the revised manuscript as **Supplementary Fig. 1d**). We observed similar performance in the no-delay version (not shown). Importantly, when parsing trials by block type, the average performance during sound trials was substantially above chance in both Odor Go and Sound-Go blocks. Same was observed for the odor trials when analyzed separately for each block type (**Reviewer Fig. 7c**, also shown as **Supplementary Fig. 1e**). In the revised manuscript, we further clarify and discuss these points (rows **145-150**).

Can the mice predict the block transition?
Expert mice perform ~270 trials/session. We chose a block size of 45 trials as a tradeoff between being short enough to afford multiple switches per session (~5) and long enough to allow for performance stabilization after each switch. Keeping a flat hazard rate for the number of trials per block was difficult under these constraints. In a subset of experiments (when performing the sniffing analysis), the length of the blocks was randomly varied between 42 to 48 trials per block (**Reviewer Fig. 8a**; also included in the revised manuscript as **Supplementary Fig. 1h**). In these experiments, mice also achieved > 80% expert performance (both during Odor Go, and Sound Go blocks) within a similar training time as in the other experiments. In these experiments, the performance for the odor and respectively the sound trials was comparable to the imaging sessions (>80%, **Reviewer Fig. 8b;** also included as **Supplementary Fig. 10e**).

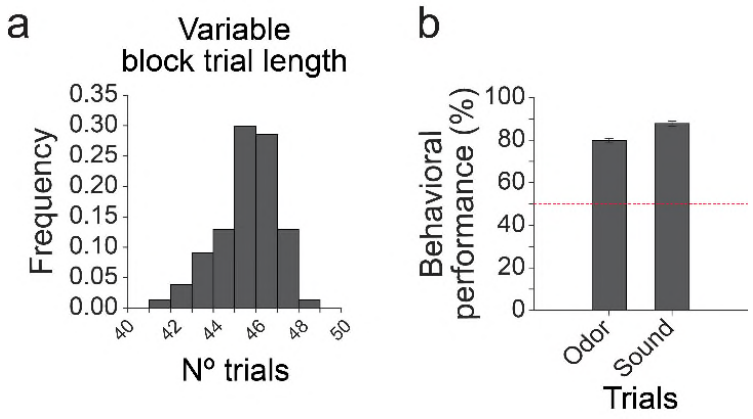

Reviewer Figure 8. a. *Distribution of number of trials per block (block length) from expert mice used for the sniffing analysis control (15 sessions, 3 mice).* ***b****. Bar plots showing average session behavioral performance for odor and sound trials of the same data (Avg±SEM).*

We reason that, despite the fixed block size, the behavioral data suggests that mice do not predict the length of the block: expert mice made mistakes at the boundary between blocks, but only after the switch point, and not before. Across sessions and mice, we observed a sudden dip in performance at the boundary between blocks (following the rule switch and not preceding it), followed on average by a gradual increase in performance (presumably informed by the change in the reward contingency rules, **Fig. 1c bottom**, **d**; **Supplementary Figs. 1c**,**i**,**j**; **Fig. 4d**). This gradual and variable increase in performance (from block to block within a given session) was tightly correlated with changes in cortical bulbar feedback activity (**Figs. 4a-e**). If mice relied on a 'noisy' estimate of the block size, one would expect mistakes to be randomly distributed around (before and after) the block boundary, which we don't observe in the data. It is also important to note that while the number of trials per block is constant, the total duration of a given block varies given the flat hazard rate ITIs (**Reviewer Fig. 7a**), and the variable inter-cue interval as a function of each trial's outcome (false alarm trials trigger additional time-outs, etc.) as explained above. Therefore, mice cannot use a simple time-keeping strategy to predict when the rule switches. In the revised manuscript, we discuss these issues accordingly (rows **152-169**).

In lieu of behavioral evidence for this, the most parsimonious explanation for how mice would be solving the task. In this case, there are a few issues: a. Authors do not know which sensory modality mice are using to solve the task for any given trial. Therefore, interpretation of behavioral relevance of neural signal is inherently limited and all claims in manuscript about sensory-domain specific learning should reflect this. For example, Page 20: “As learning of the sound-reward associations progressed, …” sound be rephrased.

As explained above, the precisely timed responses to both the odor, as well as to the sound cues despite variable ITIs indicate that mice detect and respond to both these cues. We apologize for the potential confusion. The statement on Page 20 refers specifically to the auditory-only Go/No-Go task where mice are only required to discriminate between two sound cues: one that is rewarded and the other that is not. Following the Reviewer’s suggestions, in the revised manuscript, we further clarified accordingly all statements about sensory-domain-specific learning and provided additional context to enhance clarity (rows **453-482**).

b. Even if mice are solving the task using sound only in the “sound go” blocks, and odors only in the “odor-go” blocks, because the block length is predictable, expectation (sensory prediction, reward prediction) is systematically manipulated relative to block switch. It is known that OB activity is modulated by association, memory, and attention, behavior or neural responses that depend on analysis aligned to block is confounded by this. Therefore, the claim that neural and behavioral differences across block switches are due to the sensory domain differences is not supported. This means that the main conclusion/claim of this study needs to be carefully framed, and perhaps reinterpreted completely.

As described above, mice cannot use the sound cue only in the Sound Go blocks and odor cue only in the Odor Go blocks, since the ITI is not constant, but follows a flat hazard rate, as well as due to the variable trial duration, leading to variable cue-to-cue intervals (**Reviewer Fig. 7a**). Indeed, we find robust, substantially above chance performance, in the odor trials as well as the sound trials, irrespective of the block type (**Reviewer Figs. 7b**,**c**, **8b**). Further, expert mice do not appear to predict block transitions as their performance drops consistently only after the block switch, and not preceding it (**Figs. 1c**, **d-bottom**; **Supplementary Figs. 1c**,**g**,**h**; **Fig. 6d**).

Importantly, we do not claim that the differences we observe in the feedback bouton responses across blocks are due to sensory domain differences, but rather that they emerge because of differential reward associations of the *same sensory stimuli* which update whenever the rule changes. We further hypothesize that the feedback signal is a source for such changes of activity in the bulb, which emerge as a function of reward associations, as mentioned in the Discussion section of the manuscript. We apologize for any potential confusion. In the revised manuscript, we further clarified these issues.

2. Given the confounds presented in 1), the results in Figure 3 would be a lot more completing if they were compared to neural recordings from awake passive mice over days of exposure to predictable events (same trial structure and block structure without rewards). Which part of the reported results here would be different?

As explained above, individual events (sound or odor) are not predictable given the variable ITIs and inter-cue intervals, thus mitigating the concerns about our results simply reflecting temporal expectations. Following the Reviewer’s suggestion, we performed additional experiments to investigate how activity may change in awake passive mice over days of exposure to predictable events, as detailed below.

Figure 3f-h shows that experience does matter, but given that we do not know what the mice are learning, these panels must be interpreted carefully.

Fig. 3f-h relates to the auditory-only Go/No-Go task, focusing on how cortical feedback representations evolve with task learning. In this task, in any given trial, only one of the two sound cues is presented. Of the two sound cues only one is rewarded. Trial order is chosen randomly and there is no block structure. Mice are thus exposed to both sound cues equally across days, but are only rewarded in response to licking the spout to one of the two cues. With training, mice learn this cue-reward association and refrain from licking to the unrewarded sound. To address the concern raised, in the revised manuscript, we clarified the text accordingly such as to avoid potential confusions (rows **453-459**).

Given that Varga and Wesson (2013) showed that piriform cortex responds tones in anesthetized conditions. In my view, the authors should make clear the expansion of their results beyond this study (aka tones evoke responses from piriform cortex neurons). Can the authors compare results between awake passive naïve mice vs awake passive experienced mice, and contrast these differences with results in their current manuscript (even if statements about domain-specific learning cannot be made). For example, highlighting the difference between naïve and experienced mice.

We thank the Reviewer for the suggestion. In the revised manuscript, we state that our new contributions include finding that sound responses in the cortical bulbar feedback are modulated: **1)** on a fast (seconds) time scale in expert behaving mice by sudden changes in reward contingencies within the same session (odor-sound rule reversal task), and **2)** across sessions, during learning of the stimulus-reward associations in mice engaged in the auditory-only Go/No-Go task (as observed by comparing naïve and expert sessions, rows **687-711**).

In addition, to address the Reviewer's concern, we performed new sound habituation control experiments (3 mice). We compared across 6 consecutive days the changes in the activity of cortical bulbar feedback boutons as triggered by passively experienced sound cues (A vs. B) in naïve vs. experienced mice (no water rewards were provided in these sessions, **Reviewer Fig. 9** also shown as **Supplementary Fig. 15 d-f, j**).

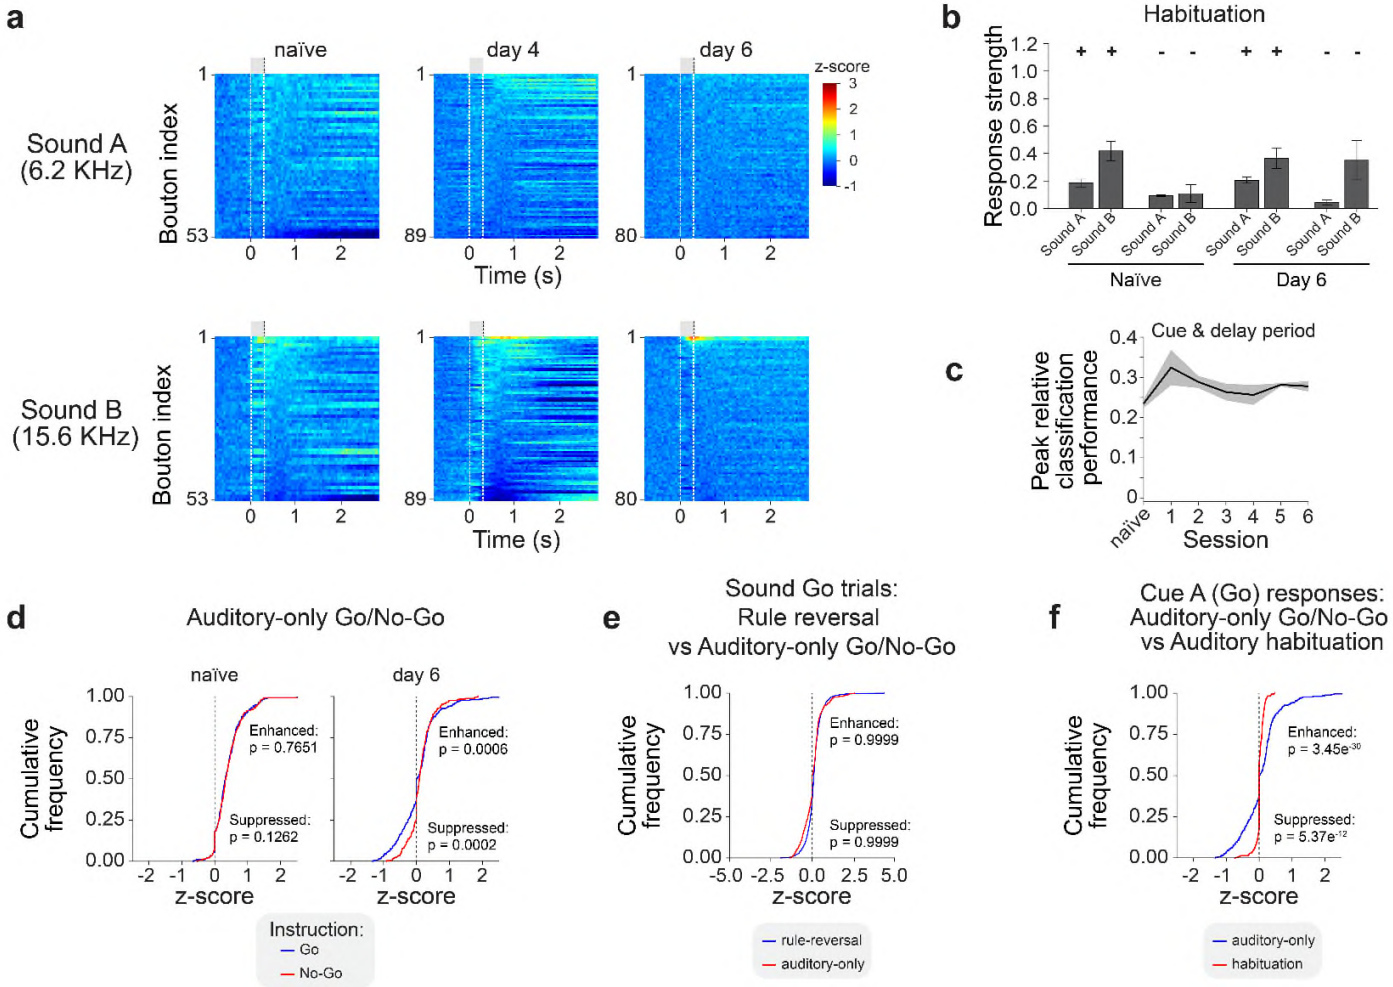

Reviewer Figure 9. a. *Average cortical feedback bouton responses in example fields of view parsed by cue ('Sound A' or 'Sound B') and across days of habituation (naïve, day 4, and day 6).* ***b.*** *Response strength differences to sound-evoked responses in the Two-sounds habituation sessions. We quantified response strength using the average z-score response during a period equivalent to the delay period ('delay like') of the behavioral tasks, sampling in each field of view the top 20 response amplitude boutons. +/- mark enhanced and respectively suppressed bouton responses.* ***c****. Peak performance of the classifier sampled from cue onset to the end of the 'delay-like' period.* ***d****. Cumulative distributions of response strength (peak amplitude during delay period) for naïve (Left; Wilcoxon rank-sum test, enhanced and suppressed: n.s.) and task-engaged (day 6, Right; Wilcoxon rank-sum test, enhanced: p < 0.001; suppressed: p < p.001) mice performing the auditory-only Go/No-Go task. Go responses are shown in blue; No-Go responses in red.* ***e****. Cumulative distributions of response strength for expert mice engaged in the odor/sound rule-reversal task (blue) and auditory-only Go/No-Go task (day 6, red).*

*Wilcoxon rank-sum test, enhanced and suppressed: n.s.. **f**. Cumulative distributions of response strength to Sound A for day 6 in mice performing the auditory-only Go/No-Go task (Go cue, blue) or passively exposed to the same sound cue (habituation, red). Wilcoxon rank-sum test, enhanced: $p < 0.0001$; suppressed: $p < 0.0001$. All panels error bars: ±SEM.*

In these experiments, we did not observe differential modulation of responses to the two sound cues (A vs. B) in a systematic manner across sessions (days 1-to-6). Feedback bouton responses were sparse and their strength lower than in the auditory-only Go/No Go task (**Reviewer Fig. 9f**, also shown as **Supplementary Fig. 15j**; Wilcoxon rank-sum test, enhanced: $p < 0.0001$; suppressed: $p < 0.0001$). On average, responses strength to each cue, did not follow a systematic trend, and appeared relatively stable across sessions. The decoding performance of classifiers (**Reviewer Fig. 9c**, also shown as **Supplementary Fig. 15f**) was flat across these habituation sessions in contrast to the increasing trend in decoding performance during learning in mice engaged in the auditory-only Go/No-Go task (**Fig. 3f**, **g**). Interestingly, we observed differences in the average cortical feedback bouton response amplitude as a function of the sound cue identity (A vs. B). Responses to one of the sound cues were on average stronger, and remained as such across days. We conclude that sound cues trigger sparse responses in the cortical bulbar feedback axons in naïve mice. These responses are shaped by learning the stimulus-reward contingency, and may also depend on other stimulus features. We discuss these new results and analyses in the revised manuscript (rows **483**-**493**). Note that while these comparisons are interesting, they are not immediately related to the main focus of our study which is aimed at investigating whether cortical bulbar feedback represents changes in stimulus-reward contingencies at rapid timescales.

3. Go/NoGo task introduces asymmetries and confounds. There is not a clear axis in the behavioral paradigm that helps to discern activity from reward/motor/sensory cues clearly. All the distinction come from the quantification of the imaging data. In this regard, authors should include a more careful comparison between EGFP signals and GCaMP signals to clearly state what is the expected noise level, and what is clear signal.

Indeed, the Go/No-Go task is asymmetric in nature and activity related to sensory/motor/reward features may be intermixed in time. We specifically included a delay period between cue and report periods to isolate the effects of the motor output and preparatory activity from the sensory and reward-related components. In addition, we performed several controls and analyses to address the same, as described below.

First, we trained mice in two versions of the odor-sound rule reversal Go/No-Go task: 'no delay' and 'delay'. In the 'delay' task, mice were trained to refrain from licking during the cue and delay periods (500 ms), and we excluded from analysis early licks trials. We analyzed separately responses in the boutons occurring during the cue vs. delay vs. report periods. This enabled us to untangle the sensory/reward expectation/motor preparatory activity in the bouton responses from activity related to the motor output *per se* (licking).

Second, we compared responses across trials with different cue/instructions conditions, but similar actions (Hits vs. False Alarms) and vice versa (hits vs. correct rejections). We find that (results presented in **Fig. 3a-d)** bouton responses occurring before the report period (licking) represent sensory information (i.e. cue identity - odor vs. sound), as well as trial outcome - hits vs. false alarms vs. correct rejections). Across fields of view, cortical feedback response amplitude during the hit and false alarm trials was generally higher than for correct rejection trials (**Supplementary Figs. 9e**,**f**). However, we also observed differences in responses of individual boutons when performing pairwise comparisons between trials of different contingencies in which mice licked the reward port. Boutons that differentially modulated their responses across conditions represented a significant fraction of the responsive population (**Fig**. **3b**; Odor hits vs. false alarms, 65.0 ± 7.5 %; Odor vs. Sound hits, 60.1 ± 11.9 %, Odor vs. Sound false alarms, 42.7 ± 12.1%). Similar to the rule-reversal task (**Figs. 3a,b**), in expert mice engaged in the auditory-only Go/No Go task, the cortical feedback bouton ensemble responses during both the cue and delay periods were different for hit vs. false alarm trials, despite the animal licking in both cases (**Supplementary Fig. 15c**). Since in all these cases, mice subsequently licked the reward port, we conclude that changes in the response of individual boutons to same sensory stimulus across reward contingencies cannot be solely attributed to preparatory motor activity.

Third, we trained an independent cohort of mice expressing EGFP instead of GCaMP to assess how licking alters our ability to determine whether a bouton is responsive or not, and explain the diversity observed in the bouton responses. Using the same signal threshold criterion, we found substantially fewer EGFP 'responsive' boutons (in the delay version) and no EGFP 'responsive' boutons (in the no-delay version) compared to GCaMP-expressing mice (**Supplementary Fig. 5e**). In addition, K-means clustering revealed substantially lower diversity

in EGFP 'responsive' boutons vs. GCaMP responsive boutons (4 clusters for EFGP data vs. 18 clusters for GCaMP delay; and 19 clusters GCaMP no-delay data, **Supplementary Fig. 6b-d**). Moreover, we found no correlation between the number of licks and the bouton response strength for the GCaMP data (**Supplementary Fig. 1l**), and no correlation between the number of licks and the percentage of responsive boutons in the GCaMP data for both the no-delay and delay versions of the task (**Supplementary Figs. 1m-o**). In the revised manuscript, we expanded and further discussed these points (rows **244-254**).

In addition, we analyzed how systematically changing the signal threshold alters the number and diversity of responsive boutons in both the GCaMP experiments and the EGFP control sessions (**Reviewer Fig. 10**). For each bouton, the signal threshold was defined using a percentile responsiveness criterion with respect to the distribution of fluorescence fluctuations during baseline, before cue onset, aggregated across all trials in a session (Methods, **Supplementary Fig. 4**). For the analyses presented in the study, we chose a signal threshold corresponding to the top 99% percentile of baseline fluorescence fluctuations to classify the enhanced bouton responses, and respectively the bottom 1% percentile for classifying the suppressed responses. Across different signal thresholds, substantially more responsive boutons were identified in the GCaMP vs. EGFP imaging data.

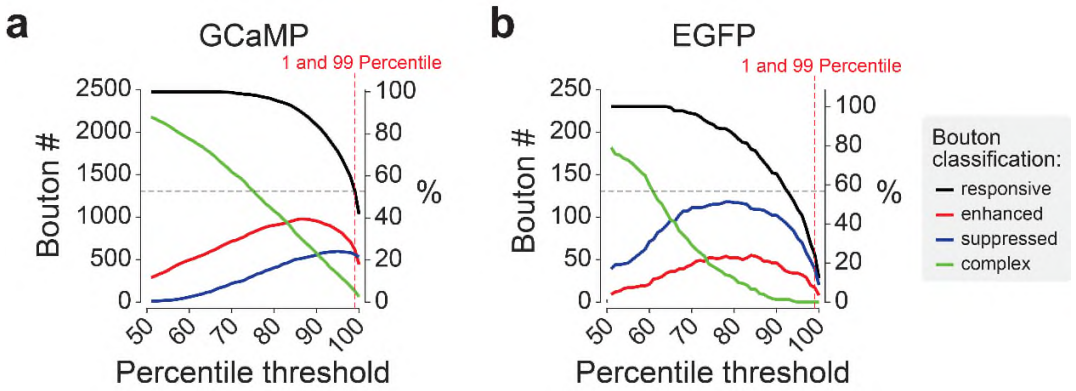

Reviewer Figure 10. *Distribution of number of feedback GCaMP (a) and EGFP (b) boutons classified as responsive (black trace) as a function of systematically varying the signal threshold; signal threshold was defined for each bouton as a percentile of the distribution of fluorescence fluctuations during the baseline periods acquired before cue onset, aggregated across all trials of a session; color traces mark different types of responses: enhanced (red), suppressed (blue) and complex (green, which displayed both enhancement and suppression response periods).*

4. It was not clear to me if the authors are able to track the same boutons over days/sessions. The manuscript would be significantly stronger if analyzes included properties of a subset of individual boutons over sessions and days.

We did not track the same boutons over days/sessions. In the study, we focus on investigating whether and how responses of individual boutons change over short time scales (within the same session) as a result of sudden and repeated changes in stimulus-reward contingency. As such, the investigation of how individual bouton responses change across sessions is beyond the scope of our study. Therefore, we think that, while potentially insightful, these technically demanding experiments are more appropriate for a separate future study.

Minor concerns:
1. In the last sentence of the abstract, it is not clear what "identity" refers to.

We apologize for the lack of clarity. It refers to the stimulus (cue) identity. We changed the text accordingly.

2. Fig 1d is very surprising. It seems very unlikely that there would be no drop in session performance at the first day of rule reversal. How do the authors explain this? There are only 6 reversals in average in one session. Perhaps a session example of a rule reversal training day might give a bit more clarity on what happens? Looking at the Methods section that details the training, it looks like there are a lot of details that are overlooked in the main body of the text. For instance, the first day of reversals animals were given automatic water and other tricks to help them achieve performance. The authors should either make this transparent or tone down the claim regarding fast learning made in Figure 1 and beyond.

We apologize for the confusion and thank the Review for the suggestion. In the revised manuscript, we explain in more detail our training procedures, focusing on Day 0 and the days preceding and following the first in-session rule reversals (rows **837-874**). During the days preceding the first in-session rule reversal ('Day 0'), mice were trained so as to reach an average performance per session of no more than 70% correct trials. We did this, aiming to ensure that mice did not get fixated on one of the two rules we employed. In a given session only one rule was used. By 'Day 0', mice experienced rule-reversals across sessions, and generally were substantially better at performing the task using one of the rules, and close to chance level for the other rule, as shown in an example early training session in **Reviewer Fig. 11** (also included in the revised manuscript as **Supplementary Fig. 1c).** The absence of a drop in the average session performance at Day 0 (**Fig. 1d**) is explained by this relatively low level of behavioral performance that is imposed by the training requirements and high animal-to-animal variability in performance at this point in the training. Once rule reversal within the same session was introduced at 'Day 0', mice retained higher performance for the rule they were more 'comfortable' prior to 'Day 0'. Across further training sessions mice also improved their behavioral performance for both rules. When they reach expert level (that we define at average 80% performance per session), they performed well across both rules, while still displaying slight biases for one of the two.

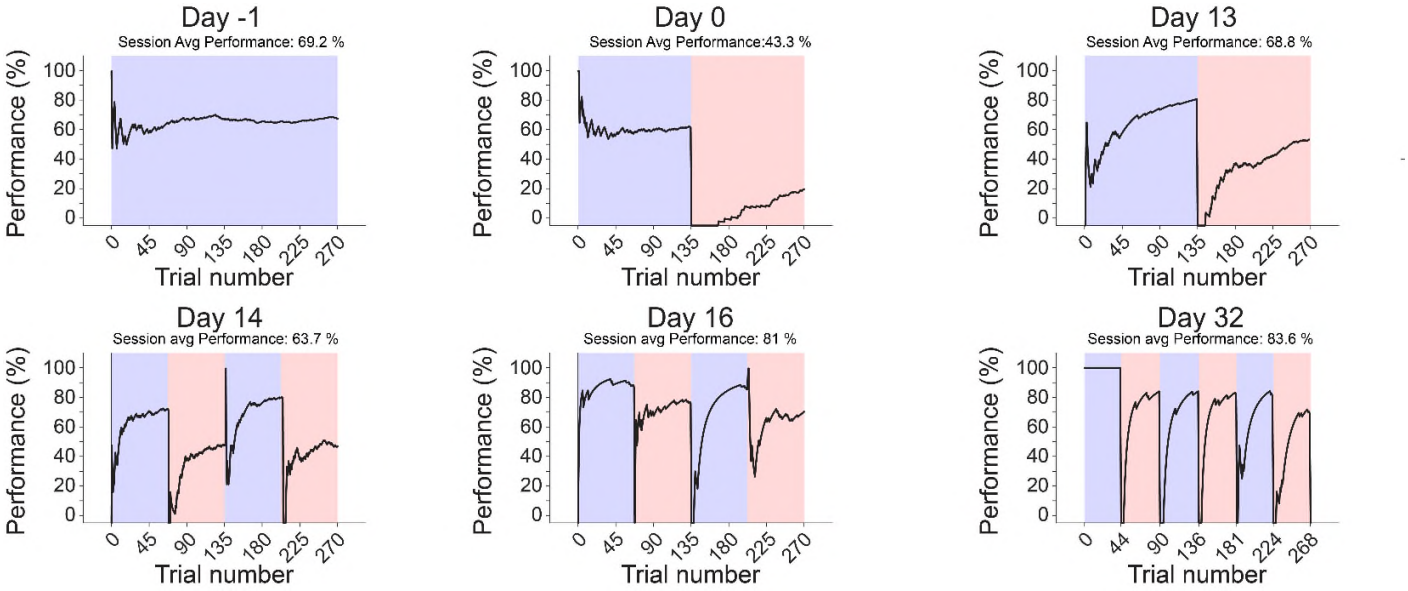

Reviewer Figure 11. *Behavioral performance for example sessions throughout training for one mouse. 'Day -1' corresponds to the last day of training using only one rule per session. 'Day 0' corresponds to the first day when we introduced rule reversal within the same session. As training progressed, the number of trials per block (Odor Go and Sound Go) was systematically decreased. Note that the average performance improved throughout the training. Different blocks are shown in different colors (Odor Go in light purple and Sound Go in pink).*

In the revised manuscript, we also clarified the procedures used for automatic (free) water delivery on Day 0, and throughout training. These were used as a means to remind the animals about the presence of water rewards, and were not counted as correct trials (rows **854-856**). We also discuss that the training procedure takes several weeks including ~ 2 weeks before introducing in-session rule reversals (prior to 'Day 0').

3. In general, the paper is missing statistical comparisons for several of the claims they make. Information such as number of trials per session, reversals, accuracy, etc. should also be including the SD of the values.

We apologize for this oversight. We went through the manuscript and included missing statistical comparisons, number of trials per session, reversal, accuracy, etc. including the SD of the reported values. We include a Table with all statistical comparisons.

4. Is the sampling of the odor locked in with the breathing? Does it take the same time to process both information?

The odor delivery was not triggered on respiration. On average, mice took similar amount of time to respond to the odor and sound cues (930.8 ± 3.7 ms for the tone and 986.4 ± 7.6 ms for the odorant from *cue* onset in the 'delay' version, N= 3 mice; 445.0 ± 2.4 vs. 443.0 ± 3.2 in the 'no-delay' version; N= 4 mice, **Fig. 1e**, **Supplementary Fig. 1b**). In the revised manuscript, we further clarified these points in the main text (row **878**).

5. Comparison between the task with delay and the no delay. It is unclear in which way the authors expected to see differences between delay and no delay task. They mention how the delay can eliminate possible confounds but that is not directly observed in any analysis or metric that they access.

We apologize for the oversight. We state in the text:

*"…To disambiguate the neural signatures of reward contingency from motion artifacts related to licking and other signals, in a subset of experiments we imposed a 500 ms delay between the cue offset and the reporting period, when water reward was available (***Fig. 1b**, **Supplementary Fig. 1a** *'no-delay' vs. 'delay' versions)."*

"…In the no-delay version of the task, decoding performance returned to baseline more rapidly than in the delay version, reflecting faster offset kinetics in the cortical bulbar feedback consistent with differences in our experimental design (**Fig. 3d**, **Supplementary Figs. 14a,b***)."*

We did not expect to see substantial differences between these two versions of the task. In our initial design, we started with the no-delay version. Once we successfully trained mice in this version of the task, we further implemented a version in which we imposed a 500 ms delay between the cue offset and the report period. As explained above, this enabled further analyses and disentangling potential confounds. In the revised manuscript, we explain the choices made in the task design accordingly (rows **876-882**).

6. Page 6. "Reaction time" for delay version of the task in not meaningful, given the cue length and enforced delay. Its misleading to use that term.

We agree. In the revised manuscript, we call it: 'report latency' (row **173; Fig. 1** legend).

7. Page 7. Licking rate really overlapping in the figure but then substantially different in the text. Include statistics and/or change the axis size to be able to perceive the subtle differences.

In the text we state: "*We observed comparable reaction times across modalities (930.8 ± 3.7 ms for the tone and 986.4 ± 7.6 ms for the odorant from cue onset in the 'delay' version…"*

The distributions of report latencies are similar for the odor and sound cues (overlapping peaks), but differ in their tails. Slight differences in average report latencies (~50 ms) are not readily apparent in the figure due to the large range of time shown (**Fig. 1e**, bottom; 50 ms average difference embedded in a several second display range). In the revised manuscript, we included an inset in **Fig. 1e** to better illustrate these results.

8. Supp 1c. Don't understand the switch frequency plot.

We apologize for the lack of clarity. In the revised manuscript, we describe this in detail. Briefly, to quantify correct switching, in former **Supplementary Fig 1c, d** (**Supplementary Fig. 1i**,**j** in the revised manuscript) we used as criterion the probability of observing five correct trials in a row, following the rule-reversal event. In other words, following the rule-reversal, we computed the percentage of cases of having five correct trials occurring one after another, starting at the trial when rule-reversal occurred and using a five-trials sliding window (**Supplementary Fig. 1i**). We further plotted this as a cumulative distribution as a function of the number of trials post switch using the same five-trials sliding window (**Supplementary Fig. 1j**). The first entry is sampled across the first five trials post-switch (trials 1-to-5), the second entry across trials 2-6, etc.

9. Page 7. Where is the analysis that Supp. Fig. 3e points at?

This is associated with the previous panels of **Supplementary Fig. 3 (a-d).** It deals with the quantification of responsiveness for all cortical feedback bouton data shown in these panels (a-d). In the revised manuscript, we

further included a new supplementary figure that explains the procedure used for determining the signal threshold (**Supplementary Fig. 4**). In the revised manuscript, **Supplementary Fig. 3** became **Supplementary Fig. 5**. We clarified in the text the steps used for classifying responsive boutons and the quantification of responses (rows **1024-1040**).

10. Page 9. Is this the best way to look at responsiveness?

We analyzed what fraction of boutons showed changes in fluorescence triggered by the odor and sound cues in the delay and no-delay versions of the task, using top 99% of the baseline fluorescence distribution as signal threshold for enhanced responses and bottom 1% of the baseline fluorescence distribution as signal threshold for suppressed responses (**Supplementary Figs. 4a-c**). In the revised version of the manuscript, we also systematically varied this threshold and assessed how number/fraction of responsive boutons varies (**Supplementary Fig. 4d**).

11. In line 193, Pg. 9 authors state that "across the cue-modulated feedback boutons … 10.1 ± 2.7 % were selectively tuned to the tone". In line 207, Pg. 10 they say: "only a small fraction of EGFP boutons passed our responsiveness criterion … 15.5 ± 7.5%; ". If we take the fraction of responsive EGFP boutons as chance level, this means that tone-modulated boutons are below chance level? Are these two statistically different?

We apologize for the lack of clarity. Indeed, we find that 10% of boutons were responding to the sound cue and not to the odor cue, but additionally also that ~30% of the feedback boutons respond to the sound, as well as the odor cues. So, in total, approximately 40% of boutons were responsive to the sound cue (**Supplementary Fig. 6a**). This is substantially above the chance (EGFP) level. In the revised version of the manuscript, we clarify these results (rows **230-234**).

12. Page 11. Isn't the fact that the boutons respond only in reward trials for sound indicates that it's mostly either reward/movement related?

Yes, indeed, the bouton responses are strongly modulated by the stimulus reward contingency. While we cannot exclude fully movement preparatory activity, as explained above, we observe differences between activity of individual trials in hit and false alarm trials (in both trial types mice lick, **Fig. 3b**). Also, we do observed differences in the boutons that respond to odor vs. sound, so at the same time, cortical feedback represents stimulus identity information as well. In the revised manuscript, we also included an analysis of hits and false alarm trials in the auditory-only Go/No-Go expert sessions (**Supplementary Fig. 15c**). Similarly, we found that the ensemble feedback responses during the cue and delay periods were different for hits vs. false alarm trials, despite the animal licking in both cases, indicating that bouton responses preceding the lick represent stimulus reward contingency and cannot be solely explained as movement preparatory activity.

13. Page 11, line 238. Where can these observation be found?

Fig. 2e (Sound trials). In the revised manuscript we reference this accordingly.

Response to reviewers

We thank the reviewers and the editorial team for their valuable and insightful support throughout the review.

Reviewer #1 (Remarks to the Author):

The authors have addressed all my concerns.

One thing is that some of the preprints that are cited have since been published in peer-reviewed journals. I have listed these citations below.

Congratulations to the authors on the beautiful study.

31. Young, H., Belbut, B., Baeta, M. & Petreanu, L. Laminar-specific cortico-cortical loops in mouse visual cortex. bioRxiv 773085 (2019) doi:10.1101/773085.
-> Now published at eLife (https://doi.org/10.7554/eLife.59551)

107. Cole, N., Harvey, M., Myers-Joseph, D., Gilra, A. & Khan, A. G. Prediction error signals in anterior cingulate cortex drive task-switching. 2022.11.27.518096 Preprint at https://doi.org/10.1101/2022.11.27.518096 (2022).
-> Now published at Nature Communications (https://doi.org/10.1038/s41467-024-51368-9)

117. Lindeman, S., Fu, X., Reinert, J. K. & Fukunaga, I. Reward contingency modulates olfactory bulb output via pathway-dependent peri-somatic inhibition. 2023.08.17.553686 Preprint at https://doi.org/10.1101/2023.08.17.553686 (2023).
-> Now published at PLOS Biology (https://doi.org/10.1371/journal.pbio.3002536)

We edited our manuscript to address the publication of the preprints mentioned by Reviewer #1. We appreciate the positive reviewer's evaluation and constructive support.

Reviewer #2 (Remarks to the Author):

My major concern was that the responses the authors reported were inconsistent with previous studies and that this might be attributable to neuromodulatory signals on the axons of piriform cortex cells that project back to the bulb. The authors argue against this based on spatial selectivity. Specifically, they argue that if the signals observed in the axon terminals are neuromodulatory, the signals recorded in nearby terminals should be correlated. However, if the signals they record are due to the activity of the neurons in the piriform cortex, then responses on boutons from the same axon will be highly correlated while responses for similarly proximal boutons on different axons should be uncorrelated.

I agree with this argument; the data the authors provide to support it are compelling. Most of my other major concerns (e.g., why responses are so different from many previous studies) hinged on the concern that the authors weren't really recording cortical responses, and so I now consider my previous concern in this regard to be resolved. I think the authors have done a commendable job addressing the concerns of all three reviewers and I have no major reservations about this manuscript. Indeed, the Federman et al. paper that was published in the meantime, and the recent paper from Kehl et al. just published in Nature, show that the piriform cortex is much more than a simple primary sensory cortex (of course it is, that is a straw man), but that it exhibits both contingency and multimodal signals. I am still a little surprised by how dense their odor responses are, but I think the authors have done everything correctly. As such, this is an important paper that should be of interest to a large audience.

We appreciate the positive reviewer's evaluation and constructive support.

Reviewer #3 (Remarks to the Author):

The authors have addressed the issues I raised during my initial review. In my opinion the manuscript is significantly strengthened. Particularly considering the recent publication from Federman et al., Nature Communications, 2024, I believe this manuscript is a timely and significant advance for the field. In particular, the additional explanation about the ITI durations is indeed helpful for ruling out a timing-based strategy that can explain behavior. Similarly, I am happy with the explanation provided about how block length estimation is also not a strong confound given the behavioral observations about failures.

We appreciate the authors for conducting extra experiments to answer this question. When comparing figure Reviewer 9c and Figure 3h, it is striking that the peak relative classification performance is already around 0.3 for the naïve session in the habituation only experiments but between 0 and 0.1 for the Auditory only experiments. Why is this the case? I would expect a performance around 0 (chance level). The authors may consider clarifying this in their final manuscript.

We edited our manuscript to address the points raised by Reviewer #3 to discuss about the performance of the classifier in the auditory-only habituation experiments (line 394). We find that in the control experiments, the cortical bulbar feedback responses and the performance of the classifier do not change in a systematic manner across days as a function of habituation, in contrast to the increase in decoding performance during reward-associated learning. However, different tones may trigger different responses in the cortical bulbar feedback, and as such the classifier's performance is above chance.

We appreciate the positive reviewer's evaluation and constructive support.