

RESEARCH ARTICLE

Temporal progression along discrete coding states during decision-making in the mouse gustatory cortex

Liam Lang^{1,2,3} , Giancarlo La Camera^{1,2,3} , Alfredo Fontanini^{1,2,3} 

1 Graduate Program in Neuroscience, Stony Brook University, Stony Brook, New York, United States of America, **2** Center for Neural Circuit Dynamics, Stony Brook University, Stony Brook, New York, United States of America, **3** Department of Neurobiology and Behavior, Stony Brook University, Stony Brook, New York, United States of America

 These authors contributed equally to this work.

* giancarlo.lacamera@stonybrook.edu (GLC); alfredo.fontanini@stonybrook.edu (AF)



Abstract

The mouse gustatory cortex (GC) is involved in taste-guided decision-making in addition to sensory processing. Rodent GC exhibits metastable neural dynamics during ongoing and stimulus-evoked activity, but how these dynamics evolve in the context of a taste-based decision-making task remains unclear. Here we employ analytical and modeling approaches to i) extract metastable dynamics in ensemble spiking activity recorded from the GC of mice performing a perceptual decision-making task; ii) investigate the computational mechanisms underlying GC metastability in this task; and iii) establish a relationship between GC dynamics and behavioral performance. Our results show that activity in GC during perceptual decision-making is metastable and that this metastability may serve as a substrate for sequentially encoding sensory, abstract cue, and decision information over time. Perturbations of the model's metastable dynamics indicate that boosting inhibition in different coding epochs differentially impacts network performance, explaining a counterintuitive effect of GC optogenetic silencing on mouse behavior.

OPEN ACCESS

Citation: Lang L, La Camera G, Fontanini A (2023) Temporal progression along discrete coding states during decision-making in the mouse gustatory cortex. *PLoS Comput Biol* 19(2): e1010865. <https://doi.org/10.1371/journal.pcbi.1010865>

Editor: Alireza Soltani, Dartmouth College, UNITED STATES

Received: August 15, 2022

Accepted: January 10, 2023

Published: February 7, 2023

Copyright: © 2023 Lang et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data and code are available on GitHub: <https://github.com/llang6/temporal-progression>.

Funding: This work was partially supported by a U01 grant from the National Institutes of Health/ National Institute of Neurological Disorders and Stroke Brain Initiative (1UF1NS115779) to G.L.C. and A.F. and by grants R01DC015234 and R01DC018227 to A.F. The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author summary

Neural circuits can operate in a dynamic regime characterized by sequences of metastable states. Each state can last for hundreds of ms and is described by correlated spiking activity across ensembles of neurons. Metastable dynamics have been shown to play a role in representing sensory information and cognitive signals. In this article, we demonstrate that activity in the gustatory cortex of mice engaged in a decision-making task is metastable. We found that some metastable states are associated with the perception of taste, while others are associated with decision processes. This neural activity, and the behavioral performance sustained by it, can be explained by a metastable network model comprising connected ensembles of spiking neurons. When mimicking optogenetic silencing conducted in previous experiments, our model simulations reveal the link between the disruption of metastable state sequences and the behavioral impairment due to the

Competing interests: The authors have declared that no competing interests exist.

optogenetic manipulation. Altogether, our results reveal the importance of metastable states in mediating taste perception related to task performance.

Introduction

Over the past decade, the gustatory cortex (GC) has emerged as a model for studying cortical dynamics. Analysis of neural activity in alert rats has revealed that ensembles of GC neurons consistently hop between semi-stable patterns of activity [1–5] termed “metastable dynamics” (for reviews, see [6, 7]). While these metastable dynamics were originally found in the context of taste processing, recent work has reported that metastability can also be associated with cognitive processes such as taste expectation [8, 9] and sensorimotor transformations [5, 10]. The latter was shown in the case of rats producing gapes—orofacial reactions signaling aversion and rejection [11]—in response to bitter stimulation. During the course of such a consummatory experience, GC neural ensembles switch from taste-encoding to gape-anticipating metastable states [5, 10]. Optogenetic manipulations of activity preceding this switch lead to a delayed onset of gapes [10], demonstrating a causal role of GC’s metastable dynamics in naturalistic consummatory behavior.

Recent studies indicated that, in addition to representing taste and naturalistic sensorimotor transformations, GC can also contribute to taste-based decisions in mice engaged in a decision-making paradigm [12]. Mice were given one of four tastants—two sweets and two bitters—and, after a brief delay period, had to register a decision by licking one of two lateral spouts. Electrophysiological recordings in mice performing this task uncovered that GC neurons encode sensory and decision-making information during the sampling and delay periods. Interestingly, and counterintuitively for a sensory region, optogenetic perturbation of GC activity during stimulus sampling had no impact on behavioral performance, whereas silencing during the delay period significantly impaired task performance [12]. While important in highlighting the role of GC in taste-based decision-making, this research raised two fundamental questions: 1) Is GC’s activity metastable during decision-making and, if so, what role does GC metastability play in perceptual decisions? 2) Why is optogenetic perturbation of GC neural activity during stimulus sampling less disruptive than during the delay period?

Here, we investigate these questions using a Hidden Markov Model (HMM) analysis of the data [1, 4, 5, 13–15] and a biophysically grounded model of GC neural activity. Using HMM, we demonstrate the existence of a temporal progression of metastable states emerging, sequentially, for coding sensory information (taste quality), internal deliberations, and finally forthcoming action. To understand the computational mechanisms underlying the genesis of metastable activity and the emergence of coding states in the context of decision-making, we developed a metastable spiking network model of GC. The model provides a mechanistic understanding of the results of our HMM analysis—in particular, the temporal progression of hidden coding states for different sensory and cognitive variables. The model is then used to explain the experimental effects of the optogenetic perturbation—in particular, why performance is more sensitive to perturbations during the delay period than during the stimulus-evoked period [12]—and investigate the effects of perturbations on metastable dynamics.

Altogether, our results show that GC supports complex taste-related decisions via sequences of metastable states that follow a precise temporal progression. The mechanism for this progression is explained by a biologically plausible model of GC, which also clarifies the counterintuitive effects of optogenetic silencing on performance, and further provides testable

predictions on the effects of intensity and timing of optogenetic perturbations on metastability and behavioral performance.

Results

Decision-making neural activity in mouse GC is metastable

Mice implanted with drivable multi-tetrode arrays in the gustatory cortex were trained on a taste-based perceptual decision-making task [12]. In this task, head-fixed, water-restricted animals sampled a tastant—sucrose, quinine, maltose, or sucrose octaacetate—from a central spout and learned to lick one of two lateral spouts based on the following rule: sucrose or quinine means lick left, and maltose or octaacetate means lick right. A correct directional lick resulted in a water reward from the lateral spout. The intentional counterbalancing of taste quality between cued directions—sucrose and maltose are sweet, while quinine and octaacetate are bitter—ensured that the animal could not rely simply on the generalization of sweet or bitter to perform well on the task, but rather had to recategorize the four stimuli according to the rule. A variable delay between sampling and availability of the lateral spouts was implemented such that the average time between the sampling/taste event (first lick to the central spout) and the decision event (first lick to the lateral spout) was 2.53 s (range: 1.40 – 6.66 s) (Fig 1A). Once mice learned the task (i.e., they reached criterion performance of at least 75% for 3 consecutive days), recording sessions began. We recorded from a total of 16 sessions in four mice, with an average of 198 trials per session (range: 75 – 292) and 5 simultaneously recorded neurons per session (range: 3 – 9).

We hypothesized that metastable dynamics support the ability of GC to code for different task-related variables, informed by recent findings revealing a crucial role for GC in sensory and decision-making aspects of this task [12], as well as work demonstrating metastability in GC [1, 4]. To test this hypothesis, we fitted a Hidden Markov Model (HMM) to our ensembles of simultaneously recorded spike trains. This method allowed us to segment the neural activity into recurrent sequences of discrete states that are not directly observable—a signature of metastable dynamics previously described in this area [1, 6, 7]. The hidden states were defined as vectors of firing rates across the simultaneously recorded neurons [1, 4, 15–18].

HMM was fitted to the spike trains of each session separately, and then used to decode, trial-by-trial, the sequence of hidden states (see [Materials and methods](#) for details). The neural activity in all sessions and trial types could be characterized by sequences of metastable states, with a median of 4.5 hidden states per session (range: 3 – 8). Examples of state sequences are shown in [Fig 1B–1D](#), ([left panel](#)). Note how state transitions occur at random times in different trials. The distribution of state durations resembled an exponential distribution with mean 577 ms (standard deviation: 602 ms; median: 338 ms; range: 50 – 5,546 ms), in keeping with previous findings in the rat GC [1, 4]. These results demonstrate that ensemble activity of mouse GC during decision-making is metastable. Remarkably, some states occurred more often in specific trial types, with their onset times obeying distinct probability distributions, as we show next.

Distinct discrete states encode task-related variables

To identify whether hidden states encoded task-related variables, state occurrences were compared across different trial types (i.e., characterized by different stimuli or actions; see [Fig 1B–1D](#), [right](#) and [S1 Fig](#)). For a state to be considered a coding state, occurrence frequencies had to differ across correct trials of the four different stimuli (Chi-squared test, $p < 0.05$). An additional comparison was performed to determine whether the probability of observing the state was significantly different for one stimulus out of the four (assessed by running pairwise

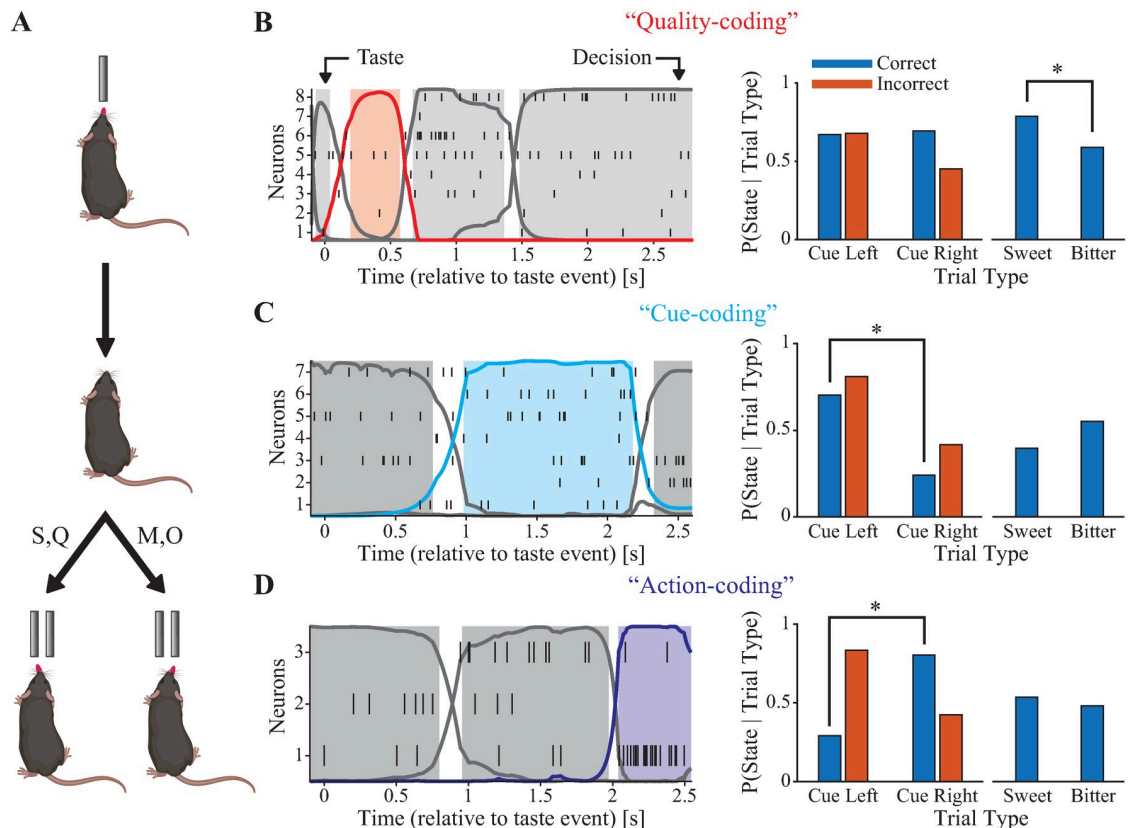


Fig 1. Decision-making activity in GC is metastable and distinct metastable states encode different task-relevant variables. **A:** Perceptual decision-making task schematic [12]. S: sucrose, Q: quinine, M: maltose, and O: sucrose octaacetate. **B-D (left):** Raster plots and HMM-decoded metastable states for in vivo spiking activity during decision-making behavior. “Taste” indicates the animal’s first lick to the central spout; “Decision” indicates the animal’s first lick to a lateral spout. **B-D (right):** Classification of hidden states. Quality-coding (B, red), Cue-coding (C, cyan), and Action-coding (D, blue) states. * indicates significant difference by Chi-squared test ($p < 0.05$, corrected for the number of decoded states in the session). Mouse graphic in **A** created with BioRender.com.

<https://doi.org/10.1371/journal.pcbi.1010865.g001>

Marascuilo post-hoc tests; see [Materials and methods](#)). If this was the case, the state was considered “Taste ID-coding.” If not, we grouped correct trials by cued direction (cue left: sucrose or quinine; cue right: maltose or octaacetate) and by taste quality (sweet: sucrose or maltose; bitter: quinine or octaacetate). If the state appeared with significantly different frequency between correct cue left and correct cue right trials, but not between correct sweet and correct bitter trials, it was considered a “Decision-coding” state. Similarly, if a state appeared with significantly different frequency between correct sweet and correct bitter trials, but not between correct cue left and correct cue right trials, it was considered a “Quality-coding” state. In the case that both of these comparisons were significantly different, the state was considered “Dual-coding” (see [Materials and methods](#) for details). Over all the 16 recording sessions we found 8 Taste ID-coding states, 3 Dual-coding states, 15 Decision-coding states, and 4 Quality-coding states. We chose to focus on Quality-coding and Decision-coding states because they have intuitive interpretations in the context of our behavioral task: Quality-coding states are associated with either sweet or bitter taste qualities, which reflect processing of sensory information. Decision-coding states are associated with either the left or right directions, which reflect decision-making activity. These states were distributed sparsely across sessions:

62.5% of sessions contained a Decision-coding state and 25% of sessions contained a Quality-coding state, with only 18.8% of sessions containing both (S1 Table, Column 1).

Decision-coding states were additionally divided into two categories: “Cue-coding” and “Action-coding” states, the former being associated with the perceptual categorization of the stimulus based on the direction it cues (i.e., sucrose or quinine → left; maltose or octaacetate → right), and the latter being associated with the execution of the action (i.e., go left; go right). This separation was based on the frequency of occurrence of states in correct and incorrect trials. Decision-coding states associated with the same cued direction in correct and incorrect trials were deemed Cue-coding since they were consistently associated with the cued direction, whereas Decision-coding states associated with opposite cued directions in correct and incorrect trials were deemed Action-coding since they were consistently associated with the chosen direction (see [Materials and methods](#) for details). We found 7 Action-coding states and 8 Cue-coding states (S1 Table, Column 1). Imposing stricter criteria at this step of the classification (i.e., requiring that sessions contain at least 10 incorrectly performed trials of each cued direction) did not qualitatively alter any of the results shown here, but decreased the number of Action- and Cue-coding states to 5 and 6, respectively.

Temporal sequences of coding states reflect task structure

Previous literature highlights the importance of temporal dynamics in GC [1, 5, 9, 19–21], with different variables being encoded at different times following stimulus delivery. To determine whether the variables associated with a taste-based, 2-alternative choice (2-AC) task are encoded sequentially, we analyzed the within-trial onset times of coding states. Consistent with the literature, we hypothesized that Quality-coding states should appear early in the trial, followed by Cue-coding and Action-coding states. To test this hypothesis, all within-trial onset times of the classified coding states were pooled across all sessions. Trial durations were “warped” so that taste delivery occurred at time 0 and decision occurred at time 1; this way timings could be compared across different trials in common time units. The distributions of coding states’ onset times in these common time units are plotted in Fig 2A. The distributions are clearly separated in time and arranged in a sequence where Quality-coding states appear before Cue-coding states, which, in turn, appear before Action-coding states. Mean onset times for Quality-coding, Cue-coding, and Action-coding states were 0.14, 0.41, and 0.67 (common time units), with inferred peaks occurring at 0.07, 0.38, and 0.87, respectively. This coding progression emerges after pooling all trials across all sessions; coding states (specifically, Quality-coding, Cue-coding, or Action-coding) were found in 1, 028 out of 3, 160 total trials, and occurrences of multiple coding states within the same trial were rare, likely due to small neural ensembles per session. However, the ordering of within-trial coding state sequences was highly reliable: out of 362 trials containing at least two different types of coding states, 313 (86%) had the proper ordering (i.e., Quality → Cue → Action). Thus, although the phenomenology is an inferred property of the aggregate pseudo-population of neurons, the results are consistent with sequential encoding in each trial unfolding from sensory information processing to perceptual categorization, and finally to decision-making and action execution.

To ensure that coding states were meaningful, and not a by-product of HMM analysis, a series of control analyses were performed. Randomly permuting trial type (i.e., stimulus) labels within correct and incorrect trials and repeating the state coding classification analysis severely reduced the probability of finding coding states in the data, as expected (4 Quality-coding states in unshuffled data vs. an average of 0.3 in shuffled datasets; 15 Decision-coding states in unshuffled data vs. none in shuffled datasets). In addition, two surrogate datasets were

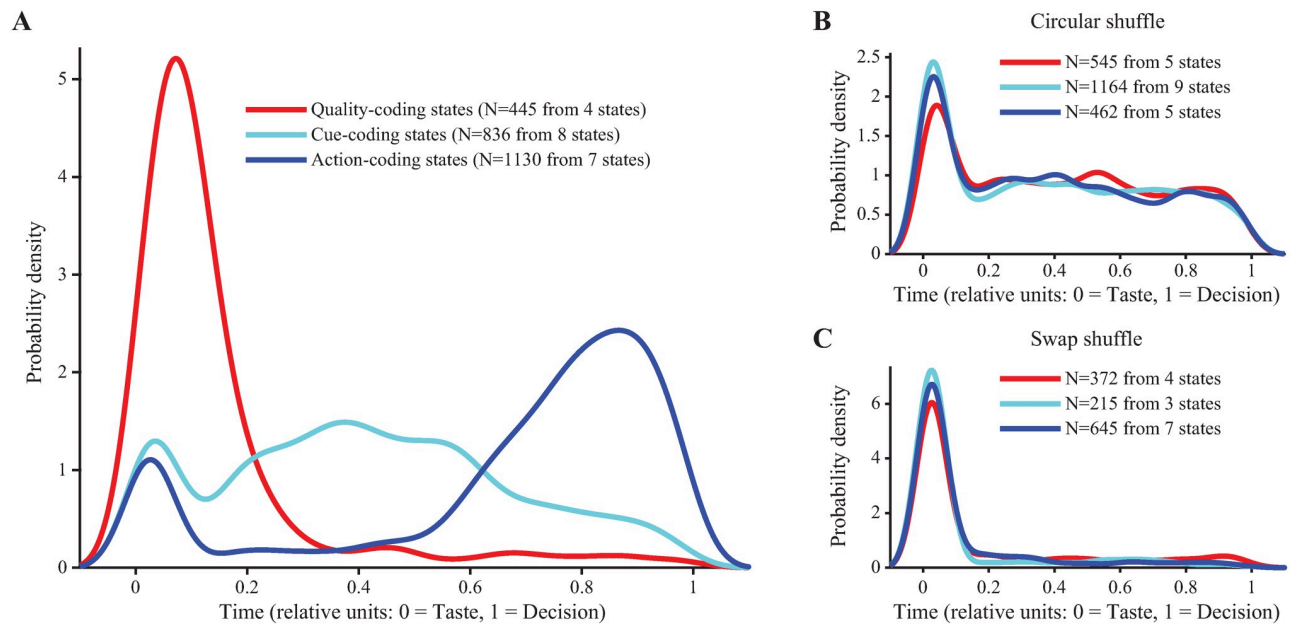


Fig 2. Coding state onset times exhibit an intuitive temporal sequence. A: Distribution of onset times of all states classified as either Quality- (red), Cue- (cyan), or Action-coding (blue) for experimental data showing a progression from coding for sensory information to abstract cue information to action information. B: Distribution of onset times for spiking data that are shuffled circularly. C: Distribution of onset times for spiking data that are swapped across time.

<https://doi.org/10.1371/journal.pcbi.1010865.g002>

generated from raw data by shuffling the spike trains circularly or by swapping segments of simultaneous activity across time (see [Materials and methods](#)). The circular shuffle preserves individual spike train autocorrelations but disrupts co-activation (hence, pairwise correlations) between neurons; the swap shuffle does the opposite, i.e., it preserves neuron co-activations but disrupts autocorrelations [22]. Both types of shuffling preserve the individual neurons' firing rates. These shuffled datasets were fitted with new HMMs and the resulting states were reclassified for their coding properties. Overall, we found only a modest reduction in the numbers of different coding states ([S1 Table](#)), which may indicate that the neural autocorrelations alone, or the pairwise correlations alone, may contain enough information to identify activity associated with task variables (e.g., if enough information is contained in the firing rates alone). This is possible since (i) the circular shuffling procedure will destroy the HMM states found in the original dataset, but can still produce new states, and (ii) randomly swapping time bins will not disrupt coordinated spiking across neurons. The states produced by shuffling, however, no longer truly reflect task variables, as confirmed by the lack of temporal information in the onset times of the coding states. In particular, the onset time distribution plots constructed for the shuffled datasets showed no clear separation among coding states of different types ([Fig 2B and 2C](#)). Thus, even though there was enough information in the shuffled datasets for the model to find coding states, the temporal structure of these states was completely disrupted.

Modeling metastable decision-making activity in GC

Results from HMM analysis demonstrate that GC relies on metastability to represent variables associated with the performance of a taste-based 2-AC task. To investigate the potential

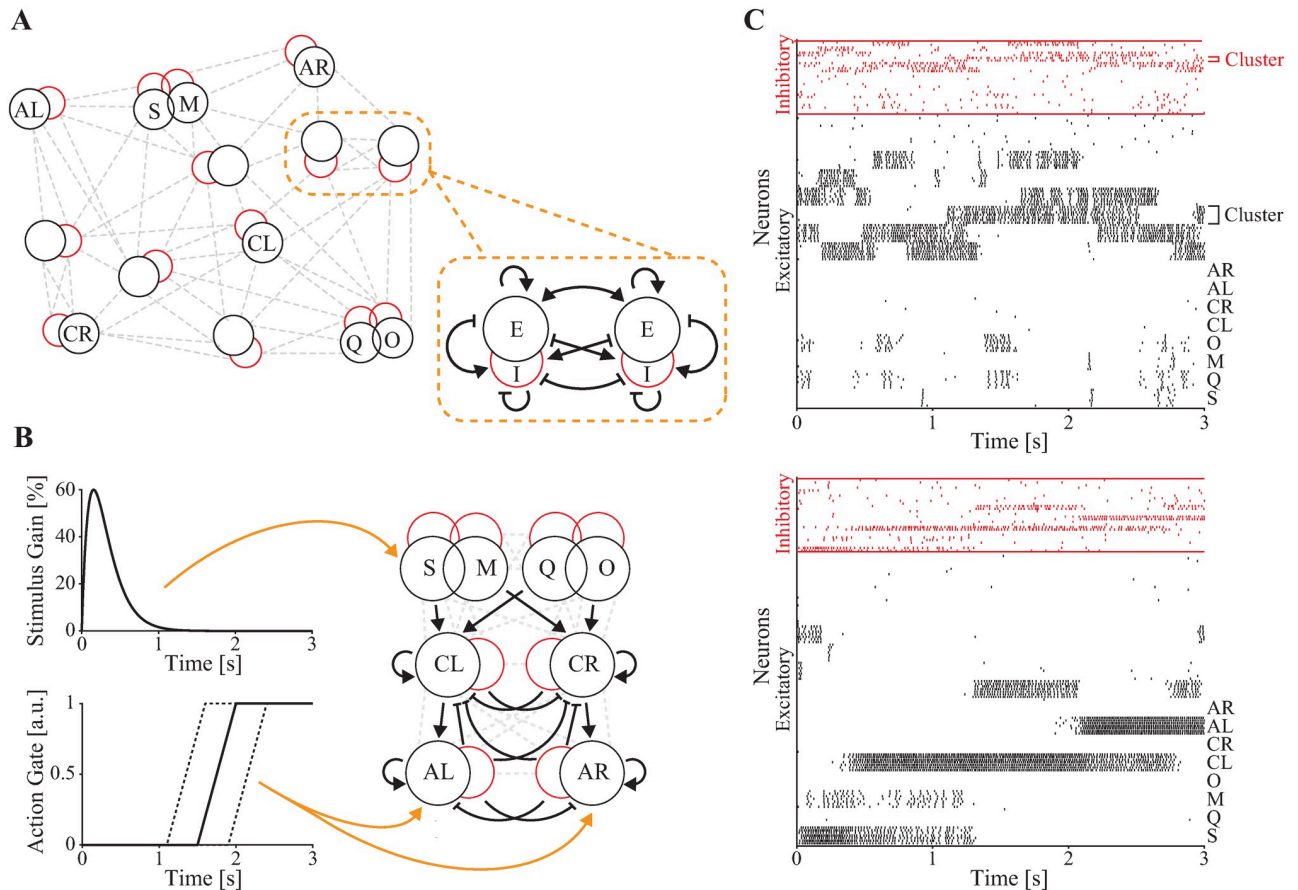


Fig 3. Clustered spiking network model for decision-making activity. **A:** General structure of the network. Neurons in the network are organized into fourteen excitatory (black) and partner inhibitory (red) clusters and connected in a highly recurrent manner. The zoomed inset indicates the generic structure of all of the connections that exist between any two excitatory clusters and their inhibitory cluster partners (pointed arrowheads are excitatory connections, flat arrowheads are inhibitory connections). **B:** Details of synaptic connections between taste, cue and action clusters, with inputs simulating the sensory stimulus (left, top) and the gating GO signal (left, bottom). **C:** Raster plot of metastable network activity in the absence of a stimulus (top) and after stimulation (bottom). Neurons are grouped by clusters (black: excitatory, red: inhibitory). S: sucrose, M: maltose, Q: quinine, O: sucrose octaacetate, CL: cue left, CR: cue right, AL: action left, and AR: action right.

<https://doi.org/10.1371/journal.pcbi.1010865.g003>

circuit-level mechanisms driving metastable activity in the context of the 2-AC task, we developed a recurrent spiking network model of GC that could replicate the key experimental findings.

Metastability in spiking network models can be achieved by a clustered architecture where synaptic connections among neurons within the same cluster are stronger than those among neurons of different clusters [4, 23–25]. Our network contained both excitatory (E) and inhibitory (I) neurons, each organized into 14 uniformly-sized clusters with stronger synaptic connections than among neurons of different clusters. Each excitatory cluster was paired with a partner inhibitory cluster via stronger synaptic connections than among neurons of non-partner E-I pairs [25] (see Fig 3A, S2A Fig, and Materials and methods for details).

The model was built to encode the crucial variables associated with the decision-making task, namely stimulus (taste), cued direction (cue), and selected direction (action). This was achieved by wiring the network such that specific properties could be assigned to 8 randomly selected excitatory clusters, coding respectively for: sucrose, quinine, maltose, sucrose

octaacetate, cue left, cue right, action left, and action right. Taste-selective clusters were characterized by having half of the neurons receive additional external input when either sucrose, quinine, maltose, or sucrose octaacetate were presented. The increased input varied in time as an alpha function (difference of two exponentials) that began at the start of the trial ([Fig 3B, top left](#)).

To replicate the experimental finding of coding states associated with a particular taste quality (i.e., sweet or bitter), clusters of neurons coding for sucrose and maltose were more strongly connected than clusters coding for, e.g., sucrose and quinine. Similarly, clusters coding for quinine and octaacetate were more strongly connected than clusters coding for, e.g., quinine and maltose. The same structure was imposed on the corresponding inhibitory partners. This structure was obtained by randomly selecting one quarter of the neurons in each partner E-I pair and connecting them with strong synaptic weights $J_{++}J_{\alpha\beta}$, with $J_{\alpha\beta}$ being a common reference synaptic weight from neurons in cluster β to neurons in cluster α ($\alpha, \beta \in \{E, I\}$). Neurons in the same taste cluster, but outside this more strongly connected subpopulation, had weaker synaptic connections (mean $J_{+}J_{\alpha\beta}$, with $J_{++} > J_{+} > 1$; see [Materials and methods](#), [S2B Fig](#), and [S3 Table](#) for model parameters). This structure mimics overlapping populations of neurons with similar response properties, and reflects the fact that many neurons in the GC [\[26\]](#) and elsewhere in the cortex [\[27\]](#) respond to multiple stimuli.

The 2 cue-selective clusters received stronger input from the appropriate pair of taste-selective clusters (i.e., connections from sucrose and quinine clusters to the “cue left” cluster and from maltose and sucrose octaacetate clusters to the “cue right” cluster were stronger compared to the cross-connections; see [Fig 3B](#) and [S3 Fig](#) for additional details); similarly the 2 action-selective clusters received stronger input from the appropriate cue-selective cluster (i.e., stronger connections from the “cue left” cluster to the “action left” cluster, and from the “cue right” cluster to the “action right” cluster). Cue- and action-selective clusters also had stronger intracluster connections than other clusters ([S3 Table](#)) to promote more “persistent” activity intended to reflect stable mental constructs.

To represent the approach of the lateral spouts in the behavior task, which prevents mice from acting on a decision until the end of the delay period, we incorporated a preparatory action gating signal in the model ([Fig 3B, bottom left](#)). The gate prevents the action-selective clusters from receiving input until a randomly chosen point in time in each trial, representing availability of the spouts. The gating mechanism is implemented as a multiplicative modifier of input that ramps linearly from 0 to 1 over 0.5 s, beginning at a random time (uniformly distributed over [1.1 s, 1.9 s]).

Altogether, the anatomy and topology of the model network gave rise to a random-looking, metastable dynamics in the absence of a stimulus ([Fig 3C, top](#)), and to a cascade of processes—from taste to cue to action—in the presence of a stimulus [\[28\]](#), due to preferential connections embedded within the highly recurrent network model ([Fig 3C, bottom](#)). The network typically responded to stimuli with high firing rates initially in the stimulus-responsive cluster and its overlapping partner, then in a cue cluster, and then in an action cluster ([Fig 3C, bottom](#)).

To infer the model’s choice on each trial, we computed firing rates for action clusters in 50 ms bins and considered a cluster to be “on” when it fired at $\geq 40\%$ of its maximum firing rate across the whole session. Behavioral scoring of trials was then evaluated following this simple rule: if the correct action cluster came on between 0.5 s and 3.0 s after stimulus onset and the incorrect action cluster did not, the trial was scored as correct; if the incorrect action cluster came on during the same window and the correct action cluster did not, the trial was scored as incorrect; if neither action cluster or both action clusters came on during this window, the decision was assigned randomly.

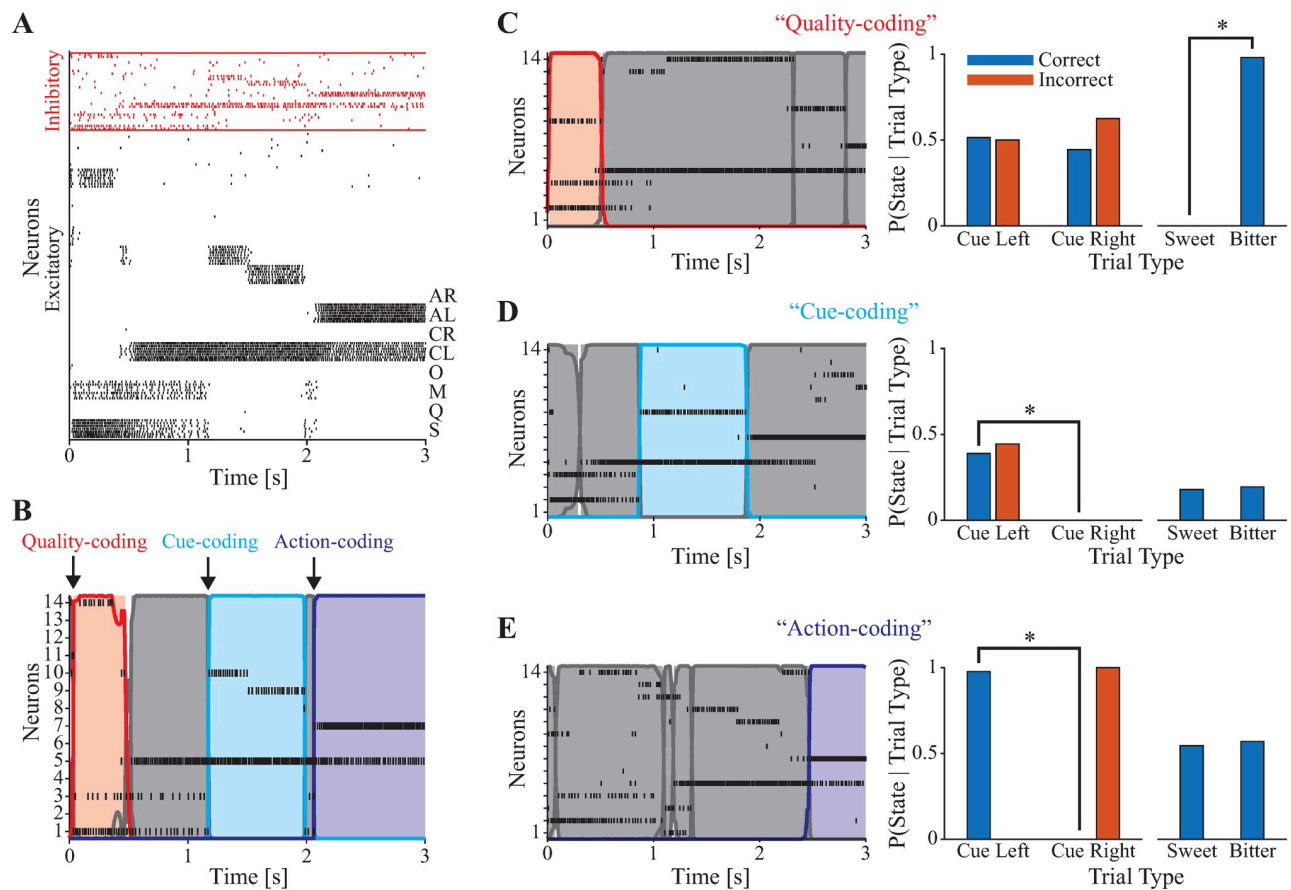


Fig 4. HMM analysis of network-simulated data. **A:** Example network activity showing sucrose-triggered response. **B:** Raster plot with overlaid HMM-decoded hidden states corresponding to the trial shown in **A**. Arrows indicate onsets of different coding states. **C-E (left):** Raster plots and HMM-decoded metastable states for simulated spiking. **C-E (right):** Classification of hidden states. Quality-coding (**B**, red), Cue-coding (**C**, cyan), and Action-coding (**D**, blue) states. * indicates significant difference by Chi-squared test ($p < 0.05$, corrected for the number of decoded states in the session).

<https://doi.org/10.1371/journal.pcbi.1010865.g004>

To test the model, we randomly generated 10 different synaptic weight matrices, which may be thought of as 10 independent network instances. For each instance of the network, we ran 100 trial simulations (25 simulations per tastant); in each simulation, a stimulus was applied and the network was allowed to evolve over 3 seconds. Over all 10 realizations, the network had an average accuracy rate of 80.1%, which is comparable to the experimental performance of mice performing the real task (82.8%; [12]). To assess whether the network model could reproduce the metastable dynamics observed experimentally, we randomly selected one neuron from each excitatory cluster to form a 14-neuron sample. The spike trains of these sampled 14-neuron ensembles were fitted with an HMM over the 100 trials (Fig 4A). Metastable states were decoded on a trial-by-trial basis as was done for the experimental data, and state coding properties were classified according to the same definitions used in Fig 1B–1D (see Fig 4C–4E).

Analysis of state onset times (Fig 5A) revealed that these distributions corresponded quite well with those observed in the experimental data. Most notably, the model captured the temporal sequence of distribution peaks observed in the data, showing a progression through time

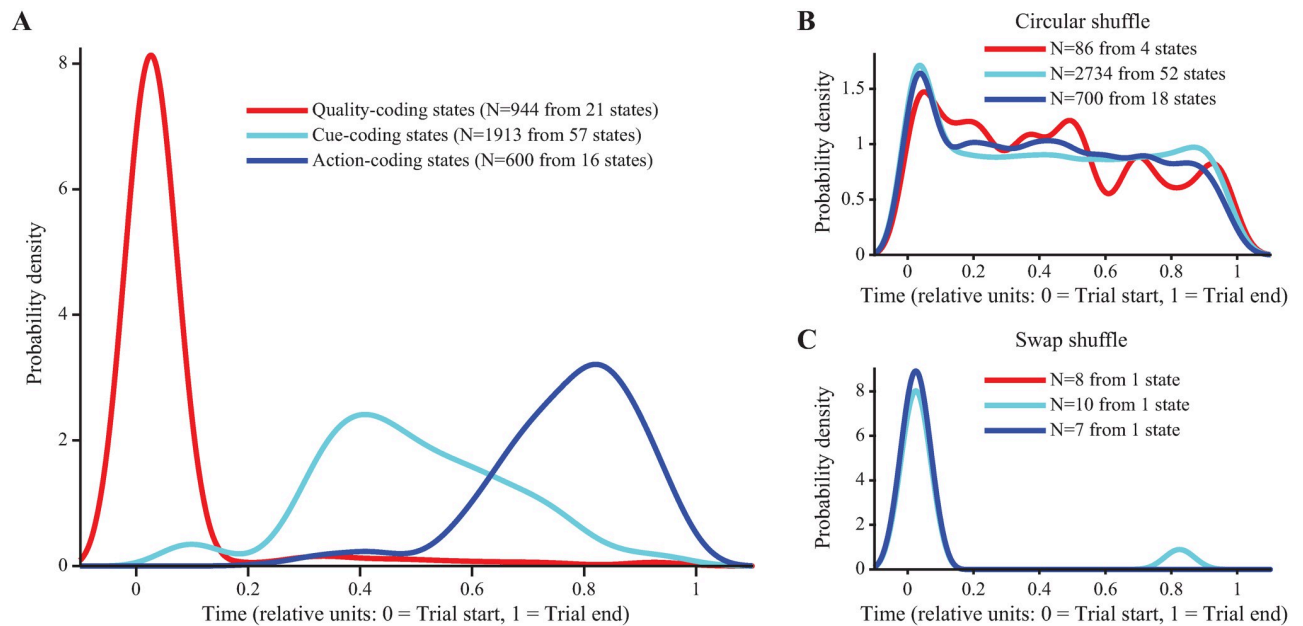


Fig 5. Coding state onset times in the simulated data qualitatively reflect those in experimental data. A: Distribution of onset times of all states classified as either Quality- (red), Cue- (cyan), or Action-coding (blue) for simulated data showing a progression from coding for sensory information to abstract cue information to action information. B: Distribution of onset times for spiking data that are shuffled circularly. C: Distribution of onset times for spiking data that are swapped across time.

<https://doi.org/10.1371/journal.pcbi.1010865.g005>

from taste-coding to abstract cue-coding and finally to action-coding. The mean onset times for Quality-coding, Cue-coding, and Action-coding states were 0.04, 0.49, and 0.74 common time units, with inferred peaks occurring at 0.03, 0.40, and 0.82, respectively. Similarly to the experimental data, no coding states with any meaningful temporal dynamics were found when shuffling the simulated data (circular or swap shuffle, see [Materials and methods](#) and [Fig 5B and 5C](#)). This finding confirms that coding the temporal structure of the task via a sequence of metastable states requires a particular network structure, with our model providing a specific link between network structure and function.

Relationship between GC dynamics and behavioral performance

To further probe the relationship between network structure and function, we subjected our model network to a series of external perturbations, inspired by the optogenetic silencing experiments carried out by [12]. Those experiments showed that task performance is unaffected by silencing GC during the “sampling period” (from 0 to 0.5 s post-taste event) but impaired by silencing GC during the “delay period” (from 0.5 to 3 s post-taste event). We modeled optogenetic silencing as a simple square pulse input on top of the baseline external current and with strength (height) measured as percent increase in the baseline. The silencing input targeted all inhibitory neurons in the network. Representative spike rasters for network activity with simulated silencing during the sampling and delay periods are shown in [Fig 6A](#). We ran 100 simulated trials for each of our 10 networks under both silencing conditions and measured task accuracy as we did before under control conditions. [S5A Fig](#) shows the results of silencing on network performance. Silencing during the delay period induced a more dramatic drop in performance than silencing during the sampling period ($p < 0.0001$;

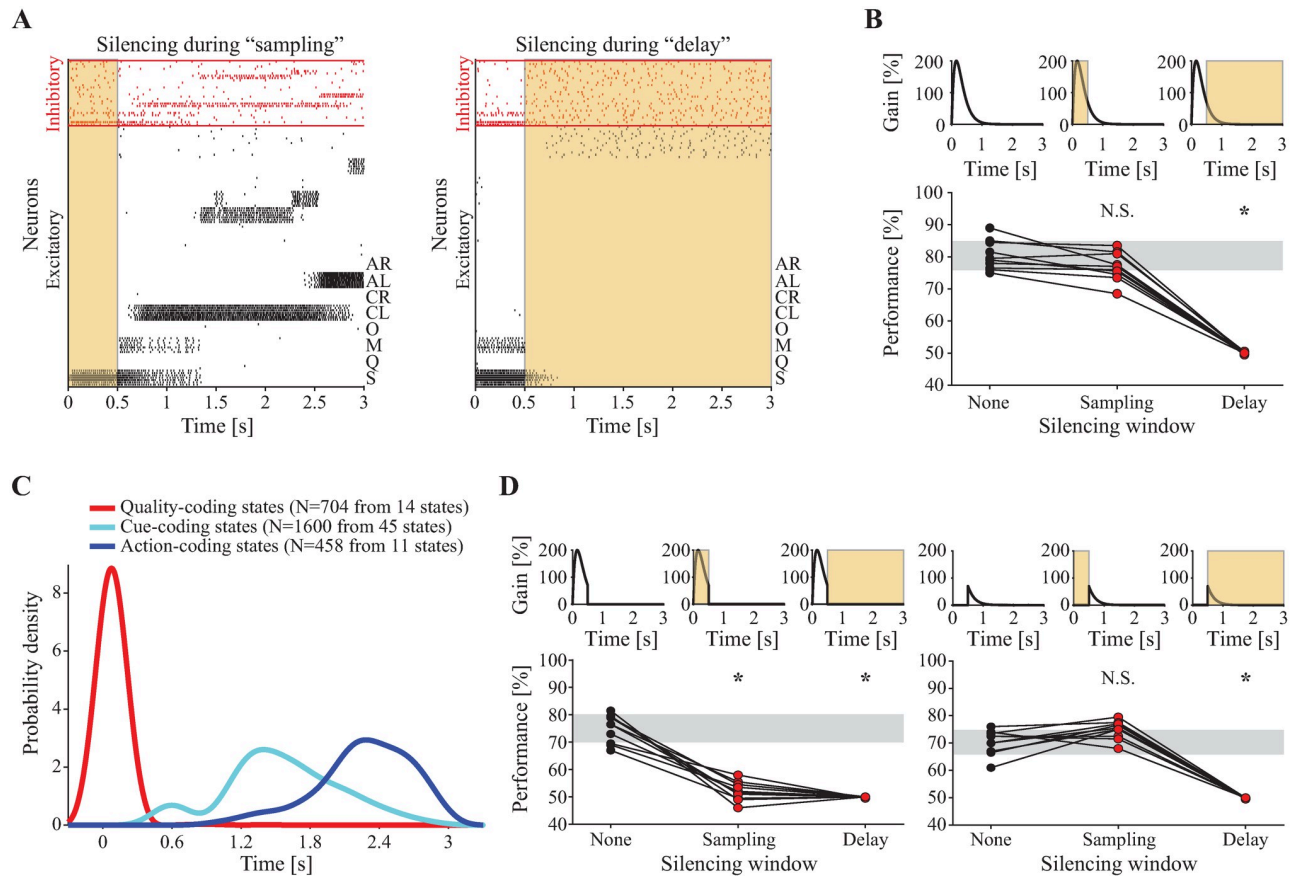


Fig 6. Effect of simulated optogenetic silencing on model performance. A: Rasters showing representative network activity when silencing is applied during "sampling" (0 to 0.5 s) and during "delay" (0.5 to 3 s). Shaded yellow regions indicate that silencing (100% increase in external baseline current for all inhibitory neurons) is on. B: Stimulus time courses and periods of silencing for 10 networks receiving full stimulus input (top) and their corresponding distributions of task accuracies (bottom). C: Distributions of onset times of coding states for 10 networks with stimulus input with gain of 200% (compare to Fig 5A, wherein the stimulus gain is 60%). D: Stimulus time courses and periods of silencing for 10 networks receiving partial stimulus input (top) and their corresponding distributions of task accuracies (bottom). Partial stimulus input was divided into Head only (left panel) and Tail only (right panel). Grey bars represent means ± 1 standard deviation for the corresponding None (baseline) condition; N.S. represents no significant difference with respect to the corresponding None; * represents significant difference with respect to the corresponding None ($p < 0.05$, Bonferroni-corrected post-hoc for 2-way within-subjects ANOVA with factors Stimulus and Silencing).

<https://doi.org/10.1371/journal.pcbi.1010865.g006>

Bonferroni-corrected post-hoc after significant 1-way within-subjects ANOVA); however, silencing during the sampling period also reduced the performance ($p = 0.0013$), contrary to our previous experimental findings [12].

These results depend on the specific parameters used to model the stimulus and the effect of silencing, parameters on which there is little or no experimental information. We therefore adjusted the taste stimulus gain and time constant to capture both the effects of silencing and our HMM results under control conditions. Increasing the stimulus decay time constant from 160 ms to 705 ms yielded performance results that matched the silencing experiments (S5C Fig), however, repeating the HMM analysis with this new stimulus revealed disrupted coding state dynamics (S5D Fig). Instead, leaving the decay time constant at 160 ms and increasing the stimulus gain from 60% to 200% captured both the silencing experiments (Fig 6B) and the temporal progression of the coding states (Fig 6C; compare with Fig 5C). Thus, with

appropriate stimulus parameters, the model can explain the silencing effects as well as the metastable dynamics of coding states found in the HMM analysis of the data.

With the effect of silencing on behavior now captured by the model, we investigated the mechanism by which silencing during the sampling period fails to impair task performance. There could be two (not mutually exclusive) mechanisms underlying the ineffectiveness of silencing during sampling: (1) the silencing input cannot overcome the simultaneous, strong stimulus input during the first half second; (2) the residual stimulus input occurring after the first half second provides sufficient information to drive network performance. We tested these hypotheses by decomposing the stimulus into two segments: the “head” that occurs during the first half second and the “tail” that occurs after the first half second. We then conducted additional simulations to measure network accuracy under control and under both silencing conditions for each component of the stimulus (Fig 6D). We compared performance values using a 2-factor within-subjects ANOVA with two 3-level categorical factors: Taste Stimulus (levels: Full, Head, and Tail) and Silencing (levels: None, Sampling, and Delay). We found a significant Stimulus-Silencing interaction and carried out post-hoc comparisons with a Bonferroni correction. Although each stimulus component was necessary for normal performance when no silencing was present (Full/None vs. Head/None, $p = 0.0026$; Full/None vs. Tail/None, $p < 0.0001$), the tail of the stimulus was both necessary and sufficient for normal performance in the presence of silencing during the sampling period (Full/Sampling vs. Head/Sampling, $p < 0.0001$; Full/Sampling vs. Tail/Sampling, $p = 0.6572$). This supports the latter of the two explanations described above—i.e., that residual stimulus input occurring after the first half second is sufficient for accurate performance. This also suggests that the dynamics of lingering stimulus information in GC can rescue the silencing-induced performance deficits after the sampling period.

To further explore this time-variable effect of silencing with a higher resolution, we took a 250 ms-wide silencing square pulse and slid its center along the 3 s trial window in 50 ms increments, computing the 10-network, 100-trial average performance at each point. Two silencing strengths were used: “Weak” (25%) and “Strong” (100%). This revealed 3 temporal windows of interest: i) sampling onset, when both weak and strong silencing have no effect on performance; ii) sampling offset and cue cluster onset, when both weak and strong silencing impair performance; and iii) mid-trial, when strong silencing impairs performance but weak silencing does not (S6 Fig). We investigated the impact of silencing in these three windows on network activity by taking a representative network and comparing its control trials (i.e., those without silencing) to silenced trials. We used the same initial conditions in all cases and 6 different combinations of silencing perturbation: Weak or Strong at the beginning of the sampling (“Beginning”: 0 to 250 ms), Weak or Strong at cue onset (“Cue onset”: 250 ms window centered on the onset time of the cue cluster in the corresponding control trial—these onset times have mean ~ 0.6 s), and Weak or Strong during the middle trial period (“Middle”: 1, 375 to 1, 625 ms). We fit an HMM to the 700 total trials of spiking activity generated by this network, and decoded its hidden states in each trial (Fig 7A). Effects on metastable states paralleled the effects on network task performance. Visually, the sequences of states under the control condition appear similar to those under silencing conditions that did not impact performance (Weak/Beginning, Strong/Beginning, and Weak/Middle) and different from those under silencing conditions that did impact performance (Weak/Cue onset, Strong/Cue onset, and Strong/Middle), wherein state transitions are much more frequent and state sequences appear muddled. This phenomenon is confirmed by computing an average pairwise state matching rate relative to control conditions (see Materials and methods) as a measure of sequence similarity to control sequences over time (Fig 7B). Classifying the hidden states further reveals a significant decrease in the fraction of trials containing a correct Action-coding

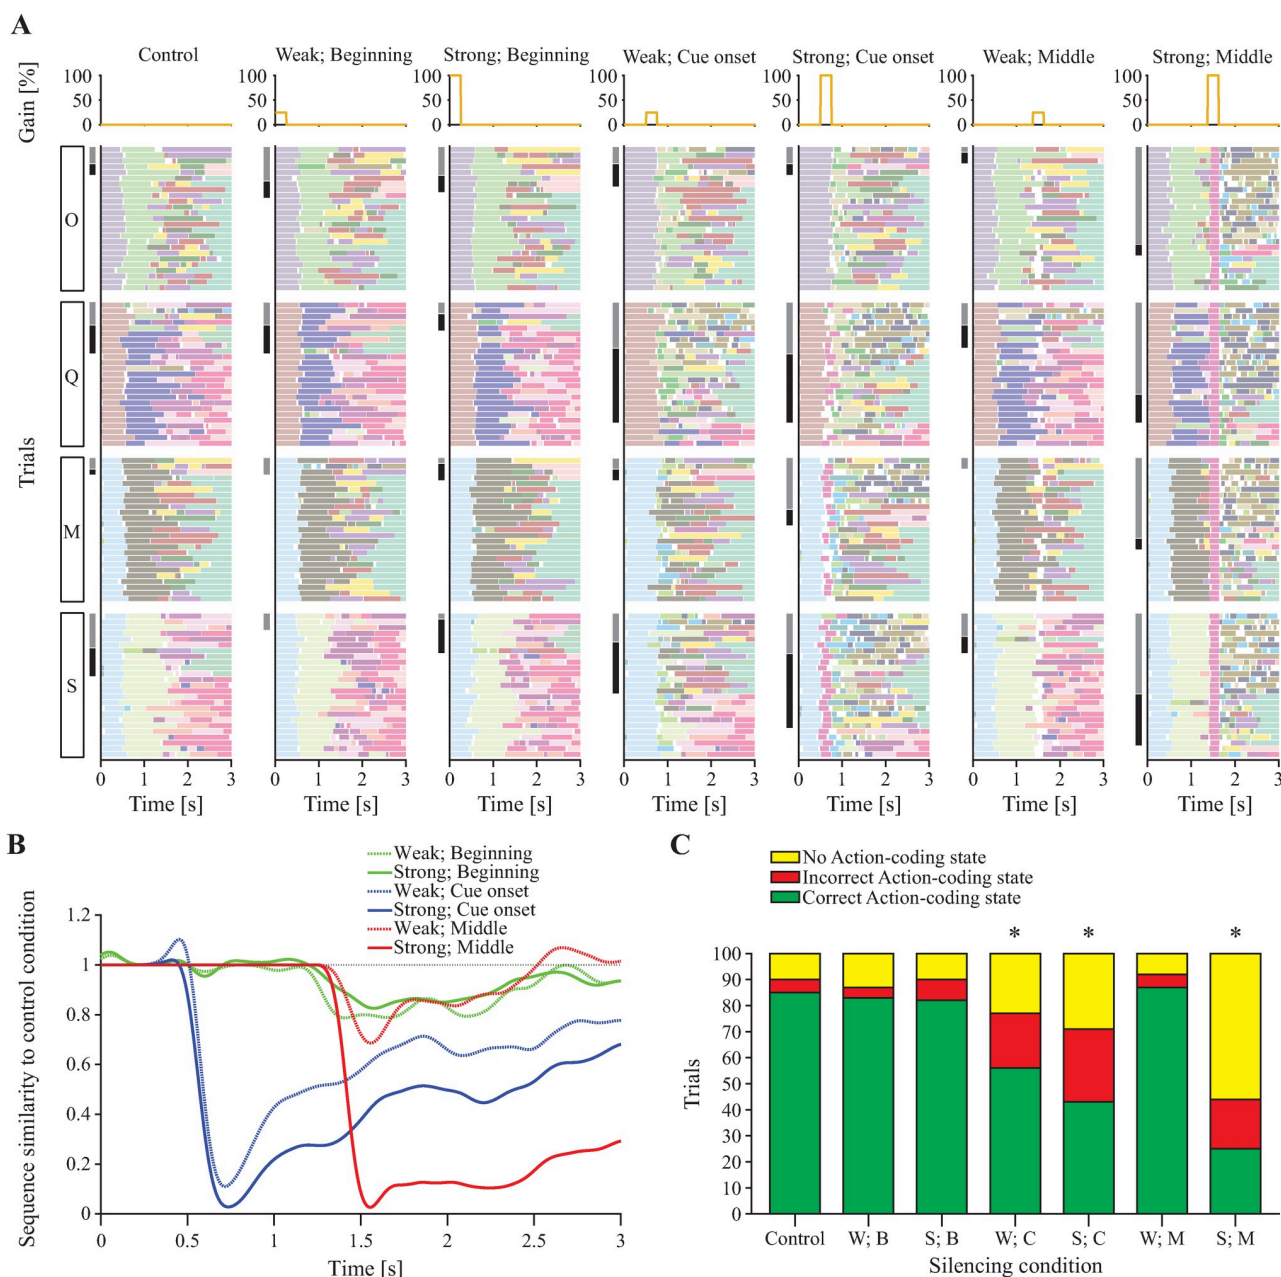


Fig 7. Effect of simulated optogenetic silencing on network activity. A: HMM-decoded states for a network subjected to various simulated optogenetic perturbation conditions. Silencing time courses for each condition are shown at the top of each column; decoded states over all 100 trials for each condition are shown below. Trials are ordered in stimulus blocks from bottom to top (S: sucrose, M: maltose, Q: quinine, O: octacetate) and are ordered in outcome blocks within each stimulus block from bottom to top (correct, incorrect, omitted; incorrect and omitted trials are indicated by black and grey shading, respectively, on the left-hand side). B: Average state sequence similarities over time. For each silencing condition in A, the corresponding curve represents how similar the state sequence is (on average) to the state sequence obtained under control (i.e., under no silencing). Each raw similarity score (ranging from 0 to 1) was normalized by the score obtained by comparing the control condition with itself (see [Materials and methods](#) for details, Eq 5). C: Numbers of trials out of 100 for each condition (W: Weak, S: Strong, B: Beginning, C: Cue onset, M: Middle) that contain correct, incorrect, or no Action-coding states, based on a coding state classification of the states from A. Action-coding states are considered correct or incorrect in each trial based on whether their direction preference (in correct trials) matches that trial's correct direction. * indicates significant difference vs. control in fraction of trials with correct Action-coding states (Chi-squared test across conditions, followed by pairwise Marascuilo post-hoc with $\alpha = 0.05$).

<https://doi.org/10.1371/journal.pcbi.1010865.g007>

state under the three silencing conditions that impair performance ($p < 0.05$ for Marascuilo post-hoc tests vs. control after significant Chi-square test across conditions; Fig 7C).

Discussion

In this study, we demonstrate that neural activity associated with taste-based decisions is metastable and temporally organized to sequentially encode information related to sensory stimuli, decision variables, and upcoming actions. We further provide a computational model that accounts for metastable dynamics in GC and its link with behavior. It further shows the relationship between silencing, metastable dynamics, and network performance. This model allows us to explain why silencing GC during sensory stimulation has a smaller impact on performance compared to perturbation during the deliberation period. These results significantly advance our understanding of GC and its role in decision-making, and confirm metastability as a dynamic regime relevant for GC's function.

Temporal dynamics of GC activity and functional roles of metastability in GC

GC neurons produce time-varying responses following gustatory stimulation [19, 29, 30]. Firing rates transition through three successive epochs, each having a different onset and encoding a distinct aspect of taste. The earliest epoch lasts for ~ 250 ms and encodes the contact of taste with the tongue. The second epoch (from ~ 250 ms to ~ 1 s) represents chemosensory information—that is, taste quality (i.e., bitter, sweet, salty, sour, and savory). Finally, the third epoch (after ~ 1 s) encodes the palatability of the stimulus and correlates with ingestive or aversive behaviors. This sequence is believed to result from recurrent interactions among GC, thalamic, and limbic areas such as the basolateral nucleus of the amygdala [8, 31–34]. These dynamics were initially uncovered through a careful analysis of single neuron firing rate modulations averaged across trials [19]. Additional studies of the ensemble activity from simultaneously recorded neurons further confirmed the existence of these dynamics at the single trial level, revealed trial-by-trial consistency of epochs, and identified epochs as metastable states with rapid onsets [1, 3, 4].

Recent experiments in rodents engaged in behavioral tasks suggest that additional epochs containing task-dependent variables could exist. For instance, recordings from GC of rodents intentionally licking a spout show preparatory activity that precedes the somatosensory epoch [35, 36]. Studies using taste-anticipating cues prove the existence in GC of anticipatory firing rate modulations, which encode general and specific expectations [8, 33, 37]. Similarly, experiments in rodents performing a decision-making task show neural activity consistent with sequences of epochs encoding decision-related variables [12, 38]. While this more recent work has been relevant in showing that GC does not just multiplex sensory variables, it has largely relied on single neuron analyses and across-trial averages.

The work presented here is the first that shows that variables associated with reward-driven decision-making are encoded by GC as metastable states. Our HMM analysis of GC ensemble activity during the four-taste, perceptual decision-making task has demonstrated that GC activity traverses sequences of metastable states encoding relevant task variables. These results are in line with prior evidence from higher order areas [15, 39–42] and are the first to show such a phenomenon in a sensory area. Our analysis further shows that, in the rodent GC, metastable states undergo a temporal progression with taste-coding states occurring first, followed by cue-coding states, and finally by action-coding states. This progression is consistent with a model in which GC first encodes sensory information (e.g., sweet vs. bitter), then decisional variables (e.g., predictive value of taste or preparatory activity associated with actions),

and then actions (e.g., lick left vs. lick right). Importantly, coding state transitions occur at variable times in different trials, allowing for pinpointing of the internal occurrence time of sensory perception or decision. This also allows for disentangling of the neural activity related to internal processes from the neural activity immediately following external triggers.

Perturbing GC dynamics

GC dynamics and their behavioral significance have been extensively tested via temporally precise optogenetic perturbations. Studies in rats receiving tastants through intraoral cannulae demonstrated that the effects of perturbations on gaping depend on the specific timing of silencing [10]. The differential impact of perturbations during specific temporal epochs was also present in mice performing a four-taste, perceptual decision-making task [12]. Optogenetic activation of parvalbumin-positive interneurons—an established method for silencing cortical activity [43]—was effective in reducing behavioral accuracy only if performed during the delay period, when cue-related, preparatory activity occurs, and not during the period of taste sampling. While these results were interpreted as evidence that GC's cognitive signals play a role in guiding taste-based decision-making, they also raised questions about the behavioral role of GC taste-coding states. In this work, we capture those results in our spiking network model and provide a possible explanation. Our results suggest that the ineffectiveness of silencing during the sampling period is best explained by residual stimulus input sufficing to drive the network after silencing, although competition between simultaneous stimulus and silencing inputs may also play a role.

Our additional silencing perturbation simulations, conducted at higher temporal resolution than explored experimentally, revealed precise points of susceptibility and resilience to silencing in the model, both in terms of task accuracy and metastable state sequences (S6 Fig). Model activity, metastable dynamics, and performance were robust against silencing when it occurred at the beginning of a trial, were disrupted by silencing when it was applied at the onset time of a cue cluster, and—depending on silencing strength—were conditionally impaired by silencing when it occurred in the middle of a trial (Fig 7). Furthermore, metastable dynamics and their ability to sequentially encode stimulus and task variables were tightly linked to network performance. The model therefore provides experimentally testable predictions for optogenetic silencing of neural decision-making activity with various intensities and windows of application.

Modeling GC: Metastability and decision-making

The observation of reliable and reproducible metastable dynamics in GC has inspired modeling efforts aimed at explaining the network mechanisms underlying this phenomenon and its functional significance. [28] showed how a spiking metastable network can account for state sequences observed during consummatory behavior (associated with ingestion or rejection of tastants); later, [4] demonstrated the existence of metastable dynamics in the absence of overt stimulation, and built a spiking network model of GC able to spontaneously produce metastable dynamics. This network has formed the basis of subsequent models that have captured a wealth of data pertaining to neural dynamics in GC [3, 9]. The model described here represents another significant advancement, in that the model can further process stimuli based on the abstract decision they instruct. The four tastants instructing decisions are categorized according to their taste quality first and their predictive value after, reproducing the temporal sequence of metastable states representing taste quality, predictive value of the taste, and action. The model also produces decisions with a similar performance to our experimental subjects.

In order to reproduce the experimental results observed in mice performing the four-taste, 2-AC task, the network had to include the following features: i) partial overlapping of clusters representing stimuli that share the same taste quality; ii) biased intercluster connections creating a flow of information from sensory clusters to cue clusters to action clusters; iii) an action gate signal priming action. These ingredients were absent in previous models and suggest the existence of possible biological counterparts in GC. Partially-overlapping clusters take into account that GC neurons may respond to multiple tastes [26] and may be the basis of stimulus generalization based on predominant taste quality. Biased intercluster connectivity may be the result of synaptic plasticity leading to the formation of neuron ensembles producing preparatory and action-related activity. The action gate was inspired by the presence of the lateral spouts moving toward the mouse, hence providing an implicit, ramping “GO” signal that may correspond either to an anticipatory input from areas like the amygdala, or to a disinhibitory signal triggered by the cue. This gate is also similar in spirit and implementation to the “urgency-gating model” described by [44]. Additional experiments are needed to test these predicted mechanisms.

The model also explains the effects of optogenetic perturbations and makes further predictions based on the strength of the perturbation and the period in which it is applied, as discussed in the previous section.

In summary, the present work shows that metastable dynamics occurs in GC during decision-making and provides a computational model explaining the genesis of metastability and its role in taste-based decisions. The ability of our model to reproduce the experimental results, together with our past body of work [3, 4, 9], suggest that the clustered inhibitory-excitatory network model may truly underlie the anatomical and functional organization of the rodent GC.

Materials and methods

Summary of behavioral procedure and electrophysiology

Methods for behavior and electrophysiology are available in [12], which presents the dataset analyzed in this study. Briefly, mice were unilaterally implanted with 8 tetrodes in the left gustatory cortex for chronic electrophysiological recording. After recovering from surgery, mice were water-restricted and trained to perform a four-taste, perceptual decision-making task to obtain water rewards. For the task, head-fixed mice sample a tastant—either sucrose, quinine, maltose, or sucrose octaacetate—from a central spout. The sampling period (time between first and last lick to the central spout) was ~ 0.5 s. After a delay period (~ 2 s) two lateral spouts advance in front of the animal’s mouth. To receive the water reward, the mouse must lick the left spout if the tastant was sucrose or quinine and must lick the right one if the tastant was maltose or sucrose octaacetate. Testing commenced once mice achieved a criterion performance of at least 75% for 3 consecutive days. Electrophysiological recordings were obtained while mice performed the task.

Raw data and preprocessing

Raw data consisted of spike trains from simultaneously recorded neurons over 21 recording sessions involving 5 mice. Within each recording session, trained mice performed an average of 190 trials (range: 75 – 292) of the behavioral task at an average accuracy of 82.8% (range: 71.0% – 93.3%). Only spikes within 100 ms before the taste event (first lick to central spout) to 100 ms after the decision event (first lick to lateral spout) were considered for each trial (we refer to these 100 ms segments before and after events as “padding”). These windows of analysis for each trial were concatenated and, prior to model fitting, neurons whose firing rates

averaged over this total window were < 2 Hz were excluded from analysis. Sessions with < 3 simultaneously recorded neurons were then excluded from analysis. This resulted in 16 recording sessions involving 4 mice, with an average of 5.4 neurons per session (median: 5.5; range: 3 – 9).

Model fitting and hidden state decoding

Hidden Markov Models (HMMs) were fit to spiking data from [12] as described in [9]. Briefly, spike sequences in a single session were binned into 2 ms bins, and each bin was assigned a symbol from 1 to $N + 1$, where N is the number of neurons and symbol s means that only the s th neuron fired in the current bin (multiple firings were rare and were replaced with a random selection of one of the firing neurons). The HMM is completely determined by the number of hidden states, M , the transition probability matrix, T , and the emission probability matrix, E . The entry T_{ij} of the $M \times M$ transition probability matrix is the probability of making a transition from state i to state j ; the entry E_{is} of the $M \times (M + 1)$ emission probability matrix is the probability of producing symbol s while in state i . This model can be considered an approximation to a Poisson-HMM in the case of very short time bins, with the exception that firing of multiple neurons in the same bin is not permitted (see [4] for details). The HMMs were fitted to the data by maximum likelihood, using the standard Baum-Welch expectation-maximization algorithm [45]. For each dataset, we fitted HMMs with values of M ranging from 2 to 50 and then chose the model with the lowest BIC score as the best model (see below). The maximal number of hidden states selected in the best model never exceeded 29. For each value of M , 10 different random initial conditions were used for T and E . For initial values of T , off-diagonal entries were close to zero and diagonal terms were close to 1. For initial values of E , entries were drawn independently from a uniform distribution between 0 and 1.

From each initial condition, the Baum-Welch algorithm was performed for 50 iterations (or until convergence at $< 10^{-10}$ tolerance on the log-likelihood, transition, and emission probabilities). Out of the 10 models, the model with the highest log-likelihood was selected as the best model for that given value of M . The best models for different M values were then compared using the Bayesian Information Criterion (BIC) to obtain the optimal model with the optimal number of states [9, 15, 42]. Specifically, the optimal model was the one that minimized the BIC score: $-2LL + [M(M - 1) + MN]\ln B$, where LL is the log-likelihood of the model given the data, M is the number of hidden states, N is the number of neurons, and B is the total number of observations in each session (i.e., the total number of time bins across all trials).

After fitting the models, we computed the posterior probabilities of each hidden state given the model and the data in each bin in single trials. A state admissibility criterion was then imposed: a state was considered detected with enough confidence only if its posterior probability was ≥ 0.8 for ≥ 50 consecutive ms ([9]; see also [1, 4, 15, 17, 18, 28, 42] for similar methods).

Classification of decoded states

For each session, all trials were associated with two labels: a stimulus label (sucrose, quinine, maltose, or sucrose octaacetate) and an outcome label (correct or incorrect). This allowed us to compute occurrence probabilities of states given the trial labels, $P(S | \text{Stimulus, Outcome})$, as the fraction of trials with the given labels in which state S was detected at least once. Each decoded state's coding properties were determined by first looking only at correct trials and comparing the state occurrence probability $P(S | \text{Stimulus, Correct})$ across all 4 stimulus labels through a Chi-squared test. If significant, all 6 pairwise post-hoc Marascuilo tests [46] were examined. If there was exactly one stimulus label for which all post-hoc tests involving it were

significant, S was classified as a Taste ID-Coding state. Otherwise, coding properties were investigated by grouping all trials into cue left vs. cue right trials (sucrose and quinine vs. maltose and sucrose octaacetate trials) and sweet vs. bitter trials (sucrose and maltose vs. quinine and sucrose octaacetate trials). The fraction of cue left, correct trials featuring S , $P(S | \text{Cue left, Correct})$, was compared to the fraction of cue right, correct trials featuring S , $P(S | \text{Cue right, Correct})$ through a second Chi-squared test, and the fraction of sweet, correct trials featuring S , $P(S | \text{Sweet, Correct})$, was compared to the fraction of bitter, correct trials featuring S , $P(S | \text{Bitter, Correct})$, through a third Chi-squared test (see S1 Fig). In each case, significance level was taken to be 0.05, corrected for the number of decoded states, N_d (i.e., $p < 1 - (1 - 0.05)^{1/N_d}$ is considered significant). Results from these tests were used to classify states as Decision-coding (significant result on second test but not on third test), Quality-coding (significant result on third test but not on second test), Dual-coding (significant results on both tests), and Non-coding (non-significant results on both tests). Decision-coding states were further divided into Cue-coding and Action-coding, as explained below. To control for chance occurrences of coding states, stimulus labels on correct trials were randomly permuted 10 times and, following each permutation, proportions of trials containing each decoded state were re-computed and the same Chi-squared tests were re-run to re-classify the states. We then compared the numbers of coding states found with this shuffling procedure to the number found from unshuffled data.

Decision-coding states were further subclassified by whether their “preferred direction” switched or remained the same between correct and incorrect trials. A coding state’s preferred direction was defined, separately for correct and incorrect trials, as the cued direction that maximized its occurrence frequency (assuming this frequency differed between trials of opposite cued directions): if $P(S | \text{Cue left}) > P(S | \text{Cue right})$ in correct and incorrect trials, Decision-coding state S codes for the left cued direction, and analogously for states coding for the right cued direction. Vice versa, if $P(S | \text{Cue left}) > P(S | \text{Cue right})$ in correct trials, but $P(S | \text{Cue left}) < P(S | \text{Cue right})$ in incorrect trials, Decision-coding state S codes for the action taken; the same conclusion is reached if $P(S | \text{Cue left}) < P(S | \text{Cue right})$ in correct trials, but $P(S | \text{Cue left}) > P(S | \text{Cue right})$ in incorrect trials. Finally, if the state S occurred equally often in cue left and cue right trials (either for correct or incorrect trials), we did not subclassify the Decision-coding state.

Randomization controls

Prior to model fitting, raw data were manipulated to form two additional surrogate data sets that served as additional controls [22, 42, 47]. For the first control, each spike train was shuffled circularly in each trial of a given session: a random number Δt from the uniform distribution on $[0, t_{\text{trial}}]$, where t_{trial} is the time length of the trial (including the padding), was added to each neuron’s spike train in each trial independently such that the portion of the spike train that now exceeded the end of the trial would “wrap back around” to the beginning (i.e., if a trial window starts at 0 and ends at t_{trial} , and the spike train $\{t_i\}$ is shifted by Δt , the new spike train is $\{t_j\} = \{t_i + \Delta t, \text{ if } t_i + \Delta t < t_{\text{trial}}; t_i + \Delta t - t_{\text{trial}}, \text{ otherwise}\}$). This procedure preserves the autocorrelations of individual neurons but alters the co-activation/pairwise correlations between neurons. For the second control, each session had its raw spiking data swap-shuffled: each trial’s neuron-by-time spike raster was segmented into 5 ms bins and then these bins/columns were randomly permuted for each trial independently. This randomization procedure preserves the co-activation/pairwise correlations between neurons but removes the autocorrelations of individual neurons. Both procedures preserve the trial-averaged firing rates of individual neurons.

Onset time analysis

For each session, lists of decoded states of each classified type (Quality-coding, Cue-coding, and Action-coding) were obtained. For each state in each list, all onset times of the state in each trial were obtained. Onset times were expressed relative to their trial's taste event ($t = 0$), with any negative times (i.e., states beginning during the left padding period) taken to be 0. To account for the fact that the inter-event interval (IEI) between the taste and decision event varied across trials (average: 2.52 s; range: 2.01 – 6.66 s), time was rescaled (or “warped”) in each trial by dividing onset times by the value of IEI in that trial. As a result, warped onset times were between 0 and 1 (where 0 is the time of the taste event and 1 is the time of the decision event) regardless of trial, neuron, session, and classification of the hidden state, and could be directly compared. From this, we obtained a list of (warped) onset times for coding states of each classified type. These lists were pooled across all sessions to obtain histograms of the onset times for coding states of each classified type (bin size: 0.05 warped time units). From the histograms, probability distributions of the onset times were estimated from linear interpolation on the peaks (100 points between any two adjacent bins). Distributions were then smoothed with a Gaussian filter (400 point width).

Spiking network model

The core of the model was a clustered excitatory/inhibitory spiking network with 4,000 excitatory (E) neurons and 994 inhibitory (I) leaky integrate-and-fire (LIF) neurons. Each neuron i had subthreshold dynamics that obeyed the differential equation in time t :

$$\frac{dV_i(t)}{dt} = -\frac{V_i(t) - V_L}{\tau_m} + \frac{I_{i,\text{ext}}(t) + I_{i,\text{syn}}(t)}{C}, \quad (1)$$

where $V_i(t)$ is neuron i 's membrane potential, V_L its equilibrium/leak potential, τ_m its membrane time constant, C its capacitance, and its total current $I_i(t)$ is the sum of its external input current $I_{i,\text{ext}}(t)$ and its synaptic current $I_{i,\text{syn}}(t)$. The synaptic current to neuron i in turn obeyed the differential equation in time t :

$$\frac{dI_{i,\text{syn}}(t)}{dt} = -\frac{I_{i,\text{syn}}(t)}{\tau_{\text{syn}}} + \frac{1}{\tau_{\text{syn}}} \sum_j J_{ij} \sum_n \delta(t - t_j^{(n)}), \quad (2)$$

where τ_{syn} is the synaptic time constant (different for excitatory and inhibitory presynaptic neurons), J_{ij} is the synaptic connection strength from presynaptic neuron j , δ is the Dirac delta function, and $t_j^{(n)}$ is a presynaptic spike time with n indexing over all spikes from a particular presynaptic neuron j . This system of differential equations was integrated numerically using the forward Euler algorithm with time step $dt = 0.05$ ms. Whenever $V_i(t)$ hit a threshold V_{th} , a spike was emitted and the membrane potential was immediately reset to a value V_r for a refractory period τ_r (see [S3 Table](#) for parameter values).

The E neurons were arranged into 14 clusters of 250 neurons each, plus 500 neurons not arranged in clusters (the “background population”). The I neurons were arranged into 14 clusters of 71 neurons each. Neurons were randomly connected with probability $P_{\alpha\beta}$, with $\alpha, \beta \in \{E, I\}$. In general, neurons in the same cluster had stronger synaptic connections than neurons belonging to separate clusters, given by $J_{E++}J_{EE}$ inside E clusters and $J_{I++}J_{II}$ inside I clusters. Each E cluster was paired with a partner I cluster. Synaptic weights between E and I neurons belonging to E-I partner clusters ($J_{I+}J_{\alpha\beta}$, with $\alpha, \beta \in \{E, I\}$) were stronger than between E and I neurons in non-partner E-I clusters ($J_{I-}J_{\alpha\beta}$, with $\alpha, \beta \in \{E, I\}$) (see [S2A Fig](#) and [S3 Table](#) for network parameters).

Eight out of the 14 excitatory clusters were randomly selected and assigned “roles” to represent crucial task information: taste clusters (sucrose, quinine, maltose, and sucrose octaacetate), cue clusters (cue left and cue right), and action clusters (action left and action right). For each taste cluster, half of the neurons inside the cluster responded to additional external input when that particular stimulus was presented to the network. All presented stimuli had the functional form $f(t) = (G/h)(e^{-t/\tau_D} - e^{-t/\tau_R})$, where G is the maximum gain, τ_D and τ_R are the decay and rise time constants, respectively, of the function $f(t)$, $h = (\tau_R/\tau_D)^{\tau_R/(\tau_D-\tau_R)} - (\tau_R/\tau_D)^{\tau_D/(\tau_D-\tau_R)}$, is a normalization constant, and the value of the function determined the percentage increase in external current at time t for selective neurons in the appropriate cluster (i.e., $I_{i,\text{ext}}(t) = I_{i,\text{ext}}^* \times (1 + f(t)/100)$ for the constant baseline external current $I_{i,\text{ext}}^*$).

Unlike generic clusters (i.e., clusters not ascribed a role), only 25% of neurons in a taste cluster were strongly connected (via synapses $J_{E++}J_{EE}$ or $J_{I++}J_{II}$, see above), while the remaining 75% were less strongly connected ($J_{E+}J_{EE}$ or $J_{I+}J_{II}$, respectively for E and I clusters, with $J_{++} > J_{+} > 1$; see [S2B Fig](#)). The subcluster of taste neurons with strong interconnections was also strongly connected with the subcluster of neurons representing the same taste quality (sweet or bitter), i.e., the sucrose and the maltose subclusters were connected with synapses equal to $J_{E++}J_{EE}$ or $J_{I++}J_{II}$ (for E or I neurons, respectively), and the same was true for the quinine and sucrose octaacetate subclusters. The rest of the connections were as for generic clusters shown in [S2A Fig](#). As shown in [S2B Fig](#), (top), it may help to think of this structure as introducing overlaps between taste clusters that share the same taste quality (i.e., sweet or bitter).

Excitatory cue clusters had stronger synaptic connections from the appropriate E taste clusters. For example, 60% of the weights from sucrose and quinine clusters to the cue left cluster and 60% of the weights from maltose and sucrose octaacetate clusters to the cue right cluster were increased by 85% ($J_{CT,E}J_{E-}J_{EE}$, see [S3 Fig](#) and [S3 Table](#) for parameter values). To facilitate longer lasting representations of cue information in the model, cue clusters had 60% of their excitatory intracluster weights scaled up by 5% ($J_{CC,E}J_{E++}J_{EE}$), and their synaptic time constants ($\tau_{\text{syn,CueE}}$, $\tau_{\text{syn,CueI}}$) increased (2.3 times greater for excitatory clusters, 1.3 times greater for their inhibitory partners). To discourage simultaneous activation of both cue clusters, 50% of the connections from inhibitory partner clusters of one cue direction to excitatory clusters of the opposite cue direction ($J_{CC,I}J_{I-}J_{EI}$) were increased by 40%.

Excitatory action clusters had potentiated connections from E neurons in cue clusters, with a preference for the appropriate direction: 50% of the weights from, e.g., cue right to action left were multiplied by 2.60 ($J_{AC,\text{In}}J_{E-}J_{EE}$), while 50% of the weights from, e.g., cue left to action left were multiplied by 2.75 ($J_{AC,\text{Cor}}J_{E-}J_{EE}$). We encouraged action clusters, similar to cue clusters, to represent information over longer time scales by scaling up 50% of their intracluster weights by 7% ($J_{AA,E}J_{E++}J_{EE}$) and increasing their synaptic time constants ($\tau_{\text{syn,ActE}}$ and $\tau_{\text{syn,ActI}}$) 3.30 times for excitatory clusters and 3.25 times for their inhibitory partners. As done with cue clusters, we discouraged co-activation of action clusters by increasing the strength of 50% of the connections from inhibitory partner clusters of one action direction to excitatory clusters of the opposite action direction three-fold ($J_{AA,I}J_{I-}J_{EI}$). Finally, to promote a transfer of information from cue to action via negative feedback, we increased the strength of 50% of the connections from inhibitory partners of action clusters to excitatory neurons in cue clusters ($J_{CA,I}J_{I-}J_{EI}$) by 40%. The full synaptic matrix for the whole model network is shown in [S4 Fig](#).

HMM analysis of the model simulations

For HMM analysis of model data, 10 different networks were used to mimic variability across sessions and/or subjects. The networks were obtained each time by sampling a new synaptic

matrix (all having the same basic structure described earlier). Although each network had random connectivity, the in-degree connectivity from each presynaptic population was fixed across the 10 networks. This enforced some degree of uniformity in firing rates across clusters and minimized the presence of bias in the performance of each network.

A simulation of each network was run for 100 trials (25 trials per taste stimulus). For each trial, each neuron started with a random initial membrane potential (normally distributed with mean 0 and standard deviation 4) and received a constant external input current. At the start of the trial, the network was presented with one of the stimuli (which modified the external input current for select neurons) and evolved over 3 seconds. We then sampled one neuron per excitatory cluster, excluding those that fired at < 2 Hz on average whenever possible. An HMM was fitted to the spike trains of the sampled neurons following the same procedure used for the HMM analysis of experimental data. Trials were scored by inferring the model's selected direction from the firing rates of its action clusters. Cluster firing rates were calculated by binning spikes in 50 ms windows, and a cluster was considered "on" when its firing rate was at least 40% of its maximum observed over the 100-trial simulation (setting a constant threshold at 34 Hz yielded exactly the same results). A trial was deemed correct if the appropriate action cluster came on after 0.5 s after stimulus onset while the opposite action cluster did not. A trial was deemed incorrect if the opposite action cluster came on during this window and the correct one did not. For an ambiguous trial, in which neither action cluster came on or both action clusters came on, we assigned a decision randomly such that the expected accuracy was 0.5, but also excluded the trial from further analysis. We then used the scored trials to classify the coding properties of the hidden states from each HMM, following the same procedure used for the experimental data. Onset times of all coding states were pooled over the 10 models, and distributions of these onset times were compared for Quality-coding states, Cue-coding states, and Action-coding states.

Simulated optogenetic silencing experiments

The same 10 networks (characterized by the same 10 synaptic weight matrices used above) were subjected to various perturbations in a series of simulations meant to mimic experimental optogenetic silencing of cortical activity [12]. Silencing was implemented as a simple, square stimulus pulse to all inhibitory neurons in the network. The height of the square pulse quantified the percent increase in baseline external current to inhibitory neurons. As a function of the square pulse's center c , its duration Δ_{sil} , and its strength κ_{sil} , the formula for the square pulse in time t read:

$$\varphi(t; c, \Delta_{\text{sil}}, \kappa_{\text{sil}}) = \kappa_{\text{sil}} \times [\Theta(t - (c - \Delta_{\text{sil}}/2)) - \Theta(t - (c + \Delta_{\text{sil}}/2))], \quad (3)$$

where Θ is the Heaviside step function. For silencing during the sampling period, $\kappa_{\text{sil}} = 100\%$, $\Delta_{\text{sil}} = 500$ ms, and $c = 250$ ms. For silencing during the delay period, $\kappa_{\text{sil}} = 100\%$, $\Delta_{\text{sil}} = 2, 500$ ms, and $c = 1, 750$ ms.

To test network performance across silencing conditions, we ran two independent ANOVAs, one for the stimulus with gain 60% and decay time constant 160 ms, and one for the stimulus with 60% and decay time constant 705 ms. These were 1-way within-subjects (repeated measures) ANOVAs with a single categorical factor, Silencing (levels: Full, Sampling, Delay). Both were significant and Bonferroni-corrected post-hoc tests were conducted (S5A and S5C Fig). Network performances across different silencing and stimulus conditions for our best stimulus (gain 200% and decay time constant 160 ms) were compared using a 2-way within-subjects ANOVA with 3-level categorical factors, Silencing (levels: None, Sampling, Delay)

and Stimulus (levels: Full, Head, Tail). The interaction was significant, and post-hoc analysis was conducted using a Bonferroni correction (Fig 6B and 6D).

We further tested perturbation strengths of $\kappa_{\text{sil}} = 25\%$ and $\kappa_{\text{sil}} = 100\%$ for a constant duration of $\Delta_{\text{sil}} = 250$ ms and variable c . Here we computed mean network accuracy as a function of the silencing center c by “sliding” the square pulse center along the trial window from 0 to 3 s in increments of $\Delta c = 50$ ms. For each value of c , we simulated 100 trials for each of the 10 networks under the given parameter set $\{\kappa_{\text{sil}}, \Delta_{\text{sil}}\}$, and then averaged the accuracy over the collected 1,000 trials.

For HMM analysis of a network with and without silencing perturbations, we chose the network whose average baseline accuracy was closest to the mean (samples shown in Fig 6B, **bottom left**). We reused the initial conditions from its 100 control trials 6 times, each time with a different silencing perturbation, either $\kappa_{\text{sil}} = 25\%$ (“Weak”) or $\kappa_{\text{sil}} = 100\%$ (“Strong”) and either $c = 125$ ms (“Beginning”), $c = t_i^{\text{Cue}}$ (“Cue onset”), or $c = 1,500$ ms (“Middle”), where t_i^{Cue} is the cue cluster onset time for the i th control trial, determined by thresholding the firing rates of the cue clusters as done for action clusters to infer network choice. In all cases $\Delta_{\text{sil}} = 250$ ms. A single HMM was then fitted to this collection of 700 trials and used to decode, trial-by-trial, the sequence of hidden states shown in Fig 7A. Classification of the coding properties of these states (for Fig 7C) was done as before by comparing state occurrence frequencies across different stimulus and outcome conditions. In each trial, an Action-coding state is called “Correct” if its direction preference in correct trials matches the current trial’s correct direction (e.g., an Action-coding state with left direction preference in correct trials is considered “Correct” when it appears in Sucrose and Quinine trials, and is considered “Incorrect” when it appears in Maltose and Octaacetate trials).

To quantify hidden state sequence similarity over time (Fig 7B), we compared the average state sequences under conditions of silencing and control (i.e., in the absence of silencing) in the following way. We discretized time into 2 ms bins (full resolution of the HMM) and, for each time bin, obtained a vector \mathbf{x} of 25 dimensions, one for each trial where the same taste stimulus was given (for example, sucrose). Each dimension contained either the current hidden HMM state in that trial, or a zero if no clear HMM state had been decoded in that trial. For example, if $\mathbf{x} = \{3, 2, 0, \dots, 1\}$, in the first trial the HMM state was 3, in the second trial it was 2, in the third trial no state had been decoded in that bin, and so on, until the last trial, where state 1 was the decoded HMM state. A vector \mathbf{x}^k was built for each silencing condition k shown at the top of Fig 7A, and an analogous vector \mathbf{x}_0 (in the same bin) was built for sequences under control (i.e., under no silencing). From each vector \mathbf{x}^k we built a measure of similarity with \mathbf{x}_0 by counting how many times two states were the same across all 25 trials. This required summing up 25×25 comparisons $\delta_{\mathbf{x}^k(t), \mathbf{x}_0(s)}$, where $t, s = 1, \dots, 25$, index the trial (i.e., the elements of vectors $\mathbf{x}^k, \mathbf{x}_0$), and $\delta_{a,b} = 1$ if $a = b$ and $\delta_{a,b} = 0$ otherwise. We performed this for all 4 taste stimuli and took the average. Since the total number of trials across all stimuli was $2,500 = 25 \times 25 \times 4$, the final measure of similarity in a single bin was

$$\langle \mathbf{x}^k, \mathbf{x}_0 \rangle = \frac{1}{2500} \sum_{\text{stim}=1}^4 \sum_{t,s=1}^{25} \delta_{\mathbf{x}^k(t), \mathbf{x}_0(s)}. \quad (4)$$

We finally normalized this measure by the similarity of \mathbf{x}_0 with itself:

$$\rho(\mathbf{x}^k, \mathbf{x}_0) = \frac{\langle \mathbf{x}^k, \mathbf{x}_0 \rangle}{\langle \mathbf{x}_0, \mathbf{x}_0 \rangle}. \quad (5)$$

We repeated this procedure for all silencing conditions in Fig 7A (top), i.e., for k standing for “Weak; Beginning”, “Strong; Beginning”, etc., and plotted the time course of each $\rho(\mathbf{x}^k, \mathbf{x}_0)$ in Fig 7B after convolving each curve with a Gaussian filter with binwidth of 150 time bins.

The hidden states shown in Fig 7B were classified for their coding properties by comparing their occurrence frequencies across trials based on stimulus and outcome (regardless of the silencing condition) using the same rules we followed previously. States classified as “Action-coding” were considered “Correct” or “Incorrect” based on their “preferred direction” in correct trials (see *Classification of decoded states*). If this direction matched the trial’s correct direction, the state was “Correct”; otherwise, it was “Incorrect.” Fractions of trials containing Correct Action-coding states were compared over all 7 conditions with a Chi-squared test. This was significant, so Marascuilo post-hoc tests were conducted to compare individual silencing conditions to the control condition (Fig 7C).

Supporting information

S1 Fig. Pipeline for classification of HMM states. Given a state S , the probability of observing S in a correctly performed trial was compared across the four trial types using a Chi-squared test. If the result was not significant (N), S was labeled “Non-coding.” Otherwise (Y), all six pairwise Marascuilo post-hoc tests were performed. If there was exactly one trial type for which all post-hoc tests involving it were significant (Y), S was labeled “Taste ID-coding.” Otherwise (N), the four trial types were grouped in two orthogonal ways: sweet (sucrose and maltose) and bitter (quinine and sucrose octaacetate), and cue left (sucrose and quinine) and cue right (maltose and sucrose octaacetate). Two additional Chi-squared tests were run: one to compare the probability of observing S in correctly performed trials between the sweet and bitter categories, and another to compare the probability of observing S in correctly performed trials between the cue left and cue right categories. These two tests are run independently, leading to four possible significance outcomes. If neither test was significant, S was labeled “Non-coding.” If both tests were significant, S was labeled “Dual-coding.” If the sweet vs. bitter test was significant and the cue left vs. cue right test was not, S was labeled “Quality-coding.” If the sweet vs. bitter test was not significant and the cue left vs. cue right test was, S was labeled “Decision-coding,” but could be further classified by examining its frequency of occurrence in correctly vs. incorrectly performed trials. If the cued direction for which S was more likely to be observed in correctly performed trials was the same for incorrectly performed trials (referred to as the “preferred direction” for S in correct and incorrect trials), S was labeled “Cue-coding.” If the preferred directions in correct and incorrect trials were opposite, S was labeled “Action-coding.”

(PDF)

S2 Fig. Synaptic structure for taste clusters. **A:** Generic synaptic structure between clusters without task roles. **Top:** Cartoon representing a pair of E-I partner clusters with their synaptic connection weights (red, flathead arrows: inhibitory connections; black, pointed arrows: excitatory connections). **Bottom:** synaptic matrix for the generic pair of E-I partner clusters (colormap corresponding to units of $\text{pA} \times \text{ms}$). Given any two excitatory clusters, $E1$ and $E2$, and their inhibitory partners, $I1$ and $I2$, synaptic connections were strong between neurons within the same cluster or belonging to partner clusters ($J_{E++}J_{EE}$ or $J_{I++}J_{\alpha\beta}$, with $\alpha, \beta \in \{E, I\}$), and weak between neurons in different clusters ($J_{E-}J_{EE}$ or $J_{I-}J_{\alpha\beta}$, with $\alpha, \beta \in \{E, I\}$). **B: Top:** Venn diagram depicting the “overlapping” structure of two E-I taste cluster pairs with the same quality (sweet or bitter). **Bottom:** synaptic matrix of “overlapping” taste clusters depicted at top (colormap corresponding to units of $\text{pA} \times \text{ms}$). Taste neurons in different clusters with the same taste quality had strongly connected overlapping subclusters ($E1 \cap E2$ and $I1 \cap I2$ in the

Venn diagram), with synaptic strength $J_{E++}J_{EE}$ or $J_{I++}J_{\alpha\beta}$, with $\alpha, \beta \in \{E, I\}$. These overlapping subclusters are shown in the bottom right corner of each taste cluster in the synaptic matrix. Taste neurons in the same cluster but outside the overlapping subclusters were connected with synaptic strength $J_{E+}J_{EE}$ or $J_{I+}J_{\alpha\beta}$, with $\alpha, \beta \in \{E, I\}$. All other neurons were connected as in the generic clusters of panel A, with $J_{E-} < 1 < J_{E+} < J_{E++}$ and $J_{I-} < 1 < J_{I+} < J_{I++}$. (PDF)

S3 Fig. Details of synaptic connections among taste, cue, and action E-I cluster pairs (compare with main Fig 3B). The main synaptic connections are depicted by curved lines (pointed arrowheads are excitatory connections, flat arrowheads are inhibitory connections). The strength of each connection is indicated next to the corresponding curved line (see S3 Table for parameter values). Synaptic weights from excitatory taste clusters to appropriate excitatory cue clusters (i.e., $S \rightarrow CL$, $M \rightarrow CR$, $Q \rightarrow CL$, and $O \rightarrow CR$) were larger by a factor $J_{CT,E}$ compared to generic weights between different clusters (e.g., $S \rightarrow CR$, $M \rightarrow CL$, $Q \rightarrow CR$, and $O \rightarrow CL$). Excitatory cue clusters had stronger intracluster connections (by a factor $J_{CC,E}$) compared to generic clusters; the synaptic weights from an inhibitory cue cluster to the opposite excitatory cue cluster were larger by a factor $J_{CC,I}$. Connections from excitatory cue clusters to correct excitatory action clusters (i.e., $CL \rightarrow AL$ and $CR \rightarrow AR$) were larger by a factor $J_{AC,Cor}$, while the connections to incorrect excitatory action clusters (i.e., $CL \rightarrow AR$ and $CR \rightarrow AL$) were larger by a factor $J_{AC,Inc}$, with $J_{AC,Cor} > J_{AC,Inc}$. Excitatory action clusters had stronger intracluster connections (scaled up by $J_{AA,E}$) and the weights from their inhibitory partners to the opposite excitatory action cluster were scaled up by $J_{AA,I}$. Connections from inhibitory action clusters to excitatory cue clusters were magnified by $J_{CA,I}$. All synaptic weight modifiers were applied with a given probability (reported in S3 Table). Connections among clusters not explicitly depicted here were as in S2 Fig (also see S4 Fig for the full synaptic weight matrix). Key: S: sucrose, M: maltose, Q: quinine, O: octaacetate, CL: cue left, CR: cue right, AL: action left, AR: action right. (PDF)

S4 Fig. Full synaptic weight matrix of the network model. The value in a block occupying row i and column j is the average synaptic weight (in units of pA×ms) from a presynaptic neuron belonging to cluster j to a postsynaptic neuron belonging to cluster i . The average synaptic values shown in the table were obtained (up to rounding) as: $P_{\alpha\beta}[f_{ov}^2 J_{\gamma++} + (1 - f_{ov}^2) J_{\gamma X}] J_{\alpha\beta}$ for connections involving taste clusters with the same taste quality (sweet or bitter), and $P_{\alpha\beta}(P_M J_M + 1 - P_M) J_{\gamma Y} J_{\alpha\beta}$ for all other connections. Here $\alpha, \beta, \gamma \in \{E, I\}$, $\gamma = E$ if and only if $\alpha = \beta = E$, $X \in \{-, +\}$, $Y \in \{-, ++\}$, $J_M \in \{J_{CT,E}, J_{CC,E}, J_{CC,I}, J_{AC,Cor}, J_{AC,Inc}, J_{AA,E}, J_{AA,I}, J_{CA,I}\}$ is a weight modifier, and $P_M \in \{0, P_{CT,E}, P_{CC,E}, P_{CC,I}, P_{AC,Cor}, P_{AC,Inc}, P_{AA,E}, P_{AA,I}, P_{CA,I}\}$ is the corresponding probability of applying the modifier ($P_M = 0$ when modifier was not applied). Key: S: sucrose cluster, M: maltose cluster, Q: quinine cluster, O: octaacetate cluster, CL: cue left cluster, CR: cue right cluster, AL: action left cluster, AR: action right cluster, 9–14: clusters without task roles, B: background (excitatory) population. (PDF)

S5 Fig. Network's performance and onset times of coding states for different parameters of the stimulus input (supplements main Fig 6). A, C: Effect of simulated silencing during sampling and delay periods on task performance for models with stimulus input with gain 60% and decay time constant 160 ms (A) and 705 ms (C), respectively. B, D: Distribution of onset times of coding states after fitting HMMs to models with stimulus input as in corresponding left panel. * indicates significant difference ($p < 0.05$) for Bonferroni-corrected post-hoc test

vs. None condition after significant within-subjects ANOVA across the 3 conditions. N.S. indicates no significant difference.

(PDF)

S6 Fig. Effect of simulated optogenetic silencing on network task performance as a function of time of silencing. Silencing was implemented as a square pulse stimulus with width 250 ms and height (strength, 100% or 25%) determining the increase in baseline external current for all inhibitory neurons in the network. Each point is an average accuracy over 1,000 trials (100 trials from each of 10 networks) for silencing centered at that point. For all points, the input stimulus had a gain of 200% and decay time constant of 160 ms. Grey region represents the mean ± 1 standard deviation of accuracy for networks with no silencing (same as in main Fig 6B). Regions-of-interest for silencing employed in main Fig 7A (Beginning, Cue onset, and Middle) are shaded in yellow.

(PDF)

S1 Table. Summary of the numbers of states found by HMM models fit to unshuffled, circularly shuffled, and swap-shuffled experimental data. Also indicated is the distribution of these states over sessions (16 sessions total). The same coding state classification procedure (see S1 Fig) was applied to the trial-by-trial decoding results from each model. “Hidden states” is the theoretical best number determined from model fitting; “Decoded states” is how many were actually found after decoding trial-by-trial.

(PDF)

S2 Table. Summary of numbers of states found by HMM models fit to unshuffled, circularly shuffled, and swap-shuffled simulation data. Also indicated is the distribution of these states over sessions (10 sessions total). The same coding state classification procedure (see S1 Fig) was applied to the trial-by-trial decoding results from each model. “Hidden states” is the theoretical best number determined from model fitting; “Decoded states” is how many were actually found after decoding trial-by-trial.

(PDF)

S3 Table. Model parameters. Symbols, values, and brief descriptions for all parameters used in the simulations of the spiking network model presented in the main text.

(PDF)

Acknowledgments

We thank Dr. Luca Mazzucato, Xiaoyu Yang, and Tianshu Li for sharing their computer code, and Dr. Arianna Maffei for helpful discussions and feedback on the manuscript.

Author Contributions

Conceptualization: Liam Lang, Giancarlo La Camera, Alfredo Fontanini.

Formal analysis: Liam Lang.

Supervision: Giancarlo La Camera, Alfredo Fontanini.

Writing – original draft: Liam Lang, Giancarlo La Camera, Alfredo Fontanini.

References

1. Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB. Natural Stimuli Evoke Dynamic Sequences of States in Sensory Cortical Ensembles. *Proc Natl Acad Sci U S A*. 2007; 104(47):18772–18777. <https://doi.org/10.1073/pnas.0705546104> PMID: 18000059

2. Moran A, Katz DB. Sensory Cortical Population Dynamics Uniquely Track Behavior across Learning and Extinction. *J Neurosci*. 2014; 34(4):1248–1257. <https://doi.org/10.1523/JNEUROSCI.3331-13.2014> PMID: 24453316
3. Mazzucato L, Fontanini A, La Camera G. Stimuli Reduce the Dimensionality of Cortical Activity. *Front Syst Neurosci*. 2016; 10:11. <https://doi.org/10.3389/fnsys.2016.00011> PMID: 26924968
4. Mazzucato L, Fontanini A, La Camera G. Dynamics of Multistable States during Ongoing and Evoked Cortical Activity. *Journal of Neuroscience*. 2015; 35(21):8214–8231. <https://doi.org/10.1523/JNEUROSCI.4819-14.2015> PMID: 26019337
5. Sadacca BF, Mukherjee N, Vladusich T, Li JX, Katz DB, Miller P. The Behavioral Relevance of Cortical Neural Ensemble Responses Emerges Suddenly. *J Neurosci*. 2016; 36(3):655–669. <https://doi.org/10.1523/JNEUROSCI.2265-15.2016> PMID: 26791199
6. La Camera G, Fontanini A, Mazzucato L. Cortical Computations via Metastable Activity. *Current Opinion in Neurobiology*. 2019; 58:37–45. <https://doi.org/10.1016/j.conb.2019.06.007> PMID: 31326722
7. Brinkman BAW, Yan H, Maffei A, Park IM, Fontanini A, Wang J, et al. Metastable Dynamics of Neural Circuits and Networks. *Appl Phys Rev*. 2022; 9(1):011313. <https://doi.org/10.1063/5.0062603> PMID: 35284030
8. Samuelsen CL, Gardner MPH, Fontanini A. Effects of Cue-Triggered Expectation on Cortical Processing of Taste. *Neuron*. 2012; 74(2):410–422. <https://doi.org/10.1016/j.neuron.2012.02.031> PMID: 22542192
9. Mazzucato L, La Camera G, Fontanini A. Expectation-Induced Modulation of Metastable Activity Underlies Faster Coding of Sensory Stimuli. *Nat Neurosci*. 2019; 22(5):787–796. <https://doi.org/10.1038/s41593-019-0364-9> PMID: 30936557
10. Mukherjee N, Wachutka J, Katz DB. Impact of Precisely-Timed Inhibition of Gustatory Cortex on Taste Behavior Depends on Single-Trial Ensemble Dynamics. *Elife*. 2019; 8. <https://doi.org/10.7554/eLife.45968> PMID: 31232693
11. Grill HJ, Norgren R. The Taste Reactivity Test. I. Mimetic Responses to Gustatory Stimuli in Neurologically Normal Rats. *Brain Res*. 1978; 143(2):263–279. [https://doi.org/10.1016/0006-8993\(78\)90569-3](https://doi.org/10.1016/0006-8993(78)90569-3) PMID: 630409
12. Vincis R, Chen K, Czarnecki L, Chen J, Fontanini A. Dynamic Representation of Taste-Related Decisions in the Gustatory Insular Cortex of Mice. *Curr Biol*. 2020; 30(10):1834–1844.e5. <https://doi.org/10.1016/j.cub.2020.03.012> PMID: 32243860
13. Radons G, Becker JD, Dülfer B, Krüger J. Analysis, Classification, and Coding of Multielectrode Spike Trains with Hidden Markov Models. *Biol Cybern*. 1994; 71(4):359–373. <https://doi.org/10.1007/BF00239623> PMID: 7948227
14. Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, Tishby N, et al. Cortical Activity Flips among Quasi-Stationary States. *Proc Natl Acad Sci U S A*. 1995; 92(19):8616–8620. <https://doi.org/10.1073/pnas.92.19.8616> PMID: 7567985
15. Ponce-Alvarez A, Nácher V, Luna R, Riehle A, Romo R. Dynamics of Cortical Neuronal Ensembles Transit from Decision Making to Storage for Later Report. *J Neurosci*. 2012; 32(35):11956–11969. <https://doi.org/10.1523/JNEUROSCI.6176-11.2012> PMID: 22933781
16. Gat I, Tishby N. Statistical Modeling of Cell-Assemblies Activities in Associative Cortex of Behaving Monkeys. In: *Neural Information Process. Syst.* vol. 5; 1993. p. 945–53.
17. Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E. Simultaneously Recorded Single Units in the Frontal Cortex Go through Sequences of Discrete and Stable States in Monkeys Performing a Delayed Localization Task. *J Neurosci*. 1996; 16(2):752–768. <https://doi.org/10.1523/JNEUROSCI.16-02-00752.1996> PMID: 8551358
18. Bollimunta A, Totten D, Ditterich J. Neural Dynamics of Choice: Single-Trial Analysis of Decision-Related Activity in Parietal Cortex. *Journal of Neuroscience*. 2012; 32(37):12684–12701. <https://doi.org/10.1523/JNEUROSCI.5752-11.2012> PMID: 22972993
19. Katz DB, Simon SA, Nicolelis MA. Dynamic and Multimodal Responses of Gustatory Cortical Neurons in Awake Rats. *J Neurosci*. 2001; 21(12):4478–4489. <https://doi.org/10.1523/JNEUROSCI.21-12-04478.2001> PMID: 11404435
20. Fontanini A, Katz DB. State-Dependent Modulation of Time-Varying Gustatory Responses. *J Neurophysiol*. 2006; 96(6):3183–3193. <https://doi.org/10.1152/jn.00804.2006> PMID: 16928791
21. Sadacca BF, Rothwax JT, Katz DB. Sodium Concentration Coding Gives Way to Evaluative Coding in Cortex and Amygdala. *J Neurosci*. 2012; 32(29):9999–10011. <https://doi.org/10.1523/JNEUROSCI.6059-11.2012> PMID: 22815514

22. Maboudi K, Ackermann E, de Jong LW, Pfeiffer BE, Foster D, Diba K, et al. Uncovering Temporal Structure in Hippocampal Output Patterns. *Elife*. 2018; 7. <https://doi.org/10.7554/eLife.34467> PMID: 29869611
23. Deco G, Hugues E. Neural Network Mechanisms Underlying Stimulus Driven Variability Reduction. *PLoS Comput Biol*. 2012; 8(3):e1002395. <https://doi.org/10.1371/journal.pcbi.1002395> PMID: 22479168
24. Litwin-Kumar A, Doiron B. Slow Dynamics and High Variability in Balanced Cortical Networks with Clustered Connections. *Nat Neurosci*. 2012; 15(11):1498–1505. <https://doi.org/10.1038/nn.3220> PMID: 23001062
25. Rostami V, Rost T, Riehle A, van Albada SJ, Nawrot MP. Spiking Neural Network Model of Motor Cortex with Joint Excitatory and Inhibitory Clusters Reflects Task Uncertainty, Reaction Times, and Variability Dynamics; 2020.
26. Jezzini A, Mazzucato L, La Camera G, Fontanini A. Processing of Hedonic and Chemosensory Features of Taste in Medial Prefrontal and Insular Networks. *J Neurosci*. 2013; 33(48):18966–18978. <https://doi.org/10.1523/JNEUROSCI.2974-13.2013> PMID: 24285901
27. Curti E, Mongillo G, La Camera G, Amit DJ. Mean Field and Capacity in Realistic Networks of Spiking Neurons Storing Sparsely Coded Random Memories. *Neural Comput*. 2004; 16(12):2597–2637. <https://doi.org/10.1162/0899766042321805> PMID: 15516275
28. Miller P, Katz DB. Stochastic Transitions between Neural States in Taste Processing and Decision-Making. *J Neurosci*. 2010; 30(7):2559–2570. <https://doi.org/10.1523/JNEUROSCI.3047-09.2010> PMID: 20164341
29. Katz DB, Nicolelis MAL, Simon SA. Gustatory Processing Is Dynamic and Distributed. *Curr Opin Neurobiol*. 2002; 12(4):448–454. [https://doi.org/10.1016/S0959-4388\(02\)00341-0](https://doi.org/10.1016/S0959-4388(02)00341-0) PMID: 12139994
30. Bouaichi CG, Vincis R. Cortical Processing of Chemosensory and Hedonic Features of Taste in Active Licking Mice. *J Neurophysiol*. 2020; 123(5):1995–2009. <https://doi.org/10.1152/jn.00069.2020> PMID: 32319839
31. Piette CE, Baez-Santiago MA, Reid EE, Katz DB, Moran A. Inactivation of Basolateral Amygdala Specifically Eliminates Palatability-Related Information in Cortical Sensory Responses. *J Neurosci*. 2012; 32(29):9981–9991. <https://doi.org/10.1523/JNEUROSCI.0669-12.2012> PMID: 22815512
32. Samuelsen CL, Gardner MPH, Fontanini A. Thalamic Contribution to Cortical Processing of Taste and Expectation. *J Neurosci*. 2013; 33(5):1815–1827. <https://doi.org/10.1523/JNEUROSCI.4026-12.2013> PMID: 23365221
33. Livneh Y, Ramesh RN, Burgess CR, Levandowski KM, Madara JC, Fenselau H, et al. Homeostatic Circuits Selectively Gate Food Cue Responses in Insular Cortex. *Nature*. 2017; 546(7660):611–616. <https://doi.org/10.1038/nature22375> PMID: 28614299
34. Lin JY, Mukherjee N, Bernstein MJ, Katz DB. Perturbation of Amygdala-Cortical Projections Reduces Ensemble Coherence of Palatability Coding in Gustatory Cortex. *Elife*. 2021; 10. <https://doi.org/10.7554/eLife.65766> PMID: 34018924
35. Stapleton JR, Lavine ML, Nicolelis MAL, Simon SA. Ensembles of Gustatory Cortical Neurons Anticipate and Discriminate between Tastants in a Single Lick. *Front Neurosci*. 2007; 1(1):161–174. <https://doi.org/10.3389/neuro.01.1.1.012.2007> PMID: 18982126
36. Dikeciligil GN, Graham DM, Park IM, Fontanini A. Layer- and Cell Type-Specific Response Properties of Gustatory Cortex Neurons in Awake Mice. *J Neurosci*. 2020; 40(50):9676–9691. <https://doi.org/10.1523/JNEUROSCI.1579-19.2020> PMID: 33172981
37. Gardner MPH, Fontanini A. Encoding and Tracking of Outcome-Specific Expectancy in the Gustatory Cortex of Alert Rats. *The Journal of neuroscience: the official journal of the Society for Neuroscience*. 2014; 34(39):13000–13017. <https://doi.org/10.1523/JNEUROSCI.1820-14.2014> PMID: 25253848
38. Fonseca E, de Lafuente V, Simon SA, Gutierrez R. Sucrose Intensity Coding and Decision-Making in Rat Gustatory Cortices. *eLife*. 2018; 7:e41152. <https://doi.org/10.7554/eLife.41152> PMID: 30451686
39. Durstewitz D, Vittoz NM, Floresco SB, Seamans JK. Abrupt Transitions between Prefrontal Neural Ensemble States Accompany Behavioral Transitions during Rule Learning. *Neuron*. 2010; 66(3):438–448. <https://doi.org/10.1016/j.neuron.2010.03.029> PMID: 20471356
40. Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J. Dynamic Coding for Cognitive Control in Prefrontal Cortex. *Neuron*. 2013; 78(2):364–375. <https://doi.org/10.1016/j.neuron.2013.01.039> PMID: 23562541
41. Rich EL, Wallis JD. Decoding Subjective Decisions from Orbitofrontal Cortex. *Nat Neurosci*. 2016; 19(7):973–980. <https://doi.org/10.1038/nn.4320> PMID: 27273768

42. Benozzo D, La Camera G, Genovesio A. Slower Prefrontal Metastable Dynamics during Deliberation Predicts Error Trials in a Distance Discrimination Task. *Cell Rep.* 2021; 35(1):108934. <https://doi.org/10.1016/j.celrep.2021.108934> PMID: 33826896
43. Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, et al. Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron.* 2014; 81(1):179–194. <https://doi.org/10.1016/j.neuron.2013.10.020> PMID: 24361077
44. Cisek P, Puskas GA, El-Murr S. Decisions in Changing Conditions: The Urgency-Gating Model. *Journal of Neuroscience.* 2009; 29(37):11560–11571. <https://doi.org/10.1523/JNEUROSCI.1844-09.2009> PMID: 19759303
45. Rabiner LR. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proc IEEE.* 1989; 77(2):257–286. <https://doi.org/10.1109/5.18626>
46. Prins J. NIST/SEMATECH e-Handbook of Statistical Methods; 2013.
47. Recanatesi S, Pereira-Obilinovic U, Murakami M, Mainen Z, Mazzucato L. Metastable Attractors Explain the Variable Timing of Stable Behavioral Action Sequences. *Neuron.* 2021; 110:139–153.e9. <https://doi.org/10.1016/j.neuron.2021.10.011> PMID: 34717794