

Hippocampal theta states are linked to reconfiguration of population activity with respect to behavioural context

Corresponding Author: Ms Laura Masaracchia

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)

The manuscript by Masaracchia et al. provides an intriguing exploration of theta oscillations in hippocampal neural ensembles. While the use of data-driven methods to distinguish between two closely related theta states (LPLT and HPHT) is methodologically sophisticated, several critical concerns arise regarding the interpretation and novelty of these findings. I have made a list of these concerns below-

The described LPLT (~7–8 Hz, lower amplitude) and HPHT (~8–9 Hz, higher amplitude) states closely parallel the classical characterization by Vanderwolf (1969) of type I (movement-dependent) and type II (immobility or arousal-dependent) theta oscillations. It is problematic that the authors neither acknowledge nor cite Vanderwolf's foundational work or subsequent literature (e.g., Bland, 1986). Clarifying how their findings differ meaningfully from these classical concepts would strengthen the manuscript.

Theta amplitude, frequency, and neuronal firing rates are well-established correlates of locomotion speed. Given that the authors report higher firing rates and increased theta amplitude and frequency in their HPHT state, a critical omission is the lack of running speed measurements or controls. Without addressing this factor explicitly, the interpretation of their "ensemble pattern modulation" as a novel mechanism remains questionable.

The manuscript presents "ensemble pattern modulation" as a form of theta modulation distinct from classical phase-based modulation (e.g., phase precession). However, without explicitly analyzing theta phase relationships or convincingly demonstrating independence from classical theta phenomena (including phase precession or known frequency-amplitude correlations during locomotion), this claim of novelty appears overstated.

The reported frequency distinction between LPLT and HPHT states is minimal (~1 Hz difference). While technically achievable using data-driven methods (HMM), the biological relevance and interpretive significance of such a minor frequency difference remain uncertain. The authors should clarify why these subtle spectral differences hold substantial functional relevance rather than reflecting minor variations within a continuous distribution.

The authors state clearly that most neurons fire during both theta states, differing only quantitatively in firing rates. This finding suggests the described states modulate neural activity primarily through known amplitude-frequency correlations rather than recruiting functionally distinct neuronal subpopulations.

Given the extremely short analysis window (32 ms), the spectral resolution inherent in the HMM-TDE approach requires clarification. Can the authors explicitly state their actual frequency resolution under these conditions? Moreover, we recommend benchmark testing of their HMM-TDE parameters on simulated data designed to generate known theta harmonics similar to those observed by Buzsaki in the 1980s (later in work by Czurko and Hirase) and shape changes as described by Terrazas, and McNaughton (2005). Such testing would help validate whether the observed states represent

distinct oscillatory modes or artifacts of their analysis method.

Importantly, no spectral decomposition method, including HMM-TDE, circumvents the fundamental limitations imposed by time-frequency tradeoff windows. Thus, it appears likely that the HMM-TDE approach is sensitive primarily to waveform shape, amplitude patterns, or transient temporal structure rather than robustly discriminating precise frequency differences (e.g., 7–8 Hz vs. 8–9 Hz). The authors should explicitly acknowledge this inherent limitation and clarify how their method accounts for or addresses this constraint.

The identified theta state switching rates (~0.1–1.2 Hz) closely match the temporal scale of task events, suggesting that these states might reflect task epoch or behavioral transitions rather than distinct, biologically meaningful oscillatory frequencies. The manuscript would benefit from explicitly analyzing and reporting the presence or absence of different theta states across distinct task epochs (odor sampling, running, waiting).

Later figures reveal prominent olfactory-driven beta bursts, reminiscent of those described by Walter Freeman (<https://pmc.ncbi.nlm.nih.gov/articles/PMC10196064/>). The authors should explicitly highlight and interpret these beta oscillatory features, integrating their occurrence and functional implications within their theta-state framework. Additionally, the exclusion of Rat 5, which exhibited pronounced olfactory-driven beta bursts, raises the possibility that this spectral discrepancy might reflect electrode placement differences within the hippocampus or nearby olfactory-related structures. Addressing electrode placement variability and anatomical factors explicitly would significantly enhance interpretability.

Figure 3 Concerns: Given the extensive number of statistical tests (neurons × time points), the use of Benjamini/Hochberg correction, while controlling the false discovery rate, may still lead to overly liberal identification of significant effects. A more stringent correction, such as Bonferroni, or explicit acknowledgment of potential limitations due to multiple comparisons would enhance rigor. Additionally, effect sizes are not clearly reported alongside p-values, making it challenging to assess the biological significance of statistically significant results. Furthermore, significant inter-subject variability raises concerns regarding the robustness and generalizability of the findings. Lastly, significant differences observed predominantly post-response may reflect motor, reward expectation, or arousal changes rather than purely cognitive or theta-driven network dynamics. Clarifying or controlling for these confounding factors would strengthen the interpretation.

In conclusion, while methodologically sound, the manuscript would benefit substantially from explicitly addressing these critical points, clarifying how the findings advance or differ from classical theta literature, examining or controlling for key confounds such as locomotion speed, spectral resolution limitations, electrode placement variability, and explicitly analyzing distinct behavioral epochs.

Reviewer #2

(Remarks to the Author)

This manuscript investigates how theta-band oscillatory states—defined by frequency and amplitude—modulate population-level neural responses during an odor-memory task in rats. The study addresses a timely and important question in systems neuroscience: how ongoing network states shape the encoding of behavioral information. The manuscript is well-written, methodologically thorough, and clearly presented. The use of both encoding (regression) and decoding approaches provides a strong framework to probe state-dependent neural coding. Overall, the authors have tackled a novel and underexplored research problem with rigor, and the findings will be of interest to a wide audience studying oscillatory dynamics and neural ensemble activity.

Below I provide a series of comments and questions, primarily minor, aimed at improving clarity and interpretability of the work:

Minor Comments

1. While the introduction highlights the role of oscillatory phase in modulating neural firing, it would benefit from a brief discussion of prior studies showing effects of amplitude and the relatively less-explored role of frequency. Framing the current study as addressing this gap would better situate its contribution.
2. The regression model's improvement is assessed relative to surrogate (shuffled) states. However, the manuscript does not report how the model performs with only firing rate (i.e., excluding state terms), or with randomized firing rates. Including these baselines would help disentangle the unique contribution of state-related information.
3. The notion of LFP frequency modulating neuronal activity is reminiscent of neuronal resonance and has been modeled by Mike Cohen as “frequency sliding”. I suggest adding discussion of these topics to better ground this manuscript in the existing literature.
4. The higher firing rate observed in the HPHT state could partially inflate LFP power due to spike bleed-through. A common method to control for this is to remove or zero out LFP segments around spike times and interpolate the signal (e.g., Jacobs et al., 2007 in a paper entitled “Brain oscillations control timing of single neurons activity in humans”). It would be helpful to clarify whether such preprocessing was performed, or how spike contamination was mitigated.
5. It is unclear how firing rates were normalized between HPHT and LPLT states. Clarifying this point is important, especially given that HPHT is associated with higher firing rates in most neurons.
6. Although the LPLT state occurs more frequently and is more associated with successful trials, regression analyses indicate that the HPHT state more strongly modulates the relationship between firing and behavioral outcome. This suggests HPHT may be more functionally informative, but the higher firing rates in HPHT could be a confound. Clarifying whether

decoding accuracy remains higher when firing rate is matched would strengthen this interpretation.

7. The authors use the first principal component of the LFP across all channels as the basis for state identification. However, it is unclear why they did not use the LFP recorded from the same electrode that isolated the spiking neuron—this would better match local field context to unit activity. Clarifying this methodological choice would be helpful.

8. Because both frequency and amplitude differ between states, it is difficult to determine whether the behavioral relevance is driven by one or both features. Since there were no "HPLT" or "LPHT" state combinations, this limits the interpretability of frequency vs. amplitude effects. The authors could acknowledge this limitation more explicitly.

9. Figure 4 shows that within-state decoding outperforms cross-state decoding, but it averages across both theta states. It would be informative to report decoding accuracy per state (i.e., LPLT-trained vs. HPHT-trained decoders) to assess whether one state supports more accurate behavioral decoding than the other.

10. The manuscript mentions that state reactivation occurred at 0.1–1.2 Hz, ruling out phase-locked switching. However, it is not clearly described how state periodicity was quantified. Please provide methodological detail on this computation.

In sum, this manuscript presents an elegant analysis of state-dependent ensemble coding using modern unsupervised modeling and decoding techniques. With some clarification and elaboration on the points above, the study will offer even greater insight into how frequency-specific network states shape population dynamics in the hippocampus.

Reviewer #3

(Remarks to the Author)

The paper by Masaracchia et al. entitled “frequency-specific theta states in the hippocampus modulate population activity with respect to behavioral context” is a statistical modeling study using previously recorded neural data on hippocampal CA1 activity in rats performing an odor-memory task. The main finding is that a decoder trained on neural spiking activity from many neurons to predict behavioral trial outcome needs to account for theta states to optimize prediction of behavioral outcome by neural activity. This is an interesting finding which is timely and important for understanding the important role of underlying brain states, on which neural dynamics correlated to behavior unfold. The paper is technically solid for the most part. However, some of the underlying assumptions, e.g., that two theta states were detected by searching for exactly two theta states, should be discussed more thoroughly. Moreover, the manuscript often lacks clarity in describing the correlational nature of the statistical modeling results. My major comments are detailed in the following sections.

Abstract:

The wording in the abstract implies that network states are primarily defined by rhythmic activity that then modulates in a causal way spiking activity. However, a network state is characterized by both rhythmic activity and spiking activity. The spiking activity is part of the network state and often part of the mechanism that gives rise to the rhythmic activity. In my opinion, it would be more accurate to describe the correlational nature of the relationship between theta states and ensemble activity.

The wording “ensemble pattern modulation” by theta state is misleading because it suggests that neural spiking activity is modulated in a causal way by theta state. However, the authors show that it is the prediction of trial outcome that is affected by neuron-state interactions. Even if the authors meant to say that neural activity statistically depends on theta state, I could not find any evidence supporting this claim. Does theta state help predict neural activity? In the results section, the authors write that “there is no clear correspondence between an active state and a specific subpopulation of neurons” and “that our analyses suggested that the states might not necessarily index distinct subpopulations of neurons”, indicating that theta state does not modulate ensemble patterns of neural activity. However, the next header reads “Theta states modulated neurons’ activity with respect to behavior”. Can the authors clarify? Also, reporting the model results for neuron-state interactions in a Table would help a lot.

If I understand the data correctly, the data demonstrate that the theta state needs to be accounted for when predicting behavior from neural activity. However, this conclusion is not stated clearly.

Introduction:

First sentence, last paragraph of Introduction: “[...] variations in frequency and amplitude affect the representation of behavioral information”. The wording implies that the authors tested a causal relationship, although only correlations were measured.

Last sentence, second paragraph of Introduction: “...suggesting that the LFP oscillatory patterns are not a mere epiphenomenon of neural activity, but serve a functional purpose.” The way the Introduction is written implies that the authors provide evidence for a causal effect of theta state on neural activity and behavior and evidence against theta state being a mere epiphenomenon. However, the nature of the modeling approach can only discover correlations, which means that the authors cannot distinguish whether the discovered statistical dependence reflects an epiphenomenon or a function. While the authors do not make an explicit claim, it is somewhat implicitly made in the tone and way the Introduction is written.

Second of last sentence of Introduction: “Our results show that the neural population activity differentially reflects success and failure depending on the theta state.” The wording is ambiguous because it can be interpreted as implying a causal relationship, where theta state was experimentally controlled/manipulated as the independent variable and correlation between neural population activity and trial outcome was tested as the dependent variable. However, the authors refer to

statistical dependence in a regression model. At this point in the Introduction, it is not completely clear what approach the authors have taken. It is not until the very end of the Discussion section that the authors finally address this important point that they do not claim an underlying unidirectional causal flow but instead are referring to statistical dependence and correlations. This disclaimer needs to be stated much earlier. It would have helped this reviewer a lot in avoiding confusion with respect to the authors' claims if this statement would have been placed in the Introduction section.

Results:

I found the order of Figures 1 and 2 confusing. One needs to understand how the theta states were defined (data presented in Figure 2) to understand the statistical models presented in Figure 1. Also, the paragraph associated with Figure 1 ends without explaining the results of the model.

Figure 1C: It would help if the authors would report the model formula and outcome (coefficients and p-values/confidence intervals) in a Table format.

Regarding the HMM with $k=2$ states: It is by design of the HMM that it identifies two states. Since the two states are at opposite poles of a spectrum (low power and low frequency vs. high power and high frequency), the question arises whether there is actually a continuum of states instead of two discrete states that continuously switch from one to another. This possibility should be discussed.

The authors frequently make statements about the population activity of neurons. However, it is unclear from reading the Results section how population activity was defined and included in the the model(s). Did the authors test population vectors based on firing rates? The Methods section provides more details, but it would help the reader to briefly summarize the way population activity was defined in the Results section.

Figure 2B: It's difficult to see which data are from the same or different rat. While the authors write that power and frequency were consistently higher in one state, statistical testing of this hypothesis would be appropriate.

Discussion:

First paragraph of Discussion: The authors write "These states did not appear to recruit different neurons but rather changed how the same population processed information". Please see my comments above on the discrepancy between this statement and other statements claiming that theta state modulates neural activity. In my opinion, a more accurate description of the main finding would be to say that the correlation of neural activity and trial outcome statistically depends on theta state.

Last sentence, second paragraph of Discussion: "Our findings show that amplitude and frequency of theta oscillations can signal different network states, which affect population-level activity". Do the authors mean activity or changes in activity/response dynamics? The authors write earlier that "there is no clear correspondence between an active state and a specific subpopulation of neurons". While there appears to be a significant neural activity*theta state interaction effect in predicting trial outcome, this is not the same as saying that theta state affects neural activity. Also, reporting the model results would be helpful.

Bottom of page 8, Discussion: statistical dependence instead of dependence?

Second last paragraph of Discussion: "(...) we do not claim an underlying unidirectional causal flow from one to the other: oscillations and neural firing are two properties of the same system that emerge together". I agree with this statement. However, the wording throughout the manuscript often suggests otherwise. There may also be a disconnect between the jargon

used by computational scientists and experimental scientists. E.g., "A modulates B" is usually interpreted as "a change in A causes a change in B" by experimentalists. I suggest to change the wording throughout the manuscript to make it clear that the authors refer to correlations between theta states and behavioral outcome-dependent changes in neural activity. Perhaps most importantly, this statement needs to be placed much earlier in the text (see my comment above).

Last paragraph of Discussion: "In conclusion, this study suggests that neural ensemble activity changes in context-specific ways depending on the state of the network as expressed by the theta oscillation, challenging static models of neural coding." In my opinion, this conclusion – in the way it is written – is not supported by the data. First, neural ensemble activity is part of the network state, and network state cannot be defined by theta state alone; second, "context-specific" is very vague and it is not clear what exactly is meant by context; third, it should be noted that it is a statistical dependence; fourth, neural ensemble activity does not change dependent on theta state. Instead, the correlation of neural activity and trial outcome changes dependent on theta state.

Reviewer #4

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Communications Biology initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

To the Authors:

The authors appeared to have evaded my primary concerns regarding movement and general behavioral outcomes, dismissing movement and locomotion as irrelevant. This claim was not met with data but with rhetoric. There was no analysis of head position, acceleration, licking, or the times of port entry or exit. Only rhetoric. I will focus my review on this issue rather than dedicating time to the entire manuscript, as I no longer believe the authors are engaging in good faith with my original comments.

We are left with a paper that uses a sophisticated HMM to tell the community: "Neurons do something (what exactly? "reconfigure" is a vague and terse description) different when animals are receiving a reward versus when they are not," a finding that is not novel and, in this context, is an expected consequence of behavior. Behavioral differences remain the most parsimonious explanation. While I can say the planets "reconfigure" around the sun, it does nothing to explain planetary orbits.

Despite the technical clarifications in the rebuttal, the authors have failed to address the central, fatal confounds that undermine the paper's core claim. Our major outstanding issues remain and are, in fact, strengthened by this review. They fall into three main categories.

Major Issue 1: The Fatal, Uncontrolled Behavioral Confound

The authors' entire claim of a cognitive "reconfiguration" rests on their ability to disprove all other simpler, behavioral explanations. They have failed to do this.

Our original review stated that the significant effects observed "predominantly post-response may reflect motor, reward expectation, or arousal changes".

The authors' entire rebuttal to this foundational criticism rests on a single, unsupported, and demonstrably false assumption:

"we first should discard the motor cause because the animals are not moving"

This assumption is in direct contradiction to the original methodological papers (Allen et al. 2016, Shahbaba et al. 2022) that this dataset is built upon. As confirmed by the provided text:

A correct response (Time 0) is immediately followed by a "water reward... delivered below the odor port."

An incorrect response (Time 0) is immediately followed by a "buzzing sound" and sequence termination.

After every trial, rats are "required to run to the end of the track opposite the odor port."

The post-response window (1-2 seconds after Time 0), which is the only period where a significant effect is found in Figure 4 for two of the animals, is therefore not a period of "attentive immobility". It is a period of guaranteed, massive, and stereotyped behavioral differences:

Condition A (Success): The neural data reflects active reward consumption (licking), reward expectation, and preparation to move (e.g., the Crespi effect, where confidence leads to faster transitions).

Condition B (Failure): The neural data reflects the absence of reward, the presence of an aversive sound, and preparation to move.

Unsurprisingly, theta and neural activity would be different in these two conditions. The authors' standardization for "gross firing rate" does nothing to control for this. The decoder is not learning a "reconfiguration" of a cognitive code; it is almost certainly just learning to distinguish the massive, non-cognitive neural and motor-artifact-driven patterns of "licking" versus "hearing a buzz."

This is a fatal behavioral confound that has been dismissed, not addressed.

Major Issue 2: The Central Finding is Statistically Weak and Selectively Reported

The authors' entire claim for "ensemble pattern reconfiguration" rests on Figure 4. This finding is statistically weak and unconvincing on multiple levels.

Selection Bias (N=3/4): The central finding is based on only 3 of the 4 subjects. The exclusion of Subject 4, while justified by the authors on a "technical" basis (low failed trials), is a form of selection bias. It presents a "significant" finding from an incomplete dataset.

No Cross-Subject Correction (Inflated Error): This is the most critical statistical flaw. The authors analyze each subject independently and claim significance at $p < 0.05$. By presenting three separate $p < 0.05$ tests as evidence for one general finding, they are inflating the family-wise error rate, giving a >14% chance that at least one subject is a false positive. Subject 3, with its sparse and weak significance, looks precisely like this Type I error. A robust finding would require a proper cross-subject correction (e.g., Bonferroni, $p < 0.05/3 = 0.017$), which would almost certainly render this N=3 finding an N=2 finding.

Arbitrary "Researcher Degrees of Freedom": The choice of a "25-timepoint" (100ms) cluster window with a "12-timepoint" stride is arbitrary and unjustified. This is a classic "researcher degree of freedom" that, in combination with Benjamini-Hochberg, can be tuned to find significance.

Major Issue 3: The "Reconfiguration" is Unsubstantiated Jargon

The paper's central claim of "ensemble pattern reconfiguration" is speculative jargon that is not demonstrated by the data.

Inference vs. Demonstration: The authors' method (Figure 4) provides only an inference that the patterns are different (i.e., the decoder fails when crossed). It provides zero demonstration of how these patterns are different.

No Visualization: At no point do the authors visualize the population code. They never describe or demonstrate what this "reconfiguration" is. Different cells? Different patterns of the same cells? We are never shown which neurons contribute differently, or what the "reconfigured" pattern even looks like.

The authors are using a vague term, "reconfiguration," to describe a difference, and that difference, as established in Issue 1, is almost certainly a behavioral artifact.

Conclusion

I did not care to continue reviewing this manuscript, as the authors did not engage in good faith with my original, primary comments and considered any further requests for them to defend their claims a waste of my time.

My concern stands. The authors have used a statistically weak (N=3), selectively reported, and arbitrarily windowed analysis to find a massive, uncontrolled behavioral confound (licking vs. no-reward). They have then obscured this finding with speculative jargon ("reconfiguration") that is not directly supported by their own evidence.

The central claim of the paper that they have found a "cognitive reconfiguration" is entirely unsubstantiated.

Reviewer #3

(Remarks to the Author)

The revised manuscript by Masaracchia et al. entitled "Frequency-specific theta states in the hippocampus reflect reconfiguration of population activity with respect to behavioural context" has substantially improved with respect to clarifying methodological aspects and underlying assumptions of the model and with respect to clarifying the conclusions of the study. The authors have addressed all my previous comments, and I have no further comments. As mentioned previously, the study is timely and of significance for the neuroscience community. The manuscript presents a very interesting finding, namely that power and frequency variations within theta oscillations are linked to the reconfiguration of neural firing at the network level so that brain state-dependent variations in population activity code for the same behavioral outcome. I support the publication of the manuscript in Communications Biology.

Reviewer #4

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Communications Biology initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Version 2:

Reviewer comments:

Reviewer #3

(Remarks to the Author)

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Response letter to reviewers' comments

The authors would like to sincerely thank the reviewers for their thorough review and helpful comments. We believe their insights have significantly contributed improving this manuscript. We have tried to address adequately all the reviewers' comments, providing clarifications, description of extra analyses and new results where appropriate. Our responses follow each comment in blue. The changes in the manuscript follow our responses in *red italics*.

Reviewers' comments:

Reviewer #1 (Remarks to the Author):

The manuscript by Masaracchia et al. provides an intriguing exploration of theta oscillations in hippocampal neural ensembles. While the use of data-driven methods to distinguish between two closely related theta states (LPLT and HPHT) is methodologically sophisticated, several critical concerns arise regarding the interpretation and novelty of these findings. I have made a list of these concerns below

- The described LPLT (~7–8 Hz, lower amplitude) and HPHT (~8–9 Hz, higher amplitude) states closely parallel the classical characterization by Vanderwolf (1969) of type I (movement-dependent) and type II (immobility or arousal-dependent) theta oscillations. It is problematic that the authors neither acknowledge nor cite Vanderwolf's foundational work or subsequent literature (e.g., Bland, 1986). Clarifying how their findings differ meaningfully from these classical concepts would strengthen the manuscript.

Thank you for your comment. We agree that Vanderwolf's, Bland's and others' work on hippocampal theta rhythms need to be acknowledged. However, we need to clarify that, while those works are in a spatial and locomotion-dependent context, the dataset and task we are investigating is non-spatial: no locomotion is involved, meaning that the rats perform most of the task in a constant spatial location (the description of the dataset can be found in Allen *et al.*, 2016, and Shahbaba *et al.*, 2022). Since the type I theta is classically observed only in connection with running, our two different theta states cannot be interpreted in this way. We have now expanded on this concept, discussing the relevance of the two types of theta within the scope of this study, in a paragraph in the Discussion section. We map our findings onto established theta theory, exploring possible sources and functional roles of our states, and clarify that a mechanistic interpretation remains tentative. Here follows the revised paragraph:

Even though the HMM that we used to detect network states is a purely data-driven model with no neurobiological assumptions or knowledge of the task, the two distinct state peaks within the theta band are reminiscent of the *two previously identified hippocampal theta modes, generated from regions projecting to the hippocampus, and related to different behavioural scenarios: Type I (~7–12 Hz, atropine-resistant) during locomotion and Type II (~4–7 Hz, atropine-sensitive) during attentive immobility*⁶⁰⁻⁶⁹. In⁴⁰, the authors speculated that the ventral dentate gyrus (DG), and in particular the ventral mossy cells (MC), may serve as a local theta generator or pacemaker. The

~~ventral MC could drive a medial septum independent theta rhythm in the dorsal hippocampus through their strong bilateral projection to this region⁴³. Thus, it is possible that the two theta states presented here have distinct origins: one coming from the medial septum⁴⁴, and the other from the ventral MC. Of note, the explanation suggested in⁴⁰ (and in REF1,2,3) had to do with changes of locomotion speed and novelty, which are different from the type of behaviour analysed here. But it is not unconceivable that the differential functions associated to these theta modes are more widespread than anticipated previously.~~ *In awake immobility, theta has been observed to transiently accelerate toward the upper band yet remaining pharmacologically of Type II (that is, sensitive to atropine)⁷⁰. The two data-driven states we observed here, however, were identified during a non-spatial odour-memory task without locomotion. Therefore, the most parsimonious interpretation is that they reflect Type II network regimes, in which coordinated shifts in theta frequency and power are associated with ensemble-level reconfiguration in reflecting trial outcome. The exact biological function of these network states, however, remains speculative. Mechanistic inference would require interventions to pinpoint the pathways driving the ensemble modulation, coupled with controlled task manipulations to establish functional relevance. Theta frequency also tracks internal states across behavioural contexts: motor preparation⁷¹, environmental novelty⁷², and higher arousal states such as risk-taking⁷³, anxiety⁷⁴, and social interaction⁷⁵. Further investigation would also be needed to explore how these states relate to cognitive processes such as task engagement, decision making or learning⁴⁷.*

- Theta amplitude, frequency, and neuronal firing rates are well-established correlates of locomotion speed. Given that the authors report higher firing rates and increased theta amplitude and frequency in their HPHT state, a critical omission is the lack of running speed measurements or controls. Without addressing this factor explicitly, the interpretation of their "ensemble pattern modulation" as a novel mechanism remains questionable.

We would like to emphasize that the task does not involve locomotion, and therefore running speed was neither measured nor relevant to the analysis. To make this explicit, we have now revised the paper, adding the keyword *non-spatial* to the already used *odour-memory* attribute to describe the task.

- The manuscript presents "ensemble pattern modulation" as a form of theta modulation distinct from classical phase-based modulation (e.g., phase precession). However, without explicitly analyzing theta phase relationships or convincingly demonstrating independence from classical theta phenomena (including phase precession or known frequency-amplitude correlations during locomotion), this claim of novelty appears overstated.

In our analyses, we had shown that the two HMM states have periodicity of the range of 0.1-1.2 Hz (**Figure 2D**), meaning that they occur with a period slower than theta frequency range. We interpreted this as the effect being unrelated to the phase of theta. To address this link explicitly, we have now computed the distribution of theta phases for each state and compared them. To do so, we first computed the theta phase of the signal across the whole recording session, then grouped the theta phase time course into two groups based on the state activation, obtaining two distributions of angles associated with the two states. We compared the distributions by means of the Watson U² test (analogous of the Mann-Whitney U test, but for circular data). We found that theta phase distribution is not significantly different

across states at least in two out of four rats, (see **Supplementary Figure 4** for a visualization of states fractional occupancy, FO, over theta phase bins). Together with the other states characterisation measures, such as lifetime and activation frequency, this suggests there is not a clear relation to theta phase. We have included these findings in the Results section, as:

*To explicitly test the link between states and theta phase, we statistically compared the distributions of theta phases between the two states (see **Methods** for more details). Our analysis revealed that the two states did not have a significantly different distribution over theta phase in two out of four subjects (Subjects 1 and 4, where the difference in mean value of the states' phase distributions was below 0.01π , and the difference in variance below 0.005π ; the p-value of the Watson U^2 test, corrected for multiple comparison with Benjamini-Hochberg method, was not significant), and small but significantly different theta phase distributions in the other two subjects (Subjects 2 and 3, with a difference in mean value of the distributions of 0.04π and 0.06π respectively, and difference in variance of 0.03π ; p-value of the Watson U^2 test, corrected for multiple comparisons with Benjamini-Hochberg method, was <0.0001). For visualization purposes, we show the states' FO across theta phase bins in **Supplementary Figure 4**. Overall, these results indicate that the states are not clearly related to the phase of theta, suggesting they are a different effect from phase precession.*

We have described this analysis in the Method section as follows:

*To explicitly link theta phase to the state time courses, we computed the theta phase of each state visit and compared the two states distribution in angles. To do so, we first extracted the phase from the LFP signal, band-pass filtered in theta band (4-12 Hz). We then grouped the theta phase time course into two groups, according to which state was active at each time point. This gave us two distributions of angles, one for each state, that we compared by means of the Watson U^2 test⁸⁵. This test is similar to the Mann-Whitney test but for circular data. To visualize the states distribution across theta phase, we divided the time courses in bins of $1/5^{\text{th}} \pi$ angle each, and computed the fractional occupancy (FO) of the states for each bin (**Supplementary Figure 4**).*

- The reported frequency distinction between LPLT and HPHT states is minimal (~ 1 Hz difference). While technically achievable using data-driven methods (HMM), the biological relevance and interpretive significance of such a minor frequency difference remain uncertain. The authors should clarify why these subtle spectral differences hold substantial functional relevance rather than reflecting minor variations within a continuous distribution.

Regarding the biological relevance of these differences, we would like to note that this is a simplified representation (an HMM representation), and that this effect may well be part of a continuous distribution instead of two categorical states. The fact that the HMM identifies two states instead of collapsing to one single state, however, indicates that this variability does exist in the data. Furthermore, the changes we observed in their relation to population activity are an indication of their functional relevance—even if we cannot fully unravel the mechanisms of the effect. We emphasize this point in Discussion:

This does not exclude that the effect we have uncovered is supported by variations within a continuous distribution, rather than by categorical distinct states. In any case, it still reveals a biologically relevant function in that changes within power and frequency of oscillatory patterns are robustly linked to nuanced changes in population activity in its relation to behaviourally meaningful variables.

Statistical testing on the differences between the states was requested by reviewer 3, and we have now added a more detailed spectral characterization of the states, including statistical analysis of the difference between the two states' power and frequency distributions (in **Supplementary Figure 2**).

-The authors state clearly that most neurons fire during both theta states, differing only quantitatively in firing rates. This finding suggests the described states modulate neural activity primarily through known amplitude-frequency correlations rather than recruiting functionally distinct neuronal subpopulations.

We want to stress that indeed, on the one hand the same neurons fire during both states (suggesting that states do not recruit different subpopulations of neurons), but on the other hand, the difference in information representation by the neural population cannot only be explained by the difference in amount of firing between the two states. We control for this confound by standardizing the neural activity per trial and per timepoint, which makes the neural population firing rate have the same mean and standard deviation for both states (see Methods, **Data preprocessing** subsection).

- Given the extremely short analysis window (32 ms), the spectral resolution inherent in the HMM-TDE approach requires clarification. Can the authors explicitly state their actual frequency resolution under these conditions? Moreover, we recommend benchmark testing of their HMM-TDE parameters on simulated data designed to generate known theta harmonics similar to those observed by Buzsaki in the 1980s (later in work by Czurko and Hirase) and shape changes as described by Terrazas, and McNaughton (2005). Such testing would help validate whether the observed states represent distinct oscillatory modes or artifacts of their analysis method. Importantly, no spectral decomposition method, including HMM-TDE, circumvents the fundamental limitations imposed by time-frequency tradeoff windows. Thus, it appears likely that the HMM-TDE approach is sensitive primarily to waveform shape, amplitude patterns, or transient temporal structure rather than robustly discriminating precise frequency differences (e.g., 7–8 Hz vs. 8–9 Hz). The authors should explicitly acknowledge this inherent limitation and clarify how their method accounts for or addresses this constraint.

Thank you for raising this concern. Just as a clarification, please note that our parameter $\text{lags}=7$ corresponds to a window of 15 time points, from -7 to +7 around each timepoint of focus, overall covering a 60ms window of the signal (sampling rate of 250Hz). We have clarified this in the HMM description in the Methods section:

- the lags L : *defining* the length of the window in which to estimate the data autocovariance (*from -L to +L around each time point*), which affects the spectral scales that the states can capture;

We would also like to clarify the working mechanisms of the HMM-TDE: at each timepoint, the HMM focuses on the autocovariance of the signal within the specified window determined by the *lags* parameter. The autocovariance contains information about the frequency of the signal. In this regard, a window of 60ms can encode changes in frequency whose period is much slower than 60ms. In the paper, we have cited our previous work (Masaracchia et al., 2023), where a thorough analysis of the sensitivity of the HMM to various aspects of the signal was conducted. There, we analysed, using simulations, the frequency resolution of HMM-TDE while varying HMM parameters as well as data characteristics. We found that the TDE observational model alone has a frequency resolution higher than 1Hz in conditions relatable to those of this experiment (i.e., similar data characteristics and same HMM settings; see Masaracchia et al., 2023, Supplementary Figure 1). Further, we included in this manuscript an analysis to explicitly test the HMM frequency resolution on data more tailored to the characteristics of the data we have here. We run the HMM on a simulated signal varying only in frequency, within 7 to 9 Hz. We show that the HMM finds two states that differ in their main frequency peak by ~1 Hz. We have added these results as **Supplementary Figure 3** and mentioned in the Results section:

*As a sanity check, we tested the sensitivity of the HMM to small spectral changes. To do so, we run the HMM on simulated brain data, varying only in frequency between 7 to 9 Hz. The model successfully distinguished two states based on these small frequency differences (**Supplementary Figure 3**).*

This analysis is described in the Methods section:

We explicitly examined the HMM-TDE ability to capture finer details in frequency. We used the same pipeline as in Masaracchia et al. (2023)⁵⁴ to simulate data varying only in frequency between 7 and 9 Hz, and run the HMM-TDE with the same settings as in this experiment (i.e., $K=2$, $L=7$, $\delta=1\text{bn}$). We characterised the states by computing their power spectrum.

To address the potential problem of waveform shape changes in the signal, we have conducted an HMM analysis on a simulated signal varying in shape (resembling the shape changes described by Terrazas et al., 2005) and frequency. This analysis was meant to test whether the HMM states identification would be preferentially driven by waveform shape or by frequency. Our results show that the HMM would be more sensitive to cycle-average frequency in a case where the signal changed both in shape and cycle-average frequency, independently. We have added the simulations and the results of this analysis as **Supplementary Figure 6**.

*To test whether the HMM states could be influenced by potential shape changes in the data (as seen in hippocampal theta literature⁵⁵), we run the HMM on simulated data, varying both in shape and frequency. Our results show that the HMM states would be mainly driven by cycle-average frequency changes in a case where the signal changed both in shape and cycle-average frequency, independently (**Supplementary Figure 6**).*

We describe the data simulation process in a new subsection of the Methods section:

Simulating oscillations changing in shape

To test whether the HMM states could be driven by changes in the oscillation shape, we generated synthetic oscillatory signals using a process based on Genephys, a generative model of electrophysiological data (publicly available at <https://github.com/vidaurre/genephys>⁸⁶), in which frequency f , waveform shape modulation m , and amplitude a evolve as independent autoregressive processes of order one around prescribed targets. This structure introduces slow stochastic wander while maintaining temporal continuity to reproduce the variability observed in neural recordings as follows:

$$\begin{aligned} f_t &= b_f f_{t-1} + \varepsilon_{f,t}, & \varepsilon_{f,t} &\sim \mathcal{N}(0, \sigma_f) \\ m_t &= b_m m_{t-1} + \varepsilon_{m,t}, & \varepsilon_{m,t} &\sim \mathcal{N}(0, \sigma_m) \\ a_t &= b_a a_{t-1} + \varepsilon_{a,t}, & \varepsilon_{a,t} &\sim \mathcal{N}(0, \sigma_a) \end{aligned}$$

Where b_f, b_m and b_a are the autoregressive weights of the process, and the various ε are noise contributions, normally distributed with zero mean and corresponding standard deviation σ . To change the shape, we modulated the phase of the sinusoid ϕ according to a nonlinear cumulative update, with the carrier frequency determining the period and the modulation parameter m distorting cycle symmetry:

$$\phi_t = \phi_{t-1} + 2\pi \frac{f_t}{F_s} (1 + m_t \cos \phi_{t-1})$$

Here F_s is the signal sampling rate. Thus, when $m = 0$, the signal reduces to a sinusoid. Non-zero values of m accelerate phase progression during one half-cycle and decelerate it during the other, producing asymmetric waveforms without altering the period.

The observed signal was obtained by modulating a sinusoid with these processes and adding Gaussian measurement noise:

$$x_t = a_t \sin(2\pi f_t + \phi_t) + \varepsilon_{x,t}, \quad \varepsilon_{x,t} \sim \mathcal{N}(0, \sigma_x)$$

To build continuous time series, we concatenated segments with randomized durations with a minimum bound of 10 seconds and rescaled to a fixed total duration of 5 minutes, ensuring heterogeneous segment lengths. We then specified frequency variation (7 - 9 Hz), shape variation with $m \in [0, 1]$, combining both independently across time. This design allowed us to control for frequency shifts, shape modulations, and their interaction on downstream analyses.

- The identified theta state switching rates (~0.1–1.2 Hz) closely match the temporal scale of task events, suggesting that these states might reflect task epoch or behavioral transitions rather than distinct, biologically meaningful oscillatory frequencies. The manuscript would benefit from explicitly analyzing and reporting the presence or absence of different theta states across distinct task epochs (odor sampling, running, waiting).

Thank you for this important insight. We have now visualized the states distribution across different task-relevant moments/epochs and added them to **Supplementary Figure 11**. We find similarities and consistencies across subjects in the states' dynamics, meaning that the LPLT state is still the most active during all task-relevant moments, and that the distribution of the states is constant during the intertrial time (when the subjects are simply waiting). We emphasise that in our regression analyses we accounted for this imbalance, and that the results are above and beyond this. We have added these to the Results section:

*We also examined the states distributions in relation to relevant task moments, such as odour sampling, response and inter-trial periods. We found consistent patterns of the states' occurrences in the various moments across subjects, i.e. the LPLT state is the most active during all task-relevant moments, and the distribution of the states is constant during the intertrial time (**Supplementary Figure 11**).*

-Later figures reveal prominent olfactory-driven beta bursts, reminiscent of those described by Walter Freeman (<https://pmc.ncbi.nlm.nih.gov/articles/PMC10196064/>). The authors should explicitly highlight and interpret these beta oscillatory features, integrating their occurrence and functional implications within their theta-state framework.

Thank you for this important observation. Our focus on theta was driven by theory (within the role of the hippocampus in processing temporal sequences) and by its spectral prominence; but differences may also exist in other frequencies. We have added a figure (**Supplementary Figure 1**) to show the relative difference in the power spectrum between the two states (i.e., the power of HPHT divided by the power of LPLT state, across the broadband frequency range), where differences in theta are the most prominent. We have also added a consideration in Results:

*As the HMM was run on the broadband signal, the states' power spectra comprised multiple frequency bands. We visualized the ratio between the two states' power spectra (**Supplementary Figure 1**) to check for proportional differences in the power spectrum across the broadband frequency range and confirmed that the largest difference between the two states was within the theta frequency range, consistently across all subjects.*

We have now included a consideration about the presence of oscillatory features within other frequency bands, including those in beta, consistent with the literature on odour tasks in rodents, in the Discussion:

Our focus on the theta band was dictated by its known relevance in the hippocampus, and because it is prominent in the data. However, the data showed oscillations across the broadband frequency spectrum, particularly in both beta and gamma bands. This is consistent with research on odour-related tasks (typically in the beta band ⁷⁶) and on temporal processing (typically associated to gamma ⁷⁷).

-Additionally, the exclusion of Rat 5, which exhibited pronounced olfactory-driven beta bursts, raises the possibility that this spectral discrepancy might reflect electrode placement differences within the hippocampus or nearby olfactory-related

structures. Addressing electrode placement variability and anatomical factors explicitly would significantly enhance interpretability.

Thank you for this important observation. As pointed out, indeed the difference in spectral properties of subject 5 may well reflect differences in the placement of the electrodes. However, because there is no detailed graphical or positional information about electrode placement in the dataset (see Allen *et al.*, 2016, where the authors write that “tetrodes were slowly advanced over a period of ~3 weeks while monitoring established electrophysiological signatures of the CA1 pyramidal cell layer”), it is impossible for us to have a definitive answer to that. We have now added this as potential explanation for the difference in power spectrum of subject 5 in the Methods section:

This important difference, *possibly due to a difference in electrodes placement (see Allen et al., 2016⁵⁰ for details in the electrode placement procedure)*, induced us to exclude subject 5 from our main analyses for consistency.

- Figure 3 Concerns: Given the extensive number of statistical tests (neurons × time points), the use of Benjamini/Hochberg correction, while controlling the false discovery rate, may still lead to overly liberal identification of significant effects. A more stringent correction, such as Bonferroni, or explicit acknowledgment of potential limitations due to multiple comparisons would enhance rigor.

We respectfully disagree on this. Having “overly liberal identification of significant effects” means having an invalid test or correction. This is not the case. False discovery rate (FDR) using the Benjamini–Hochberg procedure is a standard approach (and statistically valid) for large-scale multiple testing (Benjamini & Hochberg, 1995). Further, given that our analysis involves thousands of tests, Bonferroni correction “is at the expense of a drastic reduction in sensitivity” (Maris, “Statistical testing in electrophysiological studies”, 2011), severely inflating false negatives. This is even more the case in time series, where contiguous tests are correlated. Thus, FDR provides a better balance between Type I and Type II errors (Genovese, Lazar, & Nichols, 2002).

-Additionally, effect sizes are not clearly reported alongside p-values, making it challenging to assess the biological significance of statistically significant results. Furthermore, significant inter-subject variability raises concerns regarding the robustness and generalizability of the findings. Lastly, significant differences observed predominantly post-response may reflect motor, reward expectation, or arousal changes rather than purely cognitive or theta-driven network dynamics. Clarifying or controlling for these confounding factors would strengthen the interpretation.

Thank you for your comments. The effect sizes are reported along with the results and the p-values in both the prediction and decoding analyses. For the prediction analysis, these are in **Figure 3A**, while the corresponding p-values are in **Figure 3B**. For the decoding analysis, effect size is reported together with the decoding analysis in **Figure 4** (right y-axis), and the p-values are reported as stars ($p < 0.05$). We realized that the connection between the plots in **Figure 3** might not be straight forward, so we have now clarified this:

Figure 3A shows the effect size of including theta states in the prediction of trial outcome from neural activity (plotted the prediction error averaged across neurons, using the HMM states and randomly sampled states for comparison).

As for the generalizability of the findings, we acknowledge that the availability of LFP and neural data is a limiting factor in this experiment. We believe that testing for this effect in other datasets and in different animals would be the scope of future work. We have now added a comment regarding this limitation in the Discussion section:

We acknowledge the small number of animals analysed here as a limitation in this study. Future research may explore whether these findings generalize to other animals, brain regions, or tasks.

Regarding the potential cognitive/physiological or motor effects that might drive the changes, we first should discard the motor cause because the animals are not moving. But, of course, the spectral properties of the LFP are one level of description, which will map to other levels of description—in this case, cognitive or physiological processes. In other words, these are not confounds, they are simply the other side of the coin: no meaningful difference in theta would exist without a cognitive or physiological counterpart.

Reviewer #2 (Remarks to the Author):

This manuscript investigates how theta-band oscillatory states—defined by frequency and amplitude—modulate population-level neural responses during an odor-memory task in rats. The study addresses a timely and important question in systems neuroscience: how ongoing network states shape the encoding of behavioral information. The manuscript is well-written, methodologically thorough, and clearly presented. The use of both encoding (regression) and decoding approaches provides a strong framework to probe state-dependent neural coding. Overall, the authors have tackled a novel and underexplored research problem with rigor, and the findings will be of interest to a wide audience studying oscillatory dynamics and neural ensemble activity.

Below I provide a series of comments and questions, primarily minor, aimed at improving clarity and interpretability of the work:

Minor Comments

1. While the introduction highlights the role of oscillatory phase in modulating neural firing, it would benefit from a brief discussion of prior studies showing effects of amplitude and the relatively less-explored role of frequency. Framing the current study as addressing this gap would better situate its contribution.

Thank you for this comment. We have now added a brief paragraph in the Introduction summarizing the main focus of existing literature on the modulation of neural firing by oscillations, and including the less explored role of frequency in this relation:

Existing literature targeting the relation between oscillations and neural firing has traditionally focused on the modulation of neural firing rate by the phase and amplitude of oscillations, across animals, brain regions, frequency bands and tasks³²⁻³⁸. Although amplitude and frequency of oscillations are known to be related³⁹, less focus has been devoted to the link between frequency of oscillations and neuronal firing⁴⁰⁻⁴². Existing work in this regard investigated fluctuations in peak frequency of oscillations and their effect in modulating neural excitability⁴³.

2. The regression model's improvement is assessed relative to surrogate (shuffled) states. However, the manuscript does not report how the model performs with only firing rate (i.e., excluding state terms), or with randomized firing rates. Including these baselines would help disentangle the unique contribution of state-related information.

Thank you for this comment. We realized that this important information was missing and added a figure, showing the decoding accuracy of success from neural activity without considering the states and when shuffling neural activity. We have added it to **Supplementary Figure 13**. Decoding accuracy from neural activity is generally high, while the decoding from shuffled neural activity is, as expected, at the random chance level. We refer to these results in the Result section:

*For comparison, in **Supplementary Figure 13** we also show standard accuracy and shuffled accuracy (i.e., from the shuffled neural activity). Decoding accuracy from neural activity is generally high, while the decoding from shuffled neural activity is, as expected, at the random chance level.*

We have also briefly elaborated on the results in **Supplementary Figure 13** in the Results section “**State-related patterns of trial outcome representation share overlapping characteristics**”, explaining why we believe our state-conditioned decoding is powerful even when standard accuracy is high. The paragraph now reads:

As observed, even when the within-state decoding accuracy was significantly higher than the cross-state, the latter remained relatively high. *Moreover, standard accuracy (i.e., accuracy of the decoding from the neural population without considering the states) was also comparable to that of within-state decoding.* This is due to similarities between the two states' neural ensemble activity patterns. If the patterns were entirely distinct, cross-state accuracy would be at chance level, *and standard accuracy would also be poor*; while if they were nearly identical, cross-state accuracy would match that of within-state *(and standard accuracy would be higher than both)*. [...]

3. The notion of LFP frequency modulating neuronal activity is reminiscent of neuronal resonance and has been modeled by Mike Cohen as “frequency sliding”. I suggest adding discussion of these topics to better ground this manuscript in the existing literature.

We agree that our findings are consistent with Mike Cohen's view on the potential functional role of fluctuations in frequency of oscillations. We have now added a paragraph in the Discussion section where we frame our contribution to this topic in literature:

Our findings link the activity of single neurons and oscillations at the network level, contributing to the literature that aims at bridging different spatial scales in the brain.

Our conclusions are consistent with the concept of “frequency sliding” in the LFP data⁴³, i.e. frequency of oscillations not constant in time, but “sliding” within a small range during cognitive processes and resting state, affecting the spiking threshold of single neurons. In Cohen (2014), the author speculated that frequency sliding could serve functional purposes, as the slower frequencies could induce faster but more variable spiking time, and the faster frequencies would instead trigger slower but more precise firing. In this view, our work might help shed light in the functional mechanisms used for information encoding in the brain.

4. The higher firing rate observed in the HPHT state could partially inflate LFP power due to spike bleed-through. A common method to control for this is to remove or zero out LFP segments around spike times and interpolate the signal (e.g., Jacobs et al., 2007 in a paper entitled “Brain oscillations control timing of single neurons activity in humans”). It would be helpful to clarify whether such preprocessing was performed, or how spike contamination was mitigated.

Thank you for this important consideration. However, the approach described in Jacobs et al. (2007) could in this case insert artifacts in the signal that, in our experience, can disturb the HMM analysis. Instead, we opted for a simpler approach. We used a sharp low-pass filter with cut-off frequency at 100Hz. We describe this preprocessing in the Methods section:

To control for potential spike-LFP bleed through effects, we designed a sharp low-pass digital filter with a Kaiser window⁸⁰ to attenuate by 100% the power on the LFP signal above 100 Hz, within a stopband of 1 Hz. The low-pass filtered data $z_{<100}$ was then used for a parallel analysis and results were compared with the minimally preprocessed data z .

We then run the HMM on the low-pass filtered signal and compared the states time courses and the states characteristics to those obtained from the minimally preprocessed signal. The states characteristics in terms of power spectrum remained identical. To compare the states time courses (STC), we computed the correlation between the two different STC (from minimally preprocessed signal and from low-pass filtered signal), across the whole data, for 10 independent HMM runs. We report the correlations in **Supplementary Figure 7**, where we also show some visual examples of the different STC. We consistently found high correlations (correlation coefficient between 0.85 and 0.9). We add this consideration to the Results section:

*To check for spikes bleeding into the LFP that could potentially affect the analysis⁵⁶⁻⁵⁸, we performed a sanity check where we used a sharp low-pass filter at 100 Hz (see **Methods** for more details) to remove the potential effect of spike waveforms on the LFP⁵⁹. We then run the HMM on the low-pass filtered data and confirmed that the states’ power spectra remained identical to those from the minimally preprocessed data. We compared the HMMs by computing the correlation across the whole recordings between the state time courses obtained from the minimally preprocessed data and those obtained from the low-pass filtered data, across 10 independent HMM runs. We report the correlations, consistently between 0.85 and 0.9 across subjects, in **Supplementary Figure 7**, together with two visual examples of the two correlated state time courses, for different correlation measures. The high correlation values suggested that spike contamination effect did not affect our analyses. We later*

confirmed this by running the decoding analyses on the low-passed filtered data and comparing with the results on the minimally preprocessed signal.

We also conducted state-conditioned decoding with the HMM states obtained from the low-pass filtered data. The decoding results are not qualitatively different from those performed on the unfiltered data. We report them in **Supplementary Figure 14**, and describe them in the Results section:

*As a sanity check, we also performed state-conditioned decoding on the HMM states ran on the low-pass filtered LFP signal (see above), confirming that potential spike contamination did not interfere with ensemble pattern reconfiguration effect. We report these decoding results in **Supplementary Figure 14**, not qualitatively different from the original decoding results (those reported in **Figure 4**).*

5. It is unclear how firing rates were normalized between HPHT and LPLT states. Clarifying this point is important, especially given that HPHT is associated with higher firing rates in most neurons.

We describe the way we standardize the data in the Methods section as: activity is standardized across neurons, per trial and per timepoint. This means that for each timepoint and for each trial we consider the whole population activity, we subtract its mean and divide by its standard deviation. This way, we maintain the relationship between the neurons, but we rescale the overall activity at each timepoint and each trial. And since the states' activations are defined per timepoint and per trial, in this way each state activation will see the same mean and standard deviation in neural population activity. We have now added a clarifying sentence in the Methods section to make this connection explicit:

The data were then standardised per trial and per timepoint, across neurons, so that the population activity had the same mean and standard deviation for each trial and for each time point. *This way, we maintained the relationship between the neurons, but we rescaled the overall activity at each timepoint and each trial.* We did this so that the neurons activity was not trivially linked to the states activation (since higher average firing rate occurs when the state HPHT is active).

We also added a sentence in the Results section to clarify that the standardization of the neural population activity was done before all our prediction and decoding analyses:

As before, and for all our subsequent prediction and decoding analyses, we standardised the neural population activity across neurons, for each trial and time point, such that regardless of which state is active the gross amount of firing is the same for all trials.

6. Although the LPLT state occurs more frequently and is more associated with successful trials, regression analyses indicate that the HPHT state more strongly modulates the relationship between firing and behavioral outcome. This suggests HPHT may be more functionally informative, but the higher firing rates in HPHT could be a confound. Clarifying whether decoding accuracy remains higher when firing rate is matched would strengthen this interpretation.

As described in Methods, we standardised the firing rate per trial and per time point, to equate the mean firing rate of neurons, thereby eliminating the confound of having on average higher firing rates for the HPHT state.

As for the interpretation of the regression analysis results, we realized that our interpretation of **Figure 3C** was problematic. Here, we predict trial outcome (Y) from three regressors (X, X G and G). The weights of the X G term are not easily interpretable in terms of their sign, i.e. it is not straightforward to interpret the sign as evidence of one state being more important than the other due to possible cancelling effects between the regressors. To avoid leading to confusion, we have deleted **Figure 3C** to the final version of this manuscript.

Regarding which state might be more functionally informative, we point at the class-balanced, cross-validated, per-state decoding accuracy, that we have now added in **Supplementary Figure 12** (as requested in comment 9 below), with the caveat that their interpretation should also be approached carefully due to the different number of trials available for each state.

7. The authors use the first principal component of the LFP across all channels as the basis for state identification. However, it is unclear why they did not use the LFP recorded from the same electrode that isolated the spiking neuron—this would better match local field context to unit activity. Clarifying this methodological choice would be helpful.

Our aim was to capture a population-level effect in information processing linked to the oscillatory states of the LFP signal. Linking each group of neurons to their specific LFP electrode would restrict the analysis to the activity of a small subset of neurons at a time and miss the population-level aspect of the information representation process. As noted in Methods, the LFP signals across electrodes are anyway highly correlated, and the first PC captured most variance. We have now specified for each subject the percentage of the explained variance of the first PC in the Methods section:

We then kept only the first principal component **z**, with shape (number of timepoints by one), which explained between 60% to 75% of the variance in each subject (*subject 1: 65%, subject 2: 70%, subject 3: 60%, subject 4: 63%, subject 5: 75%*).

8. Because both frequency and amplitude differ between states, it is difficult to determine whether the behavioral relevance is driven by one or both features. Since there were no "HPLT" or "LPHT" state combinations, this limits the interpretability of frequency vs. amplitude effects. The authors could acknowledge this limitation more explicitly.

This is true. We have now added a note in the Discussion section mentioning this as a limitation in our study:

It is important to note that, since we have not identified frequency and power as independent features of our states, we cannot disentangle them, implying we are not able with this analysis to further describe whether either one or both features are the main driving characteristics of this effect.

9. Figure 4 shows that within-state decoding outperforms cross-state decoding, but it averages across both theta states. It would be informative to report decoding accuracy per state (i.e., LPLT-trained vs. HPHT-trained decoders) to assess whether one state supports more accurate behavioral decoding than the other.

We added a figure showing each state's accuracy decoding success, as well as standard accuracy (i.e., from the neurons without considering the states). We show this in [Supplementary Figure 12](#), and describe it in [Results](#):

*We visualized each individual state's accuracy, and standard accuracy (i.e., from all the trials used in the states decoding, not considering the states information), in **Supplementary Figure 12**. As observed, the accuracy changes between states. It is however important to point out that because the number of trials differs between states, these plots remain descriptive.*

10. The manuscript mentions that state reactivation occurred at 0.1–1.2 Hz, ruling out phase-locked switching. However, it is not clearly described how state periodicity was quantified. Please provide methodological detail on this computation.

We have now clarified how we computed the states periodicity in [Methods](#):

This was done by computing the power spectrum of the state time courses, i.e. the probability of each state to be active at any time point. This measure shows the frequency of activation of the states (and because the states are exclusive, they have the same activation frequency).

We have also added an analysis that links explicitly states and theta phase as per reviewer 1's request — see [Supplementary Figure 4](#).

Reviewer #3 (Remarks to the Author):

The paper by Masaracchia et al. entitled “frequency-specific theta states in the hippocampus modulate population activity with respect to behavioral context” is a statistical modeling study using previously recorded neural data on hippocampal CA1 activity in rats performing an odor-memory task. The main finding is that a decoder trained on neural spiking activity from many neurons to predict behavioral trial outcome needs to account for theta states to optimize prediction of behavioral outcome by neural activity. This is an interesting finding which is timely and important for understanding the important role of underlying brain states, on which neural dynamics correlated to behavior unfold. The paper is technically solid for the most part. However, some of the underlying assumptions, e.g., that two theta states were detected by searching for exactly two theta states, should be discussed more thoroughly. Moreover, the manuscript often lacks clarity in describing the correlational nature of the statistical modeling results. My major comments are detailed in the following sections.

- The wording in the abstract implies that network states are primarily defined by rhythmic activity that then modulates-in a causal way-spiking activity. However, a network state is characterized by both rhythmic activity and spiking activity. The

spiking activity is part of the network state and often part of the mechanism that gives rise to the rhythmic activity. In my opinion, it would be more accurate to describe the correlational nature of the relationship between theta states and ensemble activity.

Thank you for this comment. We have now changed key words both in the title and abstract, trying to avoid any implication of causal effects, in favour of the correlational aspect of this effect. The term *modulation* has now been changed with *reconfiguration*, phrases like *are important for* have now become *are correlated with*. The title of the paper now reads:

Frequency-specific theta states in the hippocampus *modulate reflect reconfiguration of* population activity with respect to behavioural context

- The wording “ensemble pattern modulation” by theta state is misleading because it suggests that neural spiking activity is modulated in a causal way by theta state.

We have now changed the word *modulation* with *reconfiguration*.

-However, the authors show that it is the prediction of trial outcome that is affected by neuron-state interactions. Even if the authors meant to say that neural activity statistically depends on theta state, I could not find any evidence supporting this claim. Does theta state help predict neural activity? In the results section, the authors write that “there is no clear correspondence between an active state and a specific subpopulation of neurons” and “that our analyses suggested that the states might not necessarily index distinct subpopulations of neurons”, indicating that theta state does not modulate ensemble patterns of neural activity. However, the next header reads “Theta states modulated neurons’ activity with respect to behavior”. Can the authors clarify? Also, reporting the model results for neuron-state interactions in a Table would help a lot.

We acknowledge our wording created confusion. Our results show that *the states don’t index distinct subpopulations of neurons*. Indeed, all neurons fire during both states (as shown in **Supplementary Figure 8**). However, this does not prevent a modulation of ensemble patterns of neural activity at some level. Our analyses show that the overall population patterns change based on which state is active. This implies that all the neurons represent success and failure in a way when one state is active, and differently when the other state is active. In our view, there can be multiple ways in which theta states could be related to neural representations. On one side there is the possibility that theta states are linked to different subsets of neurons, representing different trial outcomes. However, this scenario does not fit our data. Hence, we infer that the theta states reflect differences in patterns of activity within the same population of neurons. We have now changed the heading of the following subsection to ***Theta states reflect neural activity reconfiguration with respect to behaviour***, to be more explicit and avoid any potential implication of causal relationships.

- If I understand the data correctly, the data demonstrate that the theta state needs to be accounted for when predicting behavior from neural activity. However, this conclusion is not stated clearly.

Indeed, we aimed at showing that accounting for the theta states improves the prediction of trial outcome from neural activity. We have now clarified this concept and changed the heading of the subsection to reflect this conclusion:

[...] our interaction analysis showed that *accounting for the theta states improves the prediction of trial outcome from neural activity, meaning that* the two theta network states ~~significantly affect~~ *reflect changes in* how single neurons in a population fire in response to trial outcome, suggesting population-level effects.

Introduction:

- First sentence, last paragraph of Introduction: “[...] variations in frequency and amplitude affect the representation of behavioral information”. The wording implies that the authors tested a causal relationship, although only correlations were measured.

We have now replaced the word *affect* with *are correlated with*.

- Last sentence, second paragraph of Introduction: “...suggesting that the LFP oscillatory patterns are not a mere epiphenomenon of neural activity, but serve a functional purpose.” The way the Introduction is written implies that the authors provide evidence for a causal effect of theta state on neural activity and behavior and evidence against theta state being a mere epiphenomenon. However, the nature of the modeling approach can only discover correlations, which means that the authors cannot distinguish whether the discovered statistical dependence reflects an epiphenomenon or a function. While the authors do not make an explicit claim, it is somewhat implicitly made in the tone and way the Introduction is written.

We have now tempered the words as [...] *and suggesting that the LFP oscillatory patterns may reflect a functional purpose*.

We have also rephrased wordings throughout the paper aiming to avoid ambiguity or implications of causation.

- Second of last sentence of Introduction: “Our results show that the neural population activity differentially reflects success and failure depending on the theta state.” The wording is ambiguous because it can be interpreted as implying a causal relationship, where theta state was experimentally controlled/manipulated as the independent variable and correlation between neural population activity and trial outcome was tested as the dependent variable. However, the authors refer to statistical dependence in a regression model.

We understand the raised concern, and we have now substituted *depend on the theta state* with *whether one or the other theta state is active* to avoid any ambiguity.

-At this point in the Introduction, it is not completely clear what approach the authors have taken. It is not until the very end of the Discussion section that the authors finally address this important point that they do not claim an underlying unidirectional causal flow but instead are referring to statistical dependence and correlations. This disclaimer needs to be stated much earlier. It would have helped this reviewer a lot

in avoiding confusion with respect to the authors' claims if this statement would have been placed in the Introduction section.

We acknowledge the confusion that our phrasing can generate, and we have now added the following to the Introduction:

Of note, our analyses do not necessarily establish a causal link from the LFP to the spikes (or the other way around), but an effect across scales (meso-scale in the LFP signal and microscale in the neural spikes).

Results:

I found the order of Figures 1 and 2 confusing. One needs to understand how the theta states were defined (data presented in Figure 2) to understand the statistical models presented in Figure 1. Also, the paragraph associated with Figure 1 ends without explaining the results of the model.

Thank you for this comment. Ideally, Figure 1 should only be used as a visual reference for the methods, independently on the results of each analysis. How the states are defined is a question concerning the working mechanisms of the HMM, which for simplicity we have schematized as a black box in the figure, because that should not affect the regression and decoding analyses. We have now clarified in the text how this figure should be approached:

*We have introduced **Figure 1** to schematically visualize our method pipeline, that we summarize below. This is meant as a reference to aid the understanding of the method, rather than as a figure reporting results.*

- Figure 1C: It would help if the authors would report the model formula and outcome (coefficients and p-values/confidence intervals) in a Table format.

We have now added the model formula together with its schematic visualization in **Figure 1C**. All the models' results and corresponding p-values are reporting in the result figures. These are **Figure 3** for the regression analysis and **Figure 4** for the state-conditioned decoding.

- Regarding the HMM with $k=2$ states: It is by design of the HMM that it identifies two states. Since the two states are at opposite poles of a spectrum (low power and low frequency vs. high power and high frequency), the question arises whether there is actually a continuum of states instead of two discrete states that continuously switch from one to another. This possibility should be discussed.

Thank you for pointing this out. This concern has been partially raised also by reviewer 1 and we have now added a paragraph in the Discussion session clarifying this:

This does not exclude that the effect we have uncovered is supported by variations within a continuous distribution, rather than by categorical distinct states. In any case, it still reveals a biologically relevant function in that changes within power and frequency of oscillatory patterns are robustly linked to nuanced changes in population activity in its relation to behaviourally meaningful variables.

- The authors frequently make statements about the population activity of neurons.

However, it is unclear from reading the Results section how population activity was defined and included in the model(s). Did the authors test population vectors based on firing rates? The Methods section provides more details, but it would help the reader to briefly summarize the way population activity was defined in the Results section.

By population activity we mean the spike density computed from the firing times of each recorded neuron (spike density computed by means of a Gaussian kernel from the binary spike data, as explained in Methods). Population activity in all our analyses is hence a three-dimensional matrix of shape [number of trials by number of neurons by number of time points], where for each trial we have a vector of firing density for each neuron throughout the trial duration. We have now added this clarification at the beginning of the Results section, where this is mentioned for the first time:

*We refer to neural population activity as the firing density of all the neurons recorded for each animal, across the duration of a trial, for all the trials performed. Firing density was computed from the binary spike data by means of a Gaussian kernel (see **Methods** for more details).*

- Figure 2B: It's difficult to see which data are from the same or different rat. While the authors write that power and frequency were consistently higher in one state, statistical testing of this hypothesis would be appropriate.

Thank you for this comment. We have now changed the style of the plot to visually distinguish the subjects. We also point at **Supplementary Figure 5**, where we have reported the states characterization for each animal separately.

We have now statistically quantified the difference between the two states in their power and frequency distributions, by means of the Mann-Whitney U test. We show that for three out of four subjects, the two states differ significantly both in power and in frequency (see **Supplementary Figure 2**). We added a description of these results in the Results section as:

*We then compared statistically power and frequency across states, finding significant differences in both power and frequency distributions in three out of four subjects (**Supplementary Figure 2**).*

We describe the analysis and statistical testing in the Methods section as:

We then quantified statistically the difference between power and frequency of the two states. To do so, we computed the power spectrum of each state occurrence that was longer than 50 timepoints (minimum number of data points to compute a power spectrum, arbitrarily chosen). We extracted the maximum power and its corresponding frequency for each valid (i.e., longer than 50 timepoints) state visit, obtaining a distribution of power and frequency for each state. We finally computed whether the difference in the power and frequency distribution between states were significant by means of the Mann-Whitney U test ⁸⁴, and corrected for multiple comparisons using the Benjamini/Hochberg method ⁸⁸.

Discussion:

- First paragraph of Discussion: The authors write “These states did not appear to recruit different neurons but rather changed how the same population processed information”. Please see my comments above on the discrepancy between this statement and other statements claiming that theta state modulates neural activity. In my opinion, a more accurate description of the main finding would be to say that the correlation of neural activity and trial outcome statistically depends on theta state.

We have tried to clarify above how the states can be not linked to different neural subpopulations and still have different population activity related to the different theta states. Please see our response above for more details.

- Last sentence, second paragraph of Discussion: “Our findings show that amplitude and frequency of theta oscillations can signal different network states, which affect population-level activity”. Do the authors mean activity or changes in activity/response dynamics? The authors write earlier that “there is no clear correspondence between an active state and a specific subpopulation of neurons”. While there appears to be a significant neural activity*theta state interaction effect in predicting trial outcome, this is not the same as saying that theta state affects neural activity. Also, reporting the model results would be helpful.

We understand our statement needs clarification. As mentioned above in this letter, our results show that the two different network states are linked to different representation of success vs failure by the whole population of neurons recorded. This difference goes beyond the average firing rate of neurons, involving different population firing patterns. The sentence mentioned here has been changed to avoid implying causal relationships too:

~~Our findings show that amplitude and frequency of theta oscillations can signal different network states, which affect population-level activity.~~ *Our work uncovers a different type of link between neural population activity and oscillations, where differences in neural firing patterns representing trial outcome are reflected in changes in amplitude and frequency of oscillations within the theta band.*

- Bottom of page 8, Discussion: statistical dependence instead of dependence?

We have changed *dependence* for *co-occurrence*.

- Second last paragraph of Discussion: “(...) we do not claim an underlying unidirectional causal flow from one to the other: oscillations and neural firing are two properties of the same system that emerge together”. I agree with this statement. However, the wording throughout the manuscript often suggests otherwise. There may also be a disconnect between the jargon used by computational scientists and experimental scientists. E.g., “A modulates B” is usually interpreted as “a change in A causes a change in B” by experimentalists. I suggest to change the wording throughout the manuscript to make it clear that the authors refer to correlations between theta states and behavioral outcome-dependent changes in neural activity. Perhaps most importantly, this statement needs to be placed much earlier in the text (see my comment above).

Thank you for this comment. We have now changed the wording throughout the text, trying to avoid any ambiguity in the description of the methods and in the implications of the effect under study. This includes changing words such as *dependence* and *modulation* to *correlation* and *co-occurrence*.

- Last paragraph of Discussion: “In conclusion, this study suggests that neural ensemble activity changes in context-specific ways depending on the state of the network as expressed by the theta oscillation, challenging static models of neural coding.” In my opinion, this conclusion – in the way it is written – is not supported by the data. First, neural ensemble activity is part of the network state, and network state cannot be defined by theta state alone; second, “context-specific” is very vague and it is not clear what exactly is meant by context; third, it should be noted that it is a statistical dependence; fourth, neural ensemble activity does not change dependent on theta state. Instead, the correlation of neural activity and trial outcome changes dependent on theta state.

We have rephrased the paragraph as follows:

In conclusion, this study suggests that ~~neural ensemble activity changes in context-specific ways depending on the state of the network as expressed by the theta oscillation~~ *changes in the network state, identified here as subtle changes in oscillatory power and frequency, are correlated with how the neural population activity represents behavioural information*, thus challenging static models of neural coding²⁰. That is, the hippocampal circuit might process complex variables such as success and failure in multiple ways depending on the state of the system. This underscores the role of theta oscillations *in the hippocampus*, ~~in creating~~ *reflecting* a flexible neural *information processing* landscape, and *possibly enabling contributing to the brain's the hippocampus to* *adaptation* to changing cognitive demands^{6, 78}.

Reviewer #4 (Remarks to the Author):

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Communications Biology initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Response letter to reviewers

The authors would like to express their gratitude to the reviewers who took the time to carefully review this manuscript. In light of the renewed reviewer's concerns, we have performed a thorough revision of our manuscript. The newer version of this manuscript presents considerable changes and reframing of results and conclusions. Major changes with respect to the previous version are: a full re-analysis of our data, (controlling for motion); a change in the state-conditioned decoding algorithm, to ensure the same state balance at each time point; change of p-value correction method to the standard cluster-based method used for time series data; new figures visualizing the effects reported.

Our responses describing the new analyses follow each comment in blue. The changes made in the manuscript follow our responses in red.

We sincerely hope the reviewers will appreciate our efforts and take the time to read through our paper.

Reviewers' comments:

Reviewer #1 (Remarks to the Author):

To the Authors:

The authors appeared to have evaded my primary concerns regarding movement and general behavioral outcomes, dismissing movement and locomotion as irrelevant. This claim was not met with data but with rhetoric. There was no analysis of head position, acceleration, licking, or the times of port entry or exit. Only rhetoric. I will focus my review on this issue rather than dedicating time to the entire manuscript, as I no longer believe the authors are engaging in good faith with my original comments.

We are left with a paper that uses a sophisticated HMM to tell the community: "Neurons do something (what exactly? "reconfigure" is a vague and terse description) different when animals are receiving a reward versus when they are not," a finding that is not novel and, in this context, is an expected consequence of behavior. Behavioral differences remain the most parsimonious explanation. While I can say the planets "reconfigure" around the sun, it does nothing to explain planetary orbits.

Despite the technical clarifications in the rebuttal, the authors have failed to address the

central, fatal confounds that undermine the paper's core claim. Our major outstanding issues remain and are, in fact, strengthened by this review. They fall into three main categories.

We respectfully note that the reviewer's report of our paper's core claim is not accurate. We do not simply address how neurons differently represent success and failure, as summarized above, but how success and failure's representation changes across brain states. The different brain states are linked to different theta patterns at the meso-scale (LFP) level. If success = low theta and failure = high theta, there would not be a need for a decoder, and we would not have used an HMM to detect the states.

We, however, acknowledge the importance of controlling for movement, and we have addressed this concern in the new review of this manuscript (see point below).

Major Issue 1: The Fatal, Uncontrolled Behavioral Confound

The authors' entire claim of a cognitive "reconfiguration" rests on their ability to disprove all other simpler, behavioral explanations. They have failed to do this.

Our original review stated that the significant effects observed "predominantly post-response may reflect motor, reward expectation, or arousal changes".

The authors' entire rebuttal to this foundational criticism rests on a single, unsupported, and demonstrably false assumption:

"we first should discard the motor cause because the animals are not moving"

This assumption is in direct contradiction to the original methodological papers (Allen et al. 2016, Shahbaba et al. 2022) that this dataset is built upon. As confirmed by the provided text:

A correct response (Time 0) is immediately followed by a "water reward... delivered below the odor port."

An incorrect response (Time 0) is immediately followed by a "buzzing sound" and sequence termination.

After every trial, rats are "required to run to the end of the track opposite the odor port."

The post-response window (1-2 seconds after Time 0), which is the only period where a significant effect is found in Figure 4 for two of the animals, is therefore not a period of "attentive immobility". It is a period of guaranteed, massive, and stereotyped behavioral differences:

Condition A (Success): The neural data reflects active reward consumption (licking), reward expectation, and preparation to move (e.g., the Crespi effect, where confidence leads to faster transitions).

Condition B (Failure): The neural data reflects the absence of reward, the presence of an aversive sound, and preparation to move.

Unsurprisingly, theta and neural activity would be different in these two conditions. The authors' standardization for "gross firing rate" does nothing to control for this. The decoder is not learning a "reconfiguration" of a cognitive code; it is almost certainly just learning to distinguish the massive, non-cognitive neural and motor-artifact-driven patterns of "licking" versus "hearing a buzz."

This is a fatal behavioral confound that has been dismissed, not addressed.

Thank you for clarifying the movement-related concern. We acknowledge that movement could in principle contribute to or confound the observed effects. We had originally claimed that the states could not be associated with movement because the task is performed in a constant spatial location. Being the animals not head-fixed, however, this relationship cannot be ruled out. We have hence addressed: 1. the relationship between movement and the HMM states; 2. how much movement alone is predictive of success and failure; and 3. the effect of movement on the neural population's representation of success vs failure.

1. We first quantify the relationship between movement and the states. The states are estimated over the full LFP data recordings, including periods during the experiment when the rats run to the end of the track. For each state visit, we compute the mean and standard deviation of the animals position, X_p and Y_p (coordinates of the animals along the long and short axis of the track, respectively). This gives us a distribution in mean and standard deviation of state visits per coordinate. We then compare the distributions

statistically across states using the Mann-Whitney U test. We find that the two states do not significantly differ in mean Xp or Yp, but differ in the standard deviation of Xp and Yp in three out of four subjects, indicating that HPHT state visits encompass a wider range of positions - without movement being the only driver of state assignment. We have included these analyses in the manuscript. In the **Methods** section, a new subsection (Motion analysis and deconfounding) describes the procedure, including details and preprocessing on the position data:

Movement data

The Xp and Yp, coordinates of the animal head's position along the track- recorded at a 30 Hz sampling rate – were upsampled to the 250 Hz of the neural data, and smoothed with a gaussian kernel of width of 500ms to attenuate noise. The Xp dimension corresponds to the long axis of the track, 150 cm long, while the Yp dimension corresponds to the short axis, 9 cm wide. These coordinates are expressed in the reference system of the tracking probe, where the coordinate [0,0] is outside the track, with the Xp track between coordinates 50 to 1000 (odour and reward port at approximately Xp = 60–120 and the reward port at opposite end of the track at approximately Xp = 900-1000), and with Yp coordinates spanning between 100 and 500. Note that the exact scale in physical units (cm) should be deduced according to the calibration used in the experiment.

Characterisation of state-related movement

For each state visit (a state visit being a collection of consecutive time points assigned to one state), we computed the mean and standard deviation of the Xp and Yp position of the animal. We tested for significant differences in these distributions across states using the Mann-Whitney U test ⁸².

The results are reported in a new subsection of **Results**, “Movement does not solely explain theta state differences or the neural representation of trial outcome”, summarised in **Table 1**.

As the experiment involved freely moving animals required to run to the opposite end of the track approximately every fourth trial ^{49, 50}, we explicitly tested whether the HMM states and neural representations of trial outcome could be explained by movement. We hence computed the mean and standard deviation of the subjects' position per state visit across states. The two states did not significantly differ in mean Xp or Yp (the coordinates of the animals along the long and short axis of the track, respectively), but did differ in the standard deviation of Xp and Yp, in three out of four subjects, suggesting that HPHT state visits encompassed a wider range of positions without movement being the primary driver of state assignment (**Table 1**).

We then asked how much movement predicts success and failure, and how much movement influences the neural population representations of success and failure. We identified two types of movement, that need to be handled separately due to their different

scales (and hence the different magnitude of variance introduced in the data): locomotion and head motion.

We had previously assumed that the experimental trials are performed in a constant spatial location. However, because the rats are supposed to run to the opposite end of the track on average every 4th trial (i.e. after every failed trial and after every successful sequence), we first checked how much locomotion is actually present in our trial data. To check for locomotion, we examine the Xp position of the rats along the track (Xp being the length dimension of the track, 150cm long) within a window of -1.0 to $+2.5$ seconds around response. If the animal started running within this window, we consider the current trial (in our study, a window between -1.0 to $+2.0$ seconds around response) affected by locomotion, and discard it. To the best of our knowledge, there is no universally accepted definition of locomotion for freely moving rodent experiments, with thresholds varying substantially across studies depending on task context. We then defined locomotion as displacement along the Xp axis of the track of more than a rat's body length (~ 20 cm), within a period of 500ms, meaning that head turns or body rotations without an actual body mass translation do not count as locomotion. These movements are instead included in the "head motion" confounding element that we address separately (see below). With this criterion, we excluded from our analyses 13 trials for subject 1, 18 trials for subject 3, and 2 trials for subjects 2 and 4. We describe this in the **Methods** section:

Locomotion exclusion

No universally accepted definition of locomotion exists for freely moving rodent recording experiments, with thresholds varying substantially depending on task structure^{65,83,84}. We therefore defined an operational criterion adapted to the spatial constraints of our task. We classified a trial as locomotion-affected if the animal displaced more than approximately one body length (~ 20 cm) along the long axis of the track (Xp dimension) within any 500 ms window falling within the extended trial epoch of -1.0 to $+2.5$ s around response. This criterion corresponds to a minimum translational speed of ~ 40 cm/s along the track axis, and it excludes only clear running bouts, not local head or postural movements. Applying this criterion, we excluded 13, 2, 18, and 2 trials for subjects 1, 2, 3, and 4 respectively.

2. After discarding locomotion trials, we asked how much head motion is predictive of success and failure. To do so, we used linear classification to predict success and failure from Xp and Yp head position across time (window: -1 to $+2$ s around response). As expected, this analysis shows that head position is predictive of trial outcome, particularly 1–2 s after response (consistent with reward-related behaviour in successful trials). We show this in **Supplementary Figure 12** and describe the results in the new **Results** subsection "Movement does not solely explain theta state differences or the neural representation of trial outcome":

After checking for locomotion potentially leaking into trial data (see **Methods**), we discarded 13, 2, 18, and 2 trials for subjects 1, 2, 3, and 4 respectively. To test how much movement would explain trial outcome, we used linear classification to predict success and failure from Xp and Yp head position. Head position predicted trial outcome above chance level particularly 1–2 seconds after response - consistent with reward-related behaviour (licking) in successful trials (**Supplementary Figure 12**).

3. We then addressed the questions: a. how much does head motion affect the neural representation of success vs failure? And b: is there any more state-related information in the neural activity regarding the representation of success and failure, that is not explained by movement?

To answer these questions, we treated head motion as a confound in our analyses involving neural activity, i.e. regressing out from the neural data Xp and Yp position, and using the residuals of this regression for the analyses involving the spike data.

3a. We compared prediction accuracy of success vs failure from the neural activity as in our previous version of this study, and from the neural activity deconfounded for motion. We notice that the overall prediction accuracy is decreased but remains higher than chance level after response — indicating that head motion is a contributing factor but not the sole driver of the neural signal. We report the procedure in the **Methods**:

Head motion deconfounding

After discarding locomotion-affected trials, we quantified how much residual head motion (Xp and Yp position) could predict trial outcome. To do so, we cut the position data into trials, so to have data **M** with dimensions number of trials by two (Xp and Yp dimension) by time points. We then standardized **M** across time per position dimension, and finally used class-balanced linear classification (ridge classifier, L2 penalty = 0.001) to predict trial outcome from **M**. We then treated head motion as a confound in all neural analyses by regressing out **M** from the spike density data **X** (per neuron and per time point) using linear regression. The residuals of this regression **X_d** were used in place of the original spike density for all subsequent motion-deconfounded analyses.

The findings, shown in **Supplementary Figure 13**, are reported in the **Results** section:

To assess how much of the neural signal could be attributed to head motion, we regressed out Xp and Yp position from the spike density data. Decoding accuracy from motion-deconfounded neural activity was slightly reduced relative to the original data (i.e., neural activity not deconfounded for motion), but remained above chance level after response, indicating that head motion is a contributing factor but not the sole driver of the neural signal (**Supplementary Figure 13**) in representing trial outcome. For comparison, decoding accuracy from shuffled neural activity was instead at chance level.

3b. We then ran our state-related analyses on the deconfounded neural activity: regression analysis performed per neuron (predicting success from neural activity, state activation

and their interaction), and state-conditioned decoding (predicting success per state, from the entire neural population).

Our regression analysis is now shown in the updated **Figure 3**, together with its p-values per neuron and per timepoint (uncorrected and cluster-based, spatial-temporal corrected). See below under Major concern 2 for details about the correction method.

While revising the state-conditioned decoding analysis, we implemented a small change to our algorithm to enforce that at each time point the two per-state decoders use the same number of trials per state. This implies that at each time point we take a subset of the state with most trials and randomly sample the trials used for decoding at each bootstrap. This is the same technique we had already implemented to obtain class balance per state. The reason for this change is to ensure comparability across state decoders' weights and performance. Our results show that within-state decoding is higher than cross-state decoding for some timepoints after response in two out of three subjects. This difference, however, reaches statistical significance under our permutation testing approach and cluster-based correction only in one out of three animals. We report the decoding accuracies within their 95% CI in **Figure 4A**. We also show the per-state decoding accuracies (now comparable), within their 95% CI in **Figure 4B**, and note that no one state constantly holds the highest decoding accuracy across time, suggesting that the states do not simply capture different signal to noise ratio distinctions in the data.

We report our findings in the **Results** section:

We noted that within-state decoding accuracy was higher than cross-state decoding for some time points after response in two out of three animals. This difference remained statistically significant after permutation testing (10,000 permutations, cluster-based corrected; see **Methods**) in only one animal. We therefore report decoding accuracies together with their 95% confidence intervals in **Figure 4A**. The trend in one out of three animals supports the interpretation that the two HMM states are associated with partially distinct patterns of neural population activity with respect to trial outcome. We visualized each individual state's accuracy, and standard accuracy (standard accuracy being the prediction accuracy when decoding from all the trials used in the states decoding), in **Figure 4B**. As observed, the accuracy changes between states and across time. Of note, no one state constantly holds the highest decoding accuracy across time: if one state was always the worst performing, we could reasonably consider whether the states reflected signal to noise ratio distinctions. The fact that the highest decoding accuracy alternates among states across time suggests that the states capture different patterns that become relevant at different time points for the processing of trial outcome.

We describe the change in the algorithm in the **Methods** section:

To ensure comparability of decoder performance and weights across the two states, we additionally enforced state balance across decoders: at each time point, the two per-state decoders were constrained to use the same number of trials. When one state had more trials available at a given time point, a random subset matching the size of the smaller state was drawn

at each bootstrap repetition. This ensures that any observed differences in decoding accuracy between states are not driven by differences in trial counts.

We updated **Supplementary Figure 15**, that involved state-conditioned decoding with states on a band-filtered LFP signal. Now the new figure shows state-conditioned decoding results on the deconfounded neural activity.

Major Issue 2: The Central Finding is Statistically Weak and Selectively Reported

The authors' entire claim for "ensemble pattern reconfiguration" rests on Figure 4. This finding is statistically weak and unconvincing on multiple levels.

Selection Bias (N=3/4): The central finding is based on only 3 of the 4 subjects. The exclusion of Subject 4, while justified by the authors on a "technical" basis (low failed trials), is a form of selection bias. It presents a "significant" finding from an incomplete dataset.

No Cross-Subject Correction (Inflated Error): This is the most critical statistical flaw. The authors analyze each subject independently and claim significance at $p < 0.05$. By presenting three separate $p < 0.05$ tests as evidence for one general finding, they are inflating the family-wise error rate, giving a $>14\%$ chance that at least one subject is a false positive. Subject 3, with its sparse and weak significance, looks precisely like this Type I error. A robust finding would require a proper cross-subject correction (e.g., Bonferroni, $p < 0.05/3 = 0.017$), which would almost certainly render this N=3 finding an N=2 finding.

Arbitrary "Researcher Degrees of Freedom": The choice of a "25-timepoint" (100ms) cluster window with a "12-timepoint" stride is arbitrary and unjustified. This is a classic "researcher degree of freedom" that, in combination with Benjamini-Hochberg, can be tuned to find significance.

With due respect, selection bias arises when subjects are excluded based on the dependent variable or its correlates. Our exclusion criterion — minimum trial count for decoder reliability — is orthogonal to the neural population dynamics being measured. This is analogous to excluding a subject in an fMRI study due to excessive motion artifacts: a quality-control decision, not a results-driven one. Moreover, Subject 4 is included in all other analyses (state characterisation, regression), with exclusion limited strictly to state-conditioned decoding where the minimum trial count cannot be met.

Bonferroni correction addresses inflated Type I error across multiple tests performed on the same dataset. Our per-subject analyses are conducted on separate animals, hence separate neural populations, and are therefore considered here independent replications rather than multiple comparisons.

The reviewer rightly pointed at a researcher's degree of freedom in our choice of a time window for the within-subject corrections. We agree that this was inappropriate. The fundamental problem with our approach was to average the data within the window, effectively reducing variance and inflating results. We have now used the standard cluster-based permutation correction throughout (as described in Maris and Oostenveld, 2007). This approach controls the family-wise error rate while respecting the temporal dependency of adjacent data points, making it appropriate for both the one-dimensional temporal correction (state-conditioned decoding) and the two-dimensional spatial-temporal correction (neurons $\times$ time, regression analysis). We have explained the updated correction procedures in the **Methods** section. Specifically, we describe the cluster-based correction method under the regression analysis subsection:

The correction proceeded as follows. For each neuron and time point, we defined an effect score as the difference between the mean null MAE (averaged across permutations) and the observed MAE: a positive score indicates that the observed prediction error is lower than expected under the null, i.e. that the model genuinely outperforms chance. We computed a one-tailed p-value at each neuron–time-point as the fraction of permutations for which the null MAE was less than or equal to the observed MAE; a small p-value therefore indicates that the observed prediction was better than nearly all permutation runs. Points with $p < 0.05$ were thresholded to form a binary map, and contiguous regions in the neuron–time plane (allowing adjacency both along the time axis and across neurons) were labelled as clusters. The mass of each cluster was defined as the sum of the effect scores across all neuron–timepoint pairs within the cluster; this weights the cluster statistic by effect strength rather than cluster size alone. To build the null distribution of maximum cluster masses, the same procedure was applied to each of the 10,000 permutation runs: the permuted p-values were thresholded at $p < 0.05$, clusters were identified in the permuted data, and the maximum cluster mass was recorded. A cluster in the observed data was declared significant if its mass exceeded the 95th percentile of this null distribution (cluster-level $\alpha = 0.05$).

and specified its time-only variant under the state-conditioned decoding subsection:

Specifically, at each time point we computed the observed difference between within-state and cross-state decoding accuracy under the original HMM state assignments. A time point was considered above threshold if the observed difference exceeded the 95th percentile of the pointwise permutation null distribution (one-tailed, $\alpha = 0.05$). Clusters were defined as contiguous runs of above-threshold time points. For each cluster, a cluster-level statistic was computed as the sum of the observed differences across all time points in the cluster. To assess the significance of each cluster, we built a null distribution of maximum cluster statistics: for each of the 10,000 permutations, we applied the same thresholding procedure to the permuted

differences and recorded the maximum cluster statistic. A cluster was declared significant if its statistic exceeded the 95th percentile of this null distribution (cluster-level $\alpha = 0.05$).

Major Issue 3: The "Reconfiguration" is Unsubstantiated Jargon

The paper's central claim of "ensemble pattern reconfiguration" is speculative jargon that is not demonstrated by the data.

Inference vs. Demonstration: The authors' method (Figure 4) provides only an inference that the patterns are different (i.e., the decoder fails when crossed). It provides zero demonstration of how these patterns are different.

No Visualization: At no point do the authors visualize the population code. They never describe or demonstrate what this "reconfiguration" is. Different cells? Different patterns of the same cells? We are never shown which neurons contribute differently, or what the "reconfigured" pattern even looks like.

The authors are using a vague term, "reconfiguration," to describe a difference, and that difference, as established in Issue 1, is almost certainly a behavioral artifact.

We have now added **Supplementary Figure 14**, showing the neural activity during successful trials minus the neural activity during failed trials (averaged across trials), per neuron and per time point, for each state, as well as how this changes across states. This new figure explicitly visualizes neural activity and how the population patterns change across states.

We report this in **Results**:

We visualized the population activity patterns for successful vs failed trials (average across trials) under each state, and showed the absolute difference of these patterns across states (**Supplementary Figure 14**). This plot is intended to complement the visualization of the described effect, that we termed *ensemble pattern reconfiguration*, suggesting that the representation of success vs failure by the entire neural population changes across states.

We have also added a panel in **Figure 4: Figure 4C** shows the absolute difference in decoder weights between the two states, per neuron and per time point, highlighting which neurons contribute differently to each state's representation of trial outcome:

To visually explore the states contributions, we plot the absolute difference in decoders weights in **Figure 4C**: the plot shows that the main differences in the decoders occur for some specific neurons across time and are more pronounced when within-state accuracy is higher than cross-

state accuracy (which is also when the two states accuracies differ the most). This plot is only intended as a visualization tool and should be interpreted carefully given that significance was reached in only one out of three subjects.

The two added plots should clarify how each neuron is affected and when under each state.

Conclusion

I did not care to continue reviewing this manuscript, as the authors did not engage in good faith with my original, primary comments and considered any further requests for them to defend their claims a waste of my time.

My concern stands. The authors have used a statistically weak ($N=3$), selectively reported, and arbitrarily windowed analysis to find a massive, uncontrolled behavioral confound (licking vs. no-reward). They have then obscured this finding with speculative jargon ("reconfiguration") that is not directly supported by their own evidence.

The central claim of the paper that they have found a "cognitive reconfiguration" is entirely unsubstantiated.

Reviewer #3 (Remarks to the Author):

The revised manuscript by Masaracchia et al. entitled "Frequency-specific theta states in the hippocampus reflect reconfiguration of population activity with respect to behavioural context" has substantially improved with respect to clarifying methodological aspects and underlying assumptions of the model and with respect to clarifying the conclusions of the study. The authors have addressed all my previous comments, and I have no further comments. As mentioned previously, the study is timely and of significance for the neuroscience community. The manuscript presents a very interesting finding, namely that power and frequency variations within theta oscillations are linked to the reconfiguration of neural firing at the network level so that brain state-dependent variations in population activity code for the same behavioral outcome. I support the publication of the manuscript in Communications Biology.

Reviewer #4 (Remarks to the Author):

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Communications Biology initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

We would like to thank all the reviewers for their comments and contributions in improving this study and manuscript.

Dear Reviewers and Editors of Communications Biology,

We would like to thank you for the positive evaluation of the latest revised version of this manuscript. To comply with the editorial requirements, we have made minor changes to the manuscript. These regard:

- Shortened title, without punctuation.
- Shortened abstract, all in present tense.
- Each main figure is now submitted as a separate PDF file, without captions.
- Captions of the main figures are now reported at the end of the manuscript, under the section “Figures Captions”.
- All Supplementary Figures are now submitted in a dedicated Supplementary Information PDF file.
- All source data for the main figures have been deposited to Figshare and the doi of the folder is reported in the manuscript under the section “Data Availability”
- Data availability and Code availability sections are now distinct.
- The code used to reproduce all analyses has now been deposited to Zenodo and the doi of the folder is reported in the manuscript under the section “Code Availability”.
- Sections “Author Contributions” and “Fundings” have been added to the manuscript as requested, while the “Acknowledgements” section is now only used to acknowledge other contributions.

We remain available for any other clarification or action needed.

Thank you for your attention.

On behalf of all authors,

Laura Masaracchia