

# Behaviorally relevant sensory cortical population dynamics in the rodent taste system

A Dissertation

Presented to  
The Faculty of the Graduate School of Arts and Sciences  
Brandeis University  
Neuroscience  
Prof. Donald B Katz, Advisor

In Partial Fulfillment  
of the Requirements for the Degree  
Doctor of Philosophy

by  
Narendra Mukherjee  
August, 2019

This dissertation, directed and approved by Narendra Mukherjee's committee, has been accepted and approved by the Graduate Faculty of Brandeis University in partial fulfillment of the requirements for the degree of:

**DOCTOR OF PHILOSOPHY**

Prof. Eric Chasalow, Dean of Arts and Sciences

Dissertation Committee:

Prof. Donald B Katz, Dept. of Psychology, Brandeis University, Chair

Prof. Leslie C Griffith, Dept. of Biology, Brandeis University

Prof. Shantanu P Jadhav, Dept. of Psychology, Brandeis University

Prof. Emery N Brown, Brain & Cognitive Sciences, MIT

©Copyright by  
Narendra Mukherjee  
2019

To the fond memory of Prof. Vijay Kumar Sharma, who first introduced me to the study  
of animal behavior

# Acknowledgments

The journey from a first-year PhD student to one finishing up their thesis is usually long and winding, and can only be completed with wisdom, help and support from a substantial pool of mentors and friends, both professional and personal. The amazing band of researchers inhabiting the Neuroscience labs at Brandeis University played that role in my PhD journey, and for that I will be forever grateful to them. I am also deeply grateful to the Howard Hughes Medical Institute (HHMI) for supporting much of my PhD work at Brandeis University through the International Predoctoral Fellowship program.

In Prof. Donald B Katz (“the researcher formerly known as Don”, as his office door proclaims), I was lucky to have found a mentor, collaborator, advisor and parent rolled into one! Don is a groovy scientist (by his own admission) and groovier still as a PhD advisor - I will be eternally grateful to him for trying his very best to drive the spirit of enquiry and scientific thinking into his students’ philosophy in life. Don has taught me to be critical and circumspect (but kind) in science and otherwise, skills that I will rely on whatever I do in my life. Above all, Don has an incredible capacity to look beyond himself (and his goals) when dealing with students, and helping his students to develop the scientific (and often social) skills that they want in their future careers, academic or otherwise. Thank you Don for everything you have done over the last several years (and for not being another of those PIs toting Apple products :) ).

To the incredible humans living in the BLECh (Katz) lab at Brandeis University, I can only express my deepest gratitude. They ended up being my first (and best) supporters and critics in a foreign land, and made lab a really fun place to spend time in! Special thanks to Prof. Veronica Flores and Dr. Meredith Blankenship, my PhD classmates and labmates, for being the wonderfully skilled scientists that they are and collaborating on so much of the day-to-day science in the Katz lab. Honestly, there is a very long list of people who have

passed through the Katz lab and influenced my scientific journey, so I will name a subset for now: Anan, Joost, Jenn, Aaron, Emma, Linnea, Zhiqiang, Maria, Max, Yasmin, Daniel.

My closest collaborator (and co-author) in lab, Joseph “Joe” Wachutka has played a critical role in making all my experiments happen. Joe was the driving force behind the design and implementation of a new state-of-the-art electrophysiology system in the Katz lab that was essential in performing the studies in [chapter 2](#) (in addition to being a constant source of knowledge about American politics and society). I have also been lucky to have had the guiding patience and experience of Dr. Jian-You Lin, the energy of Bradly Stone and the enthusiasm of Abuzar Mahmood in several of my projects, collaborative or otherwise!

Outside of lab, I am thankful to the city of Boston for everything that it offers to the students and researchers living in/around it. The universities in the Boston-area have fostered a collaborative and open culture that is unseen anywhere else in the world, and I am grateful to have had the opportunity to experience it first hand. In addition, Boston’s focus on accessible transportation and bicycling facilities first led me to bike around in the city, and then transformed biking to my only (very) serious hobby! The huge support (and spread of knowledge) from the bicycling community in Boston and adjoining areas will be some of my fondest experiences in the city. The city also led me to some of my closest friends, people that I will continue to learn from and depend on - I want to express my gratitude to Sumantra, Archana, Mohit and Akshita for being with me on both the good and bad parts of the PhD journey.

Finally (and most importantly), I cannot emphasize enough the role of my lovely and wonderful partner, Shrabastee Banerjee, in pushing me to strive for the best in science and in life. Shrabastee is an incredibly skilled statistician and a gifted writer, and there is no way that this thesis (or the papers within it) would have contained the statistical analyses and prose that it does without her input! Thank you making me (finally!) finish the Acknowledgements page too, and for moving to Boston in the first place!

# Abstract

## **Behaviorally relevant sensory cortical population dynamics in the rodent taste system**

A dissertation presented to the Faculty of  
the Graduate School of Arts and Sciences of  
Brandeis University, Waltham, Massachusetts

By Narendra Mukherjee

In the studies that are part of this dissertation, we investigated the role of primary sensory taste cortex in generation of behavioral responses to taste stimuli. Starting with a general description of prior work and formulation of a roadmap for the thesis, each chapter describes a separate investigation of the connection between cortical processing of sensory stimuli, appropriately-timed behavioral outputs and associative learning based on sensory experiences. The chapters each start with their own abstract, specific to the study being reported.

# Contents

<b>Abstract</b>	<b>vii</b>
<b>Introduction</b>	<b>3</b>
<b>1 The behavioral relevance of cortical neural ensemble responses emerges suddenly</b>	<b>4</b>
1.1 Abstract . . . . .	4
1.2 Introduction . . . . .	5
1.3 Materials and Methods . . . . .	7
1.4 Results . . . . .	18
1.5 Discussion . . . . .	27
1.6 Figures . . . . .	30
<b>2 Dynamical structure of cortical taste responses revealed by precisely-timed optogenetic perturbation</b>	<b>38</b>
2.1 Abstract . . . . .	38
2.2 Introduction . . . . .	39
2.3 Materials and Methods . . . . .	41
2.4 Results . . . . .	59
2.5 Discussion . . . . .	68
2.6 Figures . . . . .	72
<b>3 The role of the gustatory cortex in incidental experience-evoked enhancement of later taste learning</b>	<b>80</b>
3.1 Abstract . . . . .	81
3.2 Introduction . . . . .	81
3.3 Materials and Methods . . . . .	83
3.4 Results . . . . .	94
3.5 Discussion . . . . .	101
3.6 Figures . . . . .	107

## CONTENTS

<b>Conclusion</b>	<b>118</b>
<b>A Python meets systems neuroscience: affordable, scalable and open-source electrophysiology in awake, behaving rodents</b>	<b>119</b>
A.1 Abstract . . . . .	119
A.2 Introduction . . . . .	120
A.3 Experimental paradigm and hardware . . . . .	122
A.4 Electrophysiology in systems neuroscience . . . . .	123
A.5 Scientific Python stack for data analysis – spike sorting . . . . .	124
A.6 Discussion . . . . .	133
A.7 Figures . . . . .	135

# Introduction

The study of neural activity in the mammalian brain and its role in animal behavior has a long and rich history in systems neuroscience. Beginning with Hubel and Wiesel’s revolutionary studies of the stimulus tuning properties of neurons in the primary visual cortex (V1, [Hubel and Wiesel \(1962\)](#)), systems neuroscience has focused on elucidating the electrophysiological properties of single neurons as the unit of behaviorally-relevant computation in the brain. These investigations have established the brain (and especially the cortex) as being composed of spatially distinct computational “modules” - with computations related to stimulus encoding, decision-making, learning and behavioral action delegated to separate brain regions.

However, this traditional, single-neuron-centric view of activity in the brain makes strong (and unrealistic) assumptions. For one, the computationally modular model of the brain has been constructed out of analyses that average the firing of individual neurons across repeated trials, and across large windows of time. Such trial-and-time averages neglect the inhomogeneous nature of neural spiking activity and characterize the firing of individual sensory cortical neurons entirely in terms of stimulus selectivity “tuning curves”. When extended to take the temporal evolution of neural firing rates into account, single neuron analyses lock trials to an externally imposed time point (e.g stimulus onset time) before averaging, thereby wiping out the influence of an internal (to the brain) behaviorally-relevant time point that might actually be modulating firing rates. Finally, both of these approaches, by using neuron-level averages, dismiss ensemble-level coordinated activity and shared variability as “noise” - noise that is becoming increasingly apparent as new experimental technologies make it commonplace to record large populations of neurons (even across brain regions) simultaneously.

In this thesis, we question the divide between sensory encoding, decision-making, learning

## CHAPTER 0. INTRODUCTION

and motor action in the brain by using the rodent taste system as a model. Unlike other, more commonly studied sensory systems (like vision and audition), the encoding of tastes is inherently connected to decision-making and action. Taste stimuli are uniquely laden with behavioral valence-related (“palatability”) information in addition to their chemical “identities” and necessitate an ingestion/egestion (or swallow/spit) decision once they land on the tongue. This makes taste processing an ideal system to study the connection between the stimulus-responsive dynamics of sensory ensembles and the motor program initiated by the animal in response to the stimulus. In addition, the changes wrought in this connection between sensory processing and motor action through learning and context can be mapped using conditioned taste aversion (CTA), a single shot learning paradigm that can reliably make a previously palatable stimulus aversive as well as shift the animal’s motor responses to this stimulus from ingestion to egestion ([Chambers \(2018\)](#)). Finally, the relatively slower time evolution of stimulus responses in the primary taste cortex (gustatory cortex, GC: 2-3s post stimulus, [Katz et al. \(2001\)](#)) compared to visual or auditory responses (<500ms post stimulus) makes it easier to build, interpret and test ensemble/network-level models of sensory activity that can connect to the animal’s eventual consumption decisions.

In [chapter 1](#), we show that peri-stimulus averaging of single neuron firing massively underestimates the speed and coherence with which GC taste responses evolve. We extend prior work using Hidden Markov Models (HMMs, [Jones et al. \(2007\)](#)) as a generative description of ensemble-level activity patterns and find that taste-responsive firing in GC is composed of a stereotypical sequence of quasi-stable states, with sudden, coherent (and nearly instantaneous) transitions between them. These ensemble activity state transitions, though stereotypical in their temporal sequence, are highly variable in their timing from trial-to-trial, and involve the coordinated participation of the large majority of recorded neurons. We show that the last ensemble activity state in this sequence, emerging anywhere between 0.5-1.8s post stimulus on individual trials, is rich in palatability(the decision variable)-related firing. We use this result to suggest a novel way of averaging taste responsive firing across repeated trials: a “peri-transition time” analysis that locks trials to the emergence of the palatability-related state and is disconnected from experimenter-imposed time points like stimulus onset. Finally, we find that the palatability-related ensemble state, despite being variably-timed, reliably precedes the onset of the animal’s ingestion-egestion related mouth movements (orofacial behaviors) by  $\sim 0.3s$  on individual trials. Taken together, these results demonstrate that population-level variability in sensory cortical taste responses, previously

## CHAPTER 0. INTRODUCTION

dismissed as “noise”, is tightly coupled to the timing of taste-reactive behaviors.

We subject the phenomenological results of [chapter 1](#) to a test of causality in [chapter 2](#). We use brief (0.5s), precisely timed optogenetic perturbations of GC population activity to investigate the role of sensory cortical neurons in driving the animal’s taste-reactive orofacial behaviors. We find that the variably-timed onset of palatability-related ensemble firing in GC is a unique “moment in time” that is behaviorally relevant but disconnected from any external, experimenter-imposed time points - perturbations of GC activity influence the timing of orofacial behaviors only if they arrive *before* this moment in time. In fact, identically timed perturbations of GC activity have drastically different effects on orofacial behavior depending on the ensemble activity state prevailing right before the perturbation began. These results stand in stark contrast to the sensory-motor divide of the modular view of processing in the brain - we show that taste processing is more akin to an attractor network, with sensory cortex being a node of this network that is not just limited to the task of stimulus encoding.

We extend these results to the realm of learning in [chapter 3](#) to show that GC neural populations, through their temporally dynamic responses, play a crucial role in keeping track of the animal’s taste experiences in addition to connecting stimulus encoding with behavioral action. We find that non-reinforced, incidental taste experiences can greatly enhance the strength of aversion learning with novel taste experiences and use targeted optogenetic inactivations to show that GC is an essential player in this process. This effect cannot be attributed to a simple change in stimulus salience wrought by GC inactivation - instead, when taken together with the results of [chapter 1](#) and [chapter 2](#), these findings lend support to the idea of taste experiences (including those that are non-reinforced) modifying the proposed attractor network (of which GC is a part) that transforms sensory processing of tastes to appropriate taste-reactive orofacial behaviors.

For the interested reader, [Appendix A](#) contains a detailed description of the open-source electrophysiological setup (both hardware and software) that was developed in-house for the experiments in this thesis.

# Chapter 1

## The behavioral relevance of cortical neural ensemble responses emerges suddenly

### Co-author contributions

The experiments and results from this chapter were published in the Journal of Neuroscience as [Sadacca et al. \(2016\)](#).

The following people contributed to this study (and were co-authors on the published paper):

1. Brian F. Sadacca: Designed research, performed research, analyzed data, wrote paper.
2. **Narendra Mukherjee**: Designed research, performed research, analyzed data.
3. Tony Vladusich: Analyzed data.
4. Jennifer X. Li: Performed research.
5. Donald B. Katz: Designed research, analyzed data, wrote paper.
6. Paul Miller: Performed research, analyzed data, wrote paper.

### 1.1 Abstract

While many laboratory-studied decisions involve a highly-trained animal identifying an ambiguous stimulus, many naturalistic decisions do not. Consumption decisions, for instance,

involve determining whether to eject or consume an already-identified stimulus in the mouth - a decision that can be made without training. By standard analyses, rodent cortical single-neuron taste responses come to predict such consumption decisions across the 500ms preceding the consumption or rejection itself; decision-related firing emerges well after stimulus identification. Analyzing single-trial ensemble activity using Hidden Markov Models (HMM), we show these decision-related cortical responses to be part of a reliable sequence of states (each defined by the firing rates within the ensemble) separated by brief state-to-state transitions, the latencies of which vary widely between trials. When we aligned data to the onset of the (late-appearing) state that dominates during the time period in which single-neuron firing is correlated to taste palatability, the apparent ramp in stimulus-aligned choice-related firing was shown to be a much more precipitous coherent jump. This jump in choice-related firing resembled a step function more than it did the output of a standard (ramping) decision-making model, and provided a robust prediction of decision latency in single trials. Together, these results demonstrate that activity related to naturalistic consumption decisions emerges nearly instantaneously in cortical ensembles.

## 1.2 Introduction

Most neuroscientific studies of decision-making involves training animals to recognize some physical property of a stimulus - to saccade or poke its nose discriminatively once the identity of a stimulus has been determined. Such studies have revealed important neural correlates of decisions: firing rates in relevant neurons begin to reflect the decision within 200ms, appearing to “ramp” to a decision threshold with a slope commensurate to task difficulty (Huk and Shadlen (2005), Kiani et al. (2008), Shadlen and Newsome (2001), but see Latimer et al. (2015)).

Taste-related decisions, in contrast, require no training: when a taste is presented, even via intra-oral cannulation, a rat necessarily makes a decision: it rates the “palatability” of the taste, deciding whether the taste object is worthy of consumption; this psychological property is exquisitely sensitive to a large range of experiential variables (Berridge and Grill (1984), Breslin et al. (1992), Fortis-Santiago et al. (2010), Galef (1986), Spector et al. (1988)), and as such can only be made subsequent to identification of stimulus identity. Thus, even easy taste-palatability decisions often have 5-10 times the latency ( $\sim 1.0$  sec) as similarly

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

easy stimulus-identification decisions ([Hanes and Schall \(1996\)](#), [Ratcliff et al. \(2003\)](#), [Uchida and Mainen \(2003\)](#)).

Our studies of this paradigm have focused on gustatory cortex (GC) and other fore-brain structures situated between basic sensory regions and the motor centers controlling the production of discriminative oral behaviors that signal that a decision has been made. Taste-related firing in GC jibes well with the above description of the decision process: identity-related firing emerges within 200 milliseconds of stimulus presentation, 800ms before the production of decision-specific orofacial behaviors. Between these two time points, GC firing reorganizes, coming to reflect the palatability decision in a manner that, according to standard analyses, begins  $\sim$ 500ms following taste delivery, becomes significant at  $\sim$ 800ms, and asymptotes prior to the emission of choice-specific behavior (e.g., [Piette et al. \(2012b\)](#), [Sadacca et al. \(2012\)](#)). That is, GC decision-related activity emerges in an apparent “ramp” preceding behavior, much as activity in LIP ramps preceding visual motion perception decisions.

In the case of GC taste responses, however, this ramp may be an artifact of across-trial averaging. Ensemble recordings in awake rats have suggested that taste responses are actually better-described as sequences of quasi-stable firing-rate “states”, with inter-state transitions that occur suddenly, but at highly variable trial-to-trial latencies ([Escola et al. \(2011\)](#), [Jones et al. \(2007\)](#), [Moran and Katz \(2014b\)](#)). Because firing-rate ramps appear artifactually in across-trial averages under such circumstances ([Miller and Katz \(2010\)](#)), it becomes reasonable to ask if, when “properly” analyzed, GC responses will predict consumption decision-making with analogous suddenness, rather than following a substantial “ramp”.

Testing this hypothesis requires analysis of simultaneously-recorded ensembles of neurons, because sudden firing rate shifts are notoriously difficult to discern when single neurons are recorded individually ([Okamoto et al. \(2007\)](#)). Here we performed precisely this analysis, applying a tool (Hidden Markov modeling - HMM) that identifies discrete states underlying ensembles of spike trains ([Abeles et al. \(1995\)](#), [Bollimunta et al. \(2012\)](#), [Jones et al. \(2007\)](#), [Kemere et al. \(2008\)](#), [Moran and Katz \(2014b\)](#), [Ponce-Alvarez et al. \(2012\)](#)) to GC taste responses. By this analysis, decision-related firing is indeed revealed to emerge in sudden firing-rate transitions, which, while varying widely in latency from trial to trial ([Jones et al. \(2007\)](#)), are more similar to transitions predicted by a step-function model than by the dominant ramping model of decision-making (the drift diffusion model, see [Hanes and Schall \(1996\)](#), [Huk and Shadlen \(2005\)](#), [Shadlen and Newsome \(2001\)](#), [Thompson et al. \(1996\)](#)).

Furthermore, the latency of this state in single trials provides an excellent prediction of the onset of choice-related behavior. Thus, these data suggest that the evidence of a decision to consume or expel a taste stimulus appears in cortex with a suddenness approaching instantaneity.

## 1.3 Materials and Methods

### 1.3.1 Experimental design

#### Subjects

Female Long–Evans rats (n=11 [2 in initial modeling, 9 with EMG]); 280–320 g at time of surgery) served as subjects in this study. Rats were maintained on a 12 h light/dark schedule and were given *ad libitum* access to food (and restricted access to water where specified). All methods complied with the Brandeis University Institutional Animal Care and Use Committee guidelines.

#### Surgery

Rats were anesthetized using an intraperitoneal injection of a ketamine/xylazine/acepromazine mixture (100 mg/kg, 5.2 mg/kg, and 1 mg/kg, respectively), with supplemental intraperitoneal injections administered as needed. The anesthetized rat was placed in a standard stereotaxic device, where its scalp was excised, and holes were bored into its skull for the insertion of 0-80 ground screws and electrodes. Multielectrode bundles (16 nichrome microwires attached to a Microdrive, [Katz et al. \(2001\)](#)) were inserted 0.5 mm above gustatory cortex (GC). Once in place, the assemblies were cemented to the skull, along with two intraoral cannulae (IOC, [Katz et al. \(2001\)](#)), using dental acrylic.

#### Passive taste administration paradigm

Three days following surgery, each animal began 2 days of adaptation to handling. Afterward, each animal was placed on a water-restriction regimen (2 hour of water/day) for 2 days, acclimatized to the experimental environment for 2 days, and adapted to 40- $\mu$ l water deliveries through the IOC for another 2 days. Once so acclimated, animals were, once per day, exposed to the experimental taste array (distilled water, four concentrations of NaCl

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

[0.01, 0.1, 0.3, 1.0M], 0.3M sucrose, and 0.001M quinine) through a manifold of fine polyimide tubes inserted to 0.5 mm past the end of one IOC (eliminating any chance of mixing), and locked onto the dental acrylic cap. All fluids (including the water rinse, which was delivered through the contralateral IOC) were delivered under slight nitrogen pressure; while delivering each taste from one side may have meant not-entirely immediate exposure of all taste buds, the pressure ensured that a brief ( $\sim 40$ ms—the ejection of taste onto the tongue was complete long before any taste-related dynamics appeared in GC responses) release of fluid resulted in extensive tongue coverage at reliably short latency (Katz et al. (2001)) and the use of a single manifold ensured essentially identical presentation of all taste stimuli.

Rats received a minimum of 10 blocks of taste deliveries (6 deliveries per block). Computer-controlled solenoid valves ejected a pseudorandomly selected taste directly into the mouth of the rat under nitrogen pressure once every 30 seconds. A water rinse was delivered through the contralateral cannula 15 s after each taste delivery. Total fluid delivered was 4.8 mL per 30 minutes of recording session, after which animals had *ad libitum* access to water for 90 minutes.

### Assessing preferences/palatability for the full array of taste stimuli

A set of rats ( $n = 4$ ) was adapted to handling and placed on a 22 hour water restriction protocol, with water provided in the home cage after handling, adaptation or testing. Testing took place in the Davis MS-160 “brief access” Lickometer rig (DiLog Instruments, Tallahassee, FL). During the first 2 days of habituation, rats were placed in the Davis rig and allowed to drink water from a single tube continuously for 30 minutes. On the last 2 days, the rat received periodic brief access (15 seconds) to one of 7 stainless steel drinking tubes on a moveable carousel, each filled with water, for 35 minutes. Finally, each rat received three 35-min testing sessions, on consecutive days, during which taste solutions (0.00M, 0.01M, 0.10M, 0.30M, and 1.0M NaCl, 0.3M sucrose, and 0.001M quinine HCl) were presented in a (blocked) randomized order. Presentations began with the automated raising of a shutter, such that the lick spout was exposed. If no lick was recorded within 60 seconds of spout presentation, the shutter closed and the tube holder moved on to the next tastant; these empty trials were dropped from the subsequent analysis. Once a lick was noted (via a low-current circuit), the solution was presented for 15 seconds (this guaranteed that lick counts were not confounded with latency to first lick), after which the shutter came back down, and

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

a 10-second interval separating each presentation began.

The average number of licks across the 15 seconds of availability, compared to that for water, provided us with a measure of preference and palatability of the tastes (Breslin et al. (1993)) without a disproportionate amount of lick-rate adaptation (Smith et al. (1992)); solutions preferred compared to water are here characterized as “palatable” whereas solutions that rats drank less than water are characterized as “aversive”. There was no decrement in lick rate across the 35 minute sessions, suggesting little influence of post-ingestive effects during taste preference assessment (data not shown).

### Electrophysiology

Neural signals were differentially recorded from GC during taste sampling, and fed into a parallel processor capable of digitizing up to 32 signals at 40 kHz simultaneously (Plexon, Dallas, TX). Discriminable action potentials ( $>3:1$  signal:noise ratio) were isolated on-line from each signal using an amplitude criterion in cooperation with a template algorithm, time-stamped records of stimulus onset, spike times, and all sampled spike waveforms were saved to disk, as was a file of discrimination parameters. All signals were then subjected to off-line re-analysis incorporating three-dimensional cluster-cutting techniques that confirmed and corrected on-line discriminations. A great deal of previous work has confirmed that this set of procedures results in the isolation of single-neuron records (Katz et al. (2001), Fontanini and Katz (2006)).

### 1.3.2 Analysis and response modeling

We obtained data from seven separate multi-neuronal recording sessions, as described above. Identified spiking events were partitioned into 10ms bins (250 bins per 2.5 seconds of post-stimulus activity) for analysis. To maintain equal samples across recording sessions, the first 10 trials of each taste delivery were included in subsequent analysis.

### Palatability correlation and taste identification

To determine the relationship between neural activity and behavior through time, we calculated the linear correlation across stimuli between each neuron’s trial-averaged activity and the palatability of the tastant as measured in the brief access task. A large number of

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

studies have established the general palatability function for these stimuli - a function that is remarkably reliable across different measures, contexts, and physiological conditions - but we used recently collected data from our lab (Sadacca et al. (2012)), which showed, as expected, the following order of palatability, from most pleasing to most aversive: sucrose, 0.1 M NaCl, 0.3 M NaCl, 0.01 M NaCl, 1.0 M NaCl, quinine HCl (the last two of these stimuli were less palatable than distilled water). The inverted-U for palatability as a function of NaCl concentration demonstrates that palatability is a measure of an animal’s response to a stimulus (which takes time to be determined) rather than a straightforward ingredient of the stimulus. That is, palatability is a measure of how much or how often an animal chooses to consume a tastant. Albeit measured in a separate battery of tests, palatability is equivalent to an average of the animal’s choices across trials, which we will correlate with an average of each neuron’s activity across trials.

Results based on this measure are thus intrinsically conservative. We neither trained our rats to produce stereotyped categorical responses nor selected trials in which those behaviors were easily recognized as correct; our inability to remove “error trials”, if it has any effect at all, serves only to reduce the strength of the correlations between neural activity and choice. The main finding from this work arises from comparisons between, on the one hand, the results of trial-averaging when each trial was aligned to stimulus delivery (producing the peri-stimulus time histogram, PSTH), and, on the other hand, analogous results observed when each trial was aligned to a state transition (producing the analogous peri-transition time histogram, PTTH, see below). Regardless of alignment, the neural firing rates for each neuron, for each 250ms window of activity, were transformed into a neural response vector of length equal to the number of tastes x the number of trials. A palatability vector of equal length was then created, with each value of the palatability vector determined by that taste stimulus’ value on Figure 1.1A. Stepping this window in increments of 10ms, we produced a time series of correlation coefficients between each neuron’s trial-averaged activity and the palatability (Figure 1.1C). The palatability or choice index,  $I(t)$ , reported is simply the mean of the square of the correlation coefficient across neurons, through time  $t$  (Figure 1.2).

To assess the availability of any taste-related information in the neural ensemble (as compared to the presence of palatability-specific information) we used a standard linear-discriminant analysis classifier to test the reliability with which evoked responses to a taste could be identified amongst responses to other tastes. Here, we binned neural responses into 50ms bins, and used a linear classifier with a leave-one-out approach, training the classifier

on all-but one trial of two pairs of evoked population taste responses (e.g. trials 1-9 of sucrose and quinine for all neurons) and tested with the ‘left-out’ trial (here, trial 10 of sucrose and quinine). This analysis was then iterated for each time bin using the mean of 5 bins of evoked response per step (250ms), and for each combination of exemplar tastes (sucrose, quinine, 0.1M NaCl, and 1.0M NaCl).

### Characterizing the rise of the palatability index.

Differences between palatability time series (e.g. PSTH vs. PTTH) for population and single-neuron data were quantified by fitting a 4 parameter sigmoid function to the palatability index:

$$I(t) = \frac{\frac{\alpha}{\beta}}{1 + e^{-\beta(t-t_0)}} + \delta \quad (1.1)$$

$\alpha$  is equal to four times the maximum slope (making  $\frac{\alpha}{4}$  our y-axis measure of slope),  $\beta$  is the “suddenness” of the increase (equal to the inverse of the period during which the slope was  $> 78\%$  of maximum),  $t_0$  is the time when maximal slope is reached, and  $\delta$  is the offset in minimum palatability correlation from zero (a baseline value). Note that the combination in the numerator,  $\frac{\alpha}{\beta}$ , denotes the overall change in palatability across the transition. To obtain consistent fits for both population and single-neuron, parameters  $\alpha$  and  $\beta$  were completely unconstrained while parameter  $t_0$  was restricted to 200ms around the time of transition, and  $\delta$  was held to a max value of 0.05. In addition to this fitting, we performed a secondary measure of transition suddenness: the time taken to transition across 40% of the total jump from minimum to maximum palatability. Values of this secondary measure broadly agreed with the fitted  $\frac{1}{\beta}$ . To make control data aligned to a different, earlier transition (either an early-dominant state or the transition immediately preceding the identified late state) comparable, the mean (early) transition time was subtracted, and the data were shifted by the mean (late) transition time.

### Hidden Markov Modeling and Model Selection

We performed HMM, limiting each analysis to simultaneously-recorded ensemble data, on the basis of its proven ability to characterize neural activity (Jones et al. (2007), Kemere et al. (2008)), and specifically on the basis of our previous demonstration that taste responses are

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

well-characterized as sequences of ensemble firing rate states with brief inter-state intervals (Jones et al. (2007), Piette et al. (2012b)). We prepared our data for HMM by first indexing each neuron in the ensemble with a scalar ( $N = 1, 2, 3 \dots M$ ). If only one neuron spiked in a time bin, we assigned the number  $N$  to that event, with  $N = 0$  corresponding to no spikes from any neuron. If more than one neuron spiked in a time bin, we randomly selected one of the spiking neurons for assignment to that bin—a highly uncommon occurrence, given the relatively low firing rates, small (10ms) bins, and small (6-14) ensembles of neurons. We used standard MATLAB packages for HM modeling. Using 25 different sets of random starting parameter values as seeds for a standard model optimization procedure (the Baum-Welch algorithm, Baum et al. (1970)). The model with maximum log likelihood, or LL, (calculated as the log probability of producing the measured spike trains given the particular model) was provisionally treated as the optimal characterization of ensemble activity, contingent to the application of a model selection technique (the Akaike Information Criterion, AIC), that penalizes the LL value according to the number of parameters in the model (the size of the emission matrix minus one column plus the size of the transition matrix minus one column). While the original LL-based models always produced better fits to the data than the AIC-based models, the penalty term provided by the evaluation of the number of free parameters militated against over-fitting the data. Typically the AIC method brought to bear on this length of time series selected four or five states as the maximally informative number, numbers that accord well with our previous empirical findings—three states max in the first 1.5s, matching the number of epochs, and then 1-2 post-consumption states (Jones et al. (2007), Piette et al. (2012b)). We produced a single HM model for each ensemble and each tastant. Given this model, we calculated, for each trial, the probability as a function of time of the ensemble being in any particular HM state.

### Post-HM Realignment

For each HMM, we determined the putative underlying state with the highest probability of occurring across all trials within a time window identified, on the basis of Figure 1.1 and previous work (Grossman et al. (2008), Katz et al. (2001), Sadacca et al. (2012)), as being the time at which rising ramps of palatability, observed using analyses keyed to stimulus delivery, reach asymptote (between 0.8-2 seconds post stimulus). These states were deemed the most likely candidate “palatability states”. We next realigned the ensemble data such

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

that the onset times of these states (calculated as the time bin at which the identified state exceeded 0.5 probability on each trial) was set to be the “zero time-point” of each trial. The few trials in which the onset of this state occurred before 0.1 seconds following stimulus onset were excluded from this (and all) analysis, as were paired trials from stimulus aligned data (to maintain equal trial number between PSTH and PTTH datasets). For two control realignments, data were also zeroed on an earlier state onset: either the state with the highest probability between 0.25 and 0.8 seconds post taste delivery (labeled ‘early’), or the state immediately prior to the calculated ‘late’ palatability state for each model (labeled ‘pre-late’).

Once the data had been realigned, we repeated the above-described palatability analyses that had been already brought to bear on “stimulus-aligned” data. To facilitate an initial visual comparison of the two analyses in [Figure 1.3](#), the time at which the correlation between palatability and neural responses became significant (860ms following stimulus delivery, [Piette et al. \(2012b\)](#), [Sadacca et al. \(2012\)](#)) in stimulus-aligned data was aligned with transition time.

### Monte Carlo Analyses

We conducted Monte Carlo analyses to deal with a complementary pair of issues having to do with the above procedures: 1) the fact that rapid transitions from low-to-high neural correlations with palatability could conceivably be produced as an artifact of HMM itself, which will identify the sharpest transition time even in data for which all “transitions” are purely stochastic and noise-based; and 2) the fact that the correlation and HMM procedures will artificially smooth even instantaneous transitions to some degree—the former because of the necessity of data windowing, and the latter because of the imprecision inherent in estimating the times of hidden state changes on the basis of relatively sparse point-process data putatively emitted as a function of the underlying state. We therefore generated three sets of 100 ‘dummy’ simulations of each trial in the real data, to evaluate these two possibilities and to test an alternative theory, the drift-diffusion model; all three sets of simulated data resulted from probabilistically chosen spike trains which approximated the inter-spike interval statistics of the real data, and by design contained essentially the same firing rates (and therefore produced similar PSTHs) observed in the real data.

**PSTH- and shuffled-based control data.** The first of these simulations was produced by assuming a neuron’s spike times in each trial to emerge purely from the temporally inhomogeneous Poisson processes observed in the across-trial averaged single neuron responses to each taste (i.e., the PSTHs); these data matched the trial-averaged real data but contained neither between-neuron coherence nor purposefully sudden rate changes—thus in these datasets the PSTHs were truly (and by definition) a valid characterization of each trial—and therefore tested the possibility that any “sharpening” of the onset of palatability-related firing was an artifact of the analysis. The second of these simulations was a trial-reshuffling of the original data for each neuron (e.g. pairing trial 1 of neuron 1 with trial 9 of neuron 2), to again disrupt between-neuron coherence while preserving the PSTH. Specifically, for each real dataset, we created 500 permutations of potential trial-neuron pairings for the ensemble, and selected the 100 that minimized the number of coherent neurons for each reshuffled trial. We subjected these control datasets to the same analysis as performed with the original dataset, statistically comparing the transition speed of control datasets with the original.

**Coherent step-function control data.** The third simulation allowed us to assess the degree to which the speed of transitions in the realigned real data differed from the theoretical maximum speed produced by instantaneous, correlated changes in firing rates, given the firing rates observed in our cortical neuron sample. For these 100 ‘dummy’ datasets, firing rates were obtained from the mean PTTH during five different response periods,  $t$ , relative to average stimulus onset: 1) a pre-stimulus period between -1s pre-stimulus and stimulus onset; 2) an early epoch period from  $t = 0$  to  $t = 200ms$ ; 3) a middle epoch period from  $t = 300ms$  to  $t = 800ms$ ; 4) a final period from  $t = 1200ms$  to  $t = 1400ms$ ; and 5) a post-stimulus period from  $t = 2000ms$  to  $t = 2800ms$ . We assumed coherent, instantaneous jumps in firing rates between these values, with simulated states lasting 1, 0.2, 0.6, 1.2, and 2.0 seconds, on average, however, the onset of the second-to-last transition was jittered, with the jitter pulled from a normal distribution with a standard deviation equal to the standard deviation of the ‘late’ transition (413ms). We again subjected this control dataset to analyses identical to those brought to bear on the real data, and compared the suddenness of the palatability index with those produced by the equivalent analyses of the original data.

**Drift diffusion model control data.** The fourth simulation was designed to assess how well a standard model of decision-making, the drift diffusion model (DDM), fits our observed data. The DDM relies upon neural integrators, which can be implemented in numerous ways, some of which require fine-tuning (Wang (2002), Wong and Wang (2006)) and others of which produce instantaneous jumps (that are not coherent across the circuit) in neural firing rates (Koulakov et al. (2002)).

The DDM is identical to the step-function control, in that it incorporates the mean firing rates, trial-to-trial variability, and firing-rate correlations across all cells, but differs in that the rates in each trial can ramp in the manner of a biased random walk. In order for the DDM to match our data, it was essential that we included a delay of 500ms before any ramping of the decision process could commence. Such an initial delay is much longer than used in standard decision-making models (Ratcliff et al. (2003)) and its inclusion already shifts the DDM toward a model more akin to the step-function control.

The DDM itself was defined by two parameters, mean time to threshold ( $T=0.6s$ ) and fraction of correct trials ( $P(\text{correct})=0.975$ ), which together constrain the “stimulus strength”,  $S$ , and the “noise level”,  $\sigma^2$ , of the decision variable for a given choice threshold (whose arbitrary value,  $a=1$ , simply scales the decision variable, so scales the mapping to firing rates) (Bogacz et al. (2006)):

$$T = \frac{a}{S} \tanh \frac{Sa}{\sigma^2} \tag{1.2}$$

$$P(\text{correct}) = \frac{1}{1 + e^{-\frac{2Sa}{\sigma^2}}}$$

The model was mapped into the firing rates of cells by assuming each cell’s firing rate was a sigmoidal function of the decision variable (the sigmoid was essential so that rates remained positive) with the sigmoid set independently for each cell to produce the cell’s pre-palatability-response firing rate when the decision variable is zero and to produce the cell’s post-palatability-response firing rate when the decision variable is at the “correct” threshold,  $+a$ . To completely define the parameters of the sigmoid for each cell we further assumed the cell’s firing rate to be bounded between zero and 110% of the greater of pre- and post-palatability-response rates.

Thus in any trial, all cells’ firing rates would coherently follow the single decision variable for that trial according to the biased random walk, but the single decision variable was mapped into a sigmoidal firing rate unique to each cell so as to best match that cell’s

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

observed rate changes. While the sigmoid functions were typically broad, they do add non-linearity, so make the DDM more step-function-like than if we assumed firing rates were a linear function of the decision variable. Such a non-linear mapping was essential in order to produce the observed firing rates in the stimulus-aligned data. This fact, combined with our incorporation of a large (500ms) delay before ramping commenced (to match the timing of the onset in ramping without disrupting other transitions) means that our test of whether the data was more “step-like” or more like a “ramping” DDM required us to first add two “step-like” features to the “ramping” DDM (the delay and the nonlinearity). Such modifications were essential in order to provide a framework in which the DDM was at all capable of fitting the observed data, but potentially they cause us to underestimate the difference between DDM-derived and ‘instantaneously-transitioning’ simulations.

**Control data comparison.** The above pair of analyses is particularly felicitous because it simplifies the difficult task of assessing statistical significance of the phenomena reported here. In each case, the proportion of dummy datasets was a direct indicator of that significance - if, for instance, the slope of the real data was higher than that of more than 190 of 200 simulations (each of which is identical in size to the real data), then in a 1-tailed test (which we use here because our hypotheses are explicitly directional) the p-value for the difference between the conditions is less than 0.05.

Finally, to directly compare how well the control datasets fit the real data, we determined the likelihood that either of the control datasets were distributions centered on the real data (z-test), and then computed the likelihood ratio ( $\Lambda$ ) between the optimal and PSTH-derived simulations, where:

$$\Lambda = 2 \ln p_{\text{OPTI}} - 2 \ln p_{\text{DDM}} \quad (1.3)$$

with significance of the likelihood ratio assumed to follow a chi-squared distribution.

### Comparison of electromyographic (EMG) and neural population data

We identified palatability-related oral behaviors (taste reactivity or TR, see [Grill and Norgren \(1978a\)](#), [Travers and Norgren \(1986\)](#)) in jaw-movement EMG signals. These behaviors are commonly understood to represent the execution of a consumption decision via the action of a single central pattern generator in the brainstem: gapes (focused upon here because they are the largest, easiest to isolate motor acts) represent the movement of fluid toward

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

the front of the tongue for ejection, while lateral tongue protrusions represent the gathering of fluid in the back of the mouth for eventual swallowing (this is true regardless of whether delivery is via IOC or licking, although the processes are non-identical and decisions are almost certainly made faster in a self-administration context, see [Samuelsen et al. \(2012\)](#)); while swallowing technically represents the end of the decision process, it provides a poor measure of when that decision is actually made (particularly for egestion/rejection decisions), for the simple reason that fluid must be gathered in the back or front of the mouth prior to swallowing or expelling.

Voltage neuromuscular signals from a bipolar electrode embedded in the anterior digastric were recorded during sampling of IOC-delivered tastes. These signals were passed through a differential amplifier (Grass Technologies: P55) and sampled at 1000 kHz, filtered offline using a two-pole Butterworth band-pass filter (300 to 500 Hz; see [Travers and Norgren \(1986\)](#)), and rectified. The EMG signatures of gapes were extracted from this signal using a two-step process: 1) We first identified all mouth movements produced in the first 2.5s following taste delivery, extracting the envelope of the filtered, rectified EMG signal using a 15 Hz low-pass filter; local peaks in the envelope were identified as mouth movements. The onset and offset of each movement were defined as the time points at which the magnitude of the envelope fell below the mean value of the envelope from the baseline period ([-1.5:0] s relative to taste delivery). 2) We then trained a quadratic classifier to identify each movement as either a “gape” or “non gape”; movement duration ( $\Delta t$  between movement onset and offset) and movement frequency (the reciprocal of the longest  $\Delta t$  between the peak of a given movement and the peak of each adjacent movement) differentiated gapes from other mouth movements, as described by ([Travers and Norgren \(1986\)](#)). These parameters were plugged into the quadratic expression, and if the resulting value was less than zero, the movement was labeled as a gape. The parameters of this automated sorting method were validated by comparison to blind coding of simultaneously-acquired video - a corpus of 1169 movements (437 of which were determined on the video to be gapes) - and found to be highly reliable. Once the gapes were identified, the time of the first gape in a bout was deemed to represent the onset of gaping - the latency to the decision, as reflected in behavior emission.

This single-trial measure of decision latency could then be directly compared to the same-trial latency to the onset of the palatability-related state, determined in an entirely independent analysis. This state was identified as described above (although the HMM was in this case calculated using 1ms bins, to maximize the temporal precision with which

transitions could be identified), and the latency of that state was calculated to be the time at which it became the most likely state. There were 2 sessions (out of 10 in which  $> 2$  neurons were simultaneously recorded) for which an interpretable HMM solution (one with clean sequences of states) could not be reached; these sessions were disregarded for purposes of this analysis. Even in the other sessions, there was a small subset of trials in which the palatability-related state became dominant twice; for this reason, we restricted our analysis such that each trial contributed only one ensemble state latency, by disregarding very early ( $< 250\text{ms}$ ) and very late ( $> 2000\text{ms}$ ) state onsets. There were also rare trials in which the most likely palatability-related state simply didn't appear; these trials were disregarded, since they didn't supply a transition time point.

The results of this brain/behavior comparison were evaluated in multiple ways. First, a simple Pearson correlation between transition and decision latencies allowed us to determine whether there was a statistically significant linear relationship between those latencies; the sign of the lag revealed whether neural transitions preceded decisions or vice-versa. We also directly compared the two distributions of decision latency, one in relation to stimulus presentation and one in relation to the neural state transition, using a  $\chi^2$  test. Since multi-rat distributions were neither normal nor similar in shape, we analyzed medians and interquartile ranges (the range of the middle 50% of scores) to ascertain whether neural state transitions predict decisions better than trial averaging: improvements in predictability should be reflected in larger concentrations of scores around a single value, a more sharply peaked distribution of latency.

## 1.4 Results

### 1.4.1 Consumption choice-related cortical activity appears only after stimulus-discriminative activity, emerging across 500ms in single-neuron analyses aligned to stimulus delivery

We initially recorded the responses of 7 primary cortical ensembles ( $10 \pm 2$  single neurons per ensemble) to deliveries of taste stimuli that varied widely in the readiness with which they are consumed (a directly choice-related property known as “palatability”, [Figure 1.1A](#)). In our (and prior) tests, sucrose and 0.1M NaCl are consumed avidly (they evoke prolonged

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

licking), whereas quinine and 1.0M NaCl are more likely to be avoided (they elicit little licking). Low NaCl concentrations were moderately pleasing and largely similar to water (not shown). The decision to consume more or less avidly (Figure 1.1A) dovetails reliably with other indicators of this naturalistic choice (see below) - sucrose on the tongue, for instance, causes swallowing-related reflex behaviors, whereas quinine on the tongue causes rejection-related behaviors (Grill and Norgren (1978a); see also below).

Figure 1.1B presents taste PSTHs (that is, single-neuron responses aligned to stimulus delivery and averaged across trials) for two cortical neurons. These responses were in good accord with previous reports (Katz et al. (2001), Piette et al. (2012b), Sadacca et al. (2012)): following initial periods of non-specific firing, firing rates became (at  $\sim 150$ -200ms post-delivery) distinct for different tastes; consistent with previous studies, these (and many other) neurons responded to multiple tastes (Katz et al. (2001), Smith-Swintosky et al. (1991), Stapleton et al. (2006), Yamamoto et al. (1984), but see also Chen et al. (2011)) in a manner that has been shown to reflect physical properties (their distinct chemical identities, concentrations, and degrees of mixture) of the stimuli (Katz et al. (2001), Maier and Katz (2013b), Sadacca et al. (2012), Yoshida and Katz (2011)).

Well after these representative responses became stimulus-specific, they can be observed to shift again (500-1100ms following stimulus delivery), at which point they came to reflect stimulus palatability - that is, the average behavioral choice (in the case of the neuron shown in ??B, the most vigorous response in this later period was to aversive stimuli; the opposite pattern was observed in the neuron shown in Figure 1.1B). To quantify this finding, we calculated moving-window correlations between the palatability of each stimulus and spike-rate responses to these same stimuli (Sadacca et al. (2012)). Figure 1.1C shows that, for both sets of responses displayed in Figure 1.1B, the correlation between palatability and firing rate was stably low and flat until approximately 0.5 seconds following stimulus delivery - well after the neurons had begun to fire distinctively to different tastes - at which time it appeared to ramp linearly upward, and to peak at or after 1.0 second.

This result was representative of the full dataset (Figure 1.2), despite the inclusion of unresponsive neurons in the analysis (which necessarily lowered the net magnitude of effect; black line). Palatability-related firing ramped upward between approximately 0.5 and 1.1 seconds after stimulus delivery, achieving significance at 0.83 seconds ( $p < 0.05$ , Tukey-Kramer test) - just prior to the average behavioral latency (Travers and Norgren (1986)): at peak, palatability accounted for 86% of the GC response variability; the firing rates of 41%

of the recorded single neurons ( $N=28$ ) were significantly linearly correlated with palatability during this response epoch (for 67% of these correlations with palatability, the correlation was negative; for more details, see [Sadacca et al. \(2012\)](#)).

Furthermore, palatability/consumption-related firing appears well after activity that is stimulus-specific (shown in grey). The emergence of palatability lags behind the emergence of pairwise discriminability by several hundred milliseconds, regardless of whether examination is keyed to: 1) the 30% or 70% points in the curves (horizontal lines); or 2) when the curves achieve significance. In fact, the discriminability curve reaches asymptote (at 100% correct according to a standard classification analysis) more than 500ms prior to the average time of choice behavior (whereas the palatability curve arrives at asymptote very close to the average time of behavior emission, see below). As these results predict, an analysis restricted to the 100-500 msec period (blue shading) reveals significantly above-chance levels of taste-specific information and essentially no linear correlation with palatability.

### 1.4.2 Single-trial emergence of consumption-related activity is more abrupt than that visible in PSTHs according to ensemble analysis

When we used HMM to re-analyze these neural data in terms of the responses of simultaneously recorded ensembles of neurons, the  $> 500$ -ms ramp of choice-related activity apparent in [Figure 1.1-Figure 1.2](#) was revealed to poorly reflect single-trial responses. The activity of cortical ensembles was instead well-characterized as reliable sequences of states (the solid lines in [Figure 1.3A](#) are state probabilities, which typically transitioned from low to high with great suddenness), in which each state was defined as a particular set of firing rates across the ensemble ([Figure 1.3A](#)).

The same sequence of states was observed in an average of 73% of the trials of any one stimulus (examples 1-3 in [Figure 1.3B](#)), and in 88% of the trials, the sequence was identical for at least 3 of the first 4 state transitions. This means that there was typically one particular state that characterized most single-trial responses in the  $> 1$  sec period (this single state was dominant in 78% of trials), the time at which palatability-related firing reached asymptote ([Figure 1.2](#)). But while the order of states was reliable, the timing of state-to-state transition latencies varied widely from trial to trial, such that the state that dominated the 1-2 sec period appeared at different latencies on different trials ([Figure 1.3C](#)). When calculated for

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

all trials of taste delivery, the median latency of state appearance was 0.8 seconds ( $\pm 20$ ms, SEM; [Figure 1.3D](#)), a good match for the time-point at which the correlation with behavioral choice became significant in across-trial and across-neuron averages ([Figure 1.1-Figure 1.2](#)). We therefore refer to the state dominant after this transition (i.e., the state that is most probable for the period between 0.8 and 2.0s in each model) as the putative “choice-related state”. The state that dominates responses before this time point appeared, on average, 0.1s after stimulus presentation, far too early to be reasonably associated with choice behavior (see below).

The trial-averaged probability of this late state (e.g., solid line in [Figure 1.3C](#)) rose slowly and linearly across the period between 0.5 and 1.0 seconds, mirroring the emergence of choice-related firing in [Figure 1.1-Figure 1.2](#), despite the fact that in single trials the state appeared suddenly. It is therefore reasonable to ask whether the slow emergence of choice-related firing itself poorly reflects single-trial ensemble activity. We tested this possibility by re-aligning trials to the onset of the transition into the choice-related state, hypothesizing that this relatively subtle realignment of the data would sharpen the transition into choice-related firing observed in PSTHs, significantly steepening the slope of the brain/behavior correlation. We calculated moving-window correlations between neural firing rates and palatability for each neuron in the re-aligned ensemble (just as had been done for data aligned to stimulus delivery, see [Figure 1.2](#)), and integrated these correlations across ensembles as a population measure of the choice-relatedness of neural activity.

After re-alignment, the correlation with choice did emerge more suddenly than the same data aligned to stimulus onset ([Figure 1.4A](#)). We quantified this finding by fitting sigmoidal curves to both stimulus-aligned and transition-aligned data, and found that the maximum transition duration (defined as  $\frac{1}{\beta}$ , a parameter that provides an unbiased estimate of the speed of transition) of the transition from low to high correlation was significantly shorter for transition-aligned data than stimulus aligned data ( $Z_{PSTH}=10.6$ ,  $p_{PSTH}>0.001$ , [Figure 1.4B](#)). The increase in suddenness of the transition was more than 3-fold (compare the black and red bars in [Figure 1.4B](#)) - that is, most of the apparent slow accumulation of palatability-related activity observed in [Figure 1.2](#) vanished without across-trial averaging.

We performed the same realignment using the two earlier state transitions (i.e., into the state just prior to the identified late-state, and into the even earlier state that reflects the initial response transient) and repeated the analysis; data aligned to the late-state transitioned more quickly than both controls ( $Z_{EARLY}=6.4$ ,  $p_{EARLY}>0.001$ ,  $Z_{PRE-LATE}=3.5$ ,

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

$p_{PRE-LATE} > 0.001$ , Figure 1.4B), neither of which sharpened the transition significantly (i.e., compared to non-realigned).

Of course, just as across-trial averaging may smooth data into gradual ramps, so HMM may identify spuriously sharp jumps in noisy data that in reality contain no coherent transitions. To evaluate the impact that this possible confound might have had on our conclusions, we produced and analyzed two control datasets: the first of these was a neuron-by-neuron reshuffling of the real data trials (which thus preserved any irregularities in spiking that might have contributed to the performance of the HMM, but disrupted trial-to-trial coherence amongst neuron pairs, Jones et al. (2007); see Figure 1.5A, red traces); the second was a direct simulation of the cortical dataset, derived from the neurons’ empirically observed PSTHs (with randomly chosen spike times); two such simulated neurons are shown in Figure 1.5A (grey traces). For each of these simulations, single-neuron spike rates closely resembled those observed in the experimental data, lacking (by design) only trial-specific dynamics.

We produced 200 versions of each control dataset, and subjected them to the same analyses brought to bear on the real data (i. e., those related to Figure 1.4). The results of these analyses are shown in Figure 1.5B-C: for both the shuffled (Figure 1.5B top) and simulated (Figure 1.5B bottom) data, the emergence of significant correlations with palatability was only slightly increased by realignment to HMM late-state onset. When fit with sigmoid functions using the same procedure brought to bear on the data in Figure 1.4A, the onset of palatability-related firing in data aligned to state transition is sharper for the real cortical ensembles than for either the PSTH-based simulation and trial-shuffled controls, an appearance confirmed by statistical analysis of suddenness (Figure 1.5C,  $p_{\text{simulated}} < 0.05$ ,  $p_{\text{shuffled}} < 0.05$  one-tailed test, 7/200 and 4/200 datasets respectively); control datasets almost never transitioned as suddenly as the real transition-aligned data. As a secondary check on the reasonableness of our fitting of suddenness, we also calculated the time required to pass from 30% to 70% of the total change in palatability. This analysis yielded quantitatively similar results ( $p_{\text{simulated}} < 0.05$ ,  $p_{\text{shuffled}} < 0.05$ , one-tailed test, 0/200 and 1/200 datasets respectively), further supporting our hypothesis that cortical neural ensembles transition suddenly into a coherent “palatability-rich” state. The true speed of this transition is obscured (i. e., slowed by a factor of 2.5) by: 1) across-trial averaging of the data; and 2) the concomitant treatment of single neurons as independent units of analysis.

### 1.4.3 The onset of palatability-related cortical firing coding is more like a step function than a “fast ramp” of firing rates.

The above results suggest that PSTH-based analyses mis-characterize the timescale of the dynamics of ingestive decisions, and that the entirety of the transition into choice-related firing is contained within, at most, a 150-200ms period - a short enough timescale to suggest that it does not represent an integrative ramp (Miller and Katz (2010), Shadlen and Newsome (2001)). It is also possible, however, that the brief period of transition reflects a genuine, albeit swift, ramping of firing rates, such as has been suggested to occur when an animal is making very easy perceptual decisions (Hanes and Schall (1996), Ratcliff et al. (2003), Uchida and Mainen (2003)), and that might be expected during similar “identification judgments” of tastes (Perez et al. (2013), Weiss and Di Lorenzo (2012)). In what follows, we contrast these possibilities.

While consumption decisions made in this context surely qualify as “easy” (see below), there are several aspects of the data already described that differ radically from that observed during the typical “fast ramps” underlying “easy decisions”. They are far too long in latency, (for instance, inappropriately variable in onset latency), and preceded by far-too lengthy epochs of taste-specific but decision-independent neural activity (see also Discussion). It is therefore reasonable to hypothesize that the underpinnings of consumption decisions observed in GC may be more ‘step-like’ than ‘ramp-like’. To more rigorously test this hypothesis, however, it is necessary to simulate data using stepping and ramping models, and then to directly, statistically compare the distribution of differences between the responses of real cortical ensembles and each simulation. We therefore produced 100 sets of simulated ensemble data reflecting the output of the standard ramping model used to describe perceptual decision-making (the drift-diffusion, or DDM, model; Figure 1.6A), in which the transition into palatability-related spiking was determined relative to a decision variable that performed a biased random walk toward a threshold. We took great pains to choose parameters that optimally matched the ramping of palatability and changes in firing rate in the real data, such that our analysis was, if anything, biased to maximize the performance of the DDM (see Methods for further details). In parallel, we produced another 100 sets of simulated ensemble data in which the determinant of that transition was an instantaneous step function, the time of which was chosen to be identical to those observed in the real data (the optimally-steep, or OPTI, model, Figure 1.6A). We then ran precisely the same analyses

on both of these sets of model-based simulations that had already been performed on the real data, and asked whether the increase in the suddenness with which palatability-related activity emerged in our real data better reflected a ramp or a step function.

The results of these analyses are clear from examination of [Figure 1.6B-C](#). First, HMM and subsequent realignment increased the suddenness of palatability-related firing for data synthesized from each model, as expected ([Figure 1.6B](#)). Second, and again as expected, the estimation and binning procedures that are an intrinsic part of HMM and correlation analysis necessarily introduced an apparent brief ramp of correlation even in OPTI simulations ([Figure 1.6B](#)), despite the fact that firing rates actually changed instantaneously. Third, transitions into palatability-related firing for realigned OPTI simulations were on average no steeper than that of real cortical ensembles - the distribution of differences between the real data and OPTI simulations was centered on zero (blue bars, [Figure 1.6C](#)) - whereas the distribution of differences between the real data and DDM simulations was shifted to the left on the x-axis of [Figure 1.6C](#) (reflecting the fact that most DDM simulations transitioned less suddenly). We again calculated the time required to pass from 30% to 70% of the total change in palatability as a validity check of this suddenness result, and again found quantitatively similar results.

Our hypothesis is that the two models should differ; therefore, we performed several tests directly comparing the distributions in [Figure 1.6C](#). First, we calculated the likelihood ratio ( $\Lambda$ ) of the probabilities of either simulation having a median transition duration equivalent to that of the real data (two-tailed Wilcoxon sign-rank). The difference between these likelihoods was significant ( $\Lambda_{OPTI-DDM} = 41.8$ ,  $p < 0.001$ ). Similarly, the two distributions of differences between the model data and the cortical ensembles were statistically different according to a two-sample Kolmogorov-Smirnov test ( $D = 0.33$ ,  $p > 0.001$ ). Despite the impossibility of instantaneous change in a real system (dynamical or otherwise, see Discussion), the step function model produces data more like our experimental data than does the integrator model. That is, data collected as rats prepare to consume or expel fluid appear closer to an instantaneous step function than a ramp.

#### 1.4.4 The sudden onset of palatability-related firing robustly predicts the latency of choice-related behavior in single trials.

The above data and analyses suggest that coherent neural ensemble activity in GC related to the making of naturalistic consumption decisions appears suddenly in single trials, with different latencies on different trials, not as an integrative ramp. Lacking from the above analysis, however, are the actual decisions that were made in the individual trials themselves, as the “palatability” measurement used so far is an average of decisions made across trials (which for the above was compiled from a separate session), and thus offers no information about the actual timing of an actual real-time choice. To further test our hypothesis that decision-making in GC is well described as occurring in a sudden transition (a “moment of insight”) with trial-specific timing, we therefore asked whether the latency of that transition predicts the making of the decision itself in that trial.

To answer this question, we collected an additional dataset (11 sessions from 9 rats) in which we simultaneously recorded the neural activity of GC ensembles and electromyographic (EMG) activity of a jaw-movement muscle (the anterior digastric). This muscle is activated to produce consumption/rejection behaviors, and thus allows us to directly correlate the occurrence of the transition into palatability-related firing and the occurrence of the behavioral decision itself. For this report we focus on the choice to reject (because it is easiest to reliably identify the occurrence of rhythmic rejection behavior, termed gapes, see [Travers and Norgren \(1986\)](#)).

If the hypothesis that neural state transitions into choice-related activity drive consumption decisions is to be supported, then the large trial-to-trial variability of these neural ensemble transitions should reliably predict the (similarly large) trial-to-trial variability of choice behavior latency. [Figure 1.7A](#) shows four consecutive quinine trials (i.e., separated only by trials of different tastes) from one session. The left edge represents the time at which the stimulus was presented, the dashed red line represents the likelihood of the palatability-related state described above, and the short vertical hash marks show individual gapes within a rhythmic bout. Note that the onset time of the choice-related state appears well linked to the time at which the decision is expressed in behavior (i.e., the appearance of the first gape), with earlier transitions to the palatability-related state occurring in the same trials that showed faster reaction times. For these four representative trials, the latency to decision could be as little as 0.5 sec after stimulus presentation, or as long as 1.2 (a 0.7-sec spread);

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

the transition into the putative choice-related state tracked this variability and predicted the behavior.

Figure 1.7B summarizes all data for the session from which these trials were culled, plotting the latency of neural state transition (x-axis) and latency of behavior (y-axis) for all 30 quinine presentations. The heavy diagonal dashed line bisecting the graph shows where trials would fall if those latencies were simultaneous. The decision was in most cases (88% of the trials) expressed only after the neural state transition occurred; in 79% (n=22) of those trials, the decision was expressed in behavior in a narrow band of time following the state transition (note the off-unity diagonal line). This means that, for this session, knowing the time of the transition made it possible to predict the time of response onset in individual trials, despite the fact that on some trials the decision was made in less than 500ms and in other trials in more than 1500ms - 59% of the trial-to-trial variability in decision time could be accounted for simply by predicting that behavior would occur 293ms after the transition was detected (i.e., the mean of the behavior - transition difference distribution).

This session, while of particularly high quality, represented the entire dataset (8 sessions—one session was removed for only having 2 single neurons, and 2 were removed because HMM failed to converge on interpretable solutions) well. Across all trials in which the operative state change occurred between 250 and 1600ms of taste presentation, the brain-behavior correlation was 0.58 ( $p < 0.001$ ), with a mean lag from ensemble transition to decision of 298ms. Moreover, the latency-from-transition data was more sharply peaked than the stimulus-aligned data ( $\chi^2 = 108.8$ ,  $p < 0.0001$ ), reflecting the fact that decision-related behavior is better predicted from ensemble transition than from stimulus presentation time (i.e., the average latency poorly characterizes taste decision-making).

But the distributions of decision latencies, transition latencies, and transition-to-decision latencies combined across rats were neither normal nor similar, and thus means and correlations provided only poor characterizations of the data. We therefore directly compared the distributions, quantifying the variability of predictions in terms of the medians and inter-quartile ranges (Figure 1.7C; the main panel shows the data normalized for the fact that the lag distribution necessarily has a larger range than the distribution of latencies related to stimulus presentation, and the inset shows the raw data). This analysis specifically revealed that behavior was much better predicted by state transitions than stimulus onset: 50% of the distribution (i.e., the inter-quartile range) of the lag between state transition times and behavioral latency was contained in an interval around the median that was less than half

the size of the interquartile range of the same distribution calculated using shuffled transition times ( $p < 0.03$ ).

Finally, additional investigation confirmed the significant relationship between the onset of decision-related behavior and the attainment of the specific cortical state under investigation. In over 80% of individual trials across the entire 11-session dataset (without any dropping of sessions), decisions were made (i.e., the first gape was produced) following the transition into a single state - a result that retained significance ( $p < 10^{-10}$ , Binomial Test,  $N=153$ ) when differences in state dwell times were controlled for. Furthermore, while gapes occurred in bouts that could last a full second, across the entire dataset some 75% of these gapes occurred in the same state ( $p < 10^{-16}$ , Binomial Test,  $N=570$ ).

Taken together, these results suggest that our single-trial ensemble coding measure of transitions into palatability-related firing does indeed reflect the making of consumption decisions and so the suddenness of such changes in neural activity is likely a reasonable measure of the suddenness with which such choices are made.

## 1.5 Discussion

In naturalistic situations, animals make many “easy” decisions in seeming moments of “insight” (Kounios et al. (2006), Kounios et al. (2008), Sternberg and Davidson (1995)) that lag far behind stimulus identification; such slow decision processes tend to be less about what the stimulus is than they are about how the animal currently plans to respond about the stimulus - emotional responses that are easily modified by context and experience - and often follow external stimulus presentation only with highly variable latencies. It is entirely possible that the latency of these decisions represent the time point at which a ramping neural function reaches some threshold (see below), but at least with regard to the consumption decisions studied here, they are reflected by coherent changes in the firing rates of ensembles of sensory cortical neurons that occur suddenly and at different latencies in different trials.

These cortical transitions are not well-described as “fast ramps”. They occur far later (three times the average latency, in fact) than similarly “easy” perceptual decisions described in the primate literature (Heekeren et al. (2004), de Lafuente and Romo (2006), Philiastides and Sajda (2006)). Long before these transitions into (linearly) decision-predictive firing occur, neural responses already indicate the identity of the stimulus - that is, the accu-

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

mulation of stimulus-related firing reaches asymptote before decision-related firing appears (Figure 1.2; see also Katz et al. (2001), Sadacca et al. (2012)). In contrast, primate perceptual decision-making ramps represent precisely these accumulations of stimulus-related firing, and are ubiquitously preceded only by information-poor transients. In the responses observed here, such transients precede even the earlier taste-specific states, and only the third “epoch” linearly predicts upcoming behavior (Katz et al. (2001), Sadacca et al. (2012), Moran and Katz (2014b)).

More to the point, cortical ensembles respond to gustatory stimuli in a manner that is more similar to simulations in which decision-related firing appears in an instantaneous transition than they are to ramping data simulated from either PSTHs or an integrator model; in fact, the emergence of choice-related firing is indistinguishable from a step function, and the timing of that transition predicts when the decision is reflected in behavior. The gradual ramps that are apparent in across-trial averages of GC responses are in fact poor reflections of single-trial ensemble activity, which “hops” into a decision-related state at different latencies on different trials; a recent publication has reached similar conclusions regarding decision-related activity in the classic primate motion detection task (Latimer et al. (2015)).

Nonetheless, it must be noted that our results are perfectly compatible with the possibility that genuine ramps occur elsewhere in the taste neuroaxis. Indeed, it is possible that integration occurs in a “downstream” region, and that the result of that integration - the “motor plan” - reaches sensory cortex (in which firing then reflects the crossing of the decision threshold) via feedback. Although our work on other taste-responsive regions (central amygdala, basolateral amygdala, and lateral hypothalamus) has thus far failed to turn up any obvious evidence for such a site of integration (Fontanini et al. (2009), Sadacca et al. (2012), Li et al. (2013)), it is likely that a decision requiring 500-1000ms of processing time brings into play a relatively distributed circuit. This assertion is consistent with the extant literature, which makes it clear that the most basic ingestive/egestive responses can be supported by the rodent brainstem alone (Grill and Norgren (1978b)), but that in the intact rat the forebrain, including GC, is a part of palatability-related decision-making circuitry (Berridge and Valenstein (1991), Kiefer and Orr (1992), Moraga-Amaro et al. (2014), Sasamoto et al. (1990), Schafe and Bernstein (1998), Shammah-Lagnado et al. (1992), Stehberg et al. (2011), Travers et al. (1997), Zhang and Sasamoto (1990)).

In fact, it is clear that feedback from basolateral amygdala (BLA) is vital for the cortical pro-

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

cess described here: BLA responses become palatability-related early; while these responses are far too early (800 ms prior to behavior, [Fontanini et al. \(2009\)](#)) to be reasonably thought of as reflecting the completion of a decision-making process, they are nonetheless necessary for this process, in that inactivation of BLA eliminates most decision-related activity in GC ([Piette et al. \(2012b\)](#)). Thus, while we have no evidence that the GC responses described here represent “efference copy” per se, they undoubtedly make vital use of “top-down” feedback. In that regard, our results are analogous to those reported by [Romo et al. \(2002\)](#), who observed responses in secondary somatosensory cortex that, across 0.5s of post-stimulus time, transitioned from being sensory-to response-related; these responses, too, appear not to reflect obvious efference from motor planning regions, but were conjectured to require feedback from areas involved in executive processing.

It is also worth considering the possibility that consumption decisions are better thought of under the aegis of purely non-integrating models - sometimes called “hopping” or “jumping” models. Such models have been used to explain the neural underpinnings of many other “post-perceptual” choice phenomena ([Deco et al. \(2007\)](#), [Insabato et al. \(2010\)](#), [Kemere et al. \(2008\)](#), [Rolls et al. \(2010\)](#)), as well as bistable perception ([Deco et al. \(2007\)](#), [Moreno-Bote et al. \(2011\)](#)), “changes of mind” ([Bollimunta et al. \(2012\)](#)), and Bayesian inference through sampling ([Moreno-Bote et al. \(2007\)](#), [Moreno-Bote et al. \(2011\)](#)), and have been shown to outperform perfect integrator models under certain conditions ([Miller and Katz \(2010\)](#)). Evidence suggesting that we might be recording from a network of this sort is found in the fact that the latency of firing rate transitions is much more variable than the transition width itself; this aspect of the data runs counter to models employing integration of evidence, but is a reliable feature of attractor-hopping models ([Miller and Katz \(2010\)](#)).

But even if we entertain the possibility that the consumption decisions examined here work according to an attractor-hopping model, we would not want to argue that integration doesn’t occur in most studied contexts. Perceptual decision-making, for instance, naturally lends itself to integration, as novel stimulus information is acquired with each passing moment ([Bowman et al. \(2012\)](#), [Halpern and Tapper \(1971\)](#), [Kepecs et al. \(2006\)](#), [Kiani et al. \(2008\)](#), [Shadlen and Newsome \(2001\)](#), [Stapleton et al. \(2006\)](#), [Uchida and Mainen \(2003\)](#)). Furthermore, the weeks to months of training required for perceptual decision-making tasks (training that is not necessary for consumption decisions) allows for the extensive fine-tuning of connection strengths that characterize ([Wang \(2002\)](#)) and optimize integrator circuits ([Bogacz et al. \(2006\)](#), [Miller and Katz \(2010\)](#), [Miller and Katz \(2013\)](#)).

## CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

Finally, even if an “attractor hopping” model best explains decision-making in the consumption context, it is simplistic to think that real neuronal networks actually transition with perfect instantaneity from one state to another. Conduction delays and other physical properties will slow transitions; furthermore, dynamics arise on a slower timescale than that of single neuron responses when an attractor-based circuit transitions between two states. Thus, while we observed no significant differences between the real data and “step-function simulations”, the real transitions probably take a certain finite amount of time, reflecting intrinsic circuit dynamics. Trial-specific multiple neuron analyses (Kass et al. (2005), Lawhern et al. (2010)), such as HMM (Abeles et al. (1995), Escola et al. (2011), Jones et al. (2007), Kemere et al. (2008), Ponce-Alvarez et al. (2012)) - even with its simplifications of uncorrelated firing within states and precise simultaneity of firing rate changes across whole ensembles - and models inspired by them to reproduce substantial across-trial variability in network activity, constitute a large advance in the direction of understanding naturalistic decisions as they are made “in the moment”. Our data contribute to a growing literature indicating that temporally distinct, discrete states of neural activity play an important role in cortical sensory coding and production of behavioral responses.

### 1.6 Figures

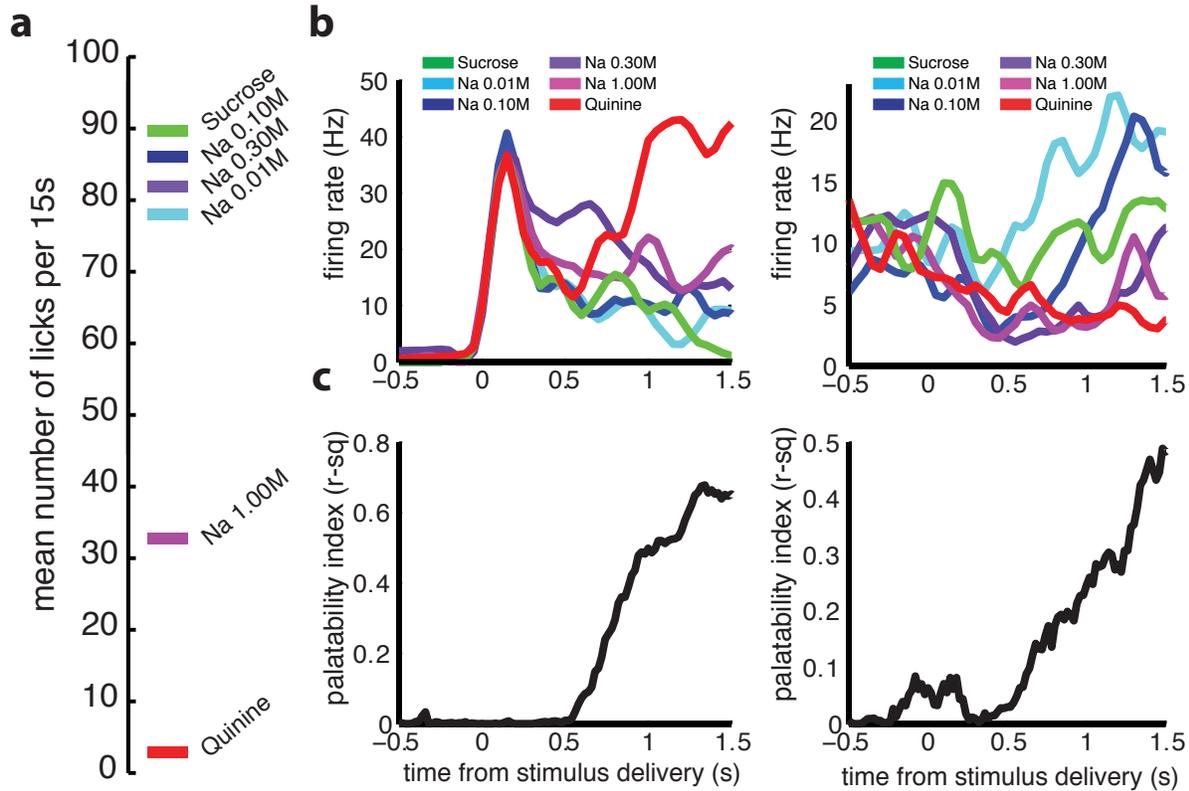


Figure 1.1: **Late-onset cortical activity reflects behavioral responses to the stimuli. A:** The amount that rats choose to consume (in terms of mean number of licks assessed in a brief access task, Breslin et al. (1993)) of each stimulus varies as a function of both the chemical identity and concentration of that stimulus. This schematic distribution serves as the basis for assessing the correlation with behavioral (consumption-related) choice. **B:** The trial-averaged, stimulus-aligned responses of two cortical neurons. Typical of such responses, initial activity was not stimulus-specific, selectivity emerging only after  $\sim 200$ ms (Katz et al. (2001), Piette et al. (2012b), Sadacca et al. (2012)). Still later, for each neuron, the response became obviously choice-related, with the strongest responses for this neuron to the most aversive stimuli and the weakest to the most palatable stimuli. **C:** The apparent emergence of choice-related responses in **B** is confirmed using a moving-window analysis of the linear correlation between the numbers of spikes in **B** and the behavior pattern in **A**.

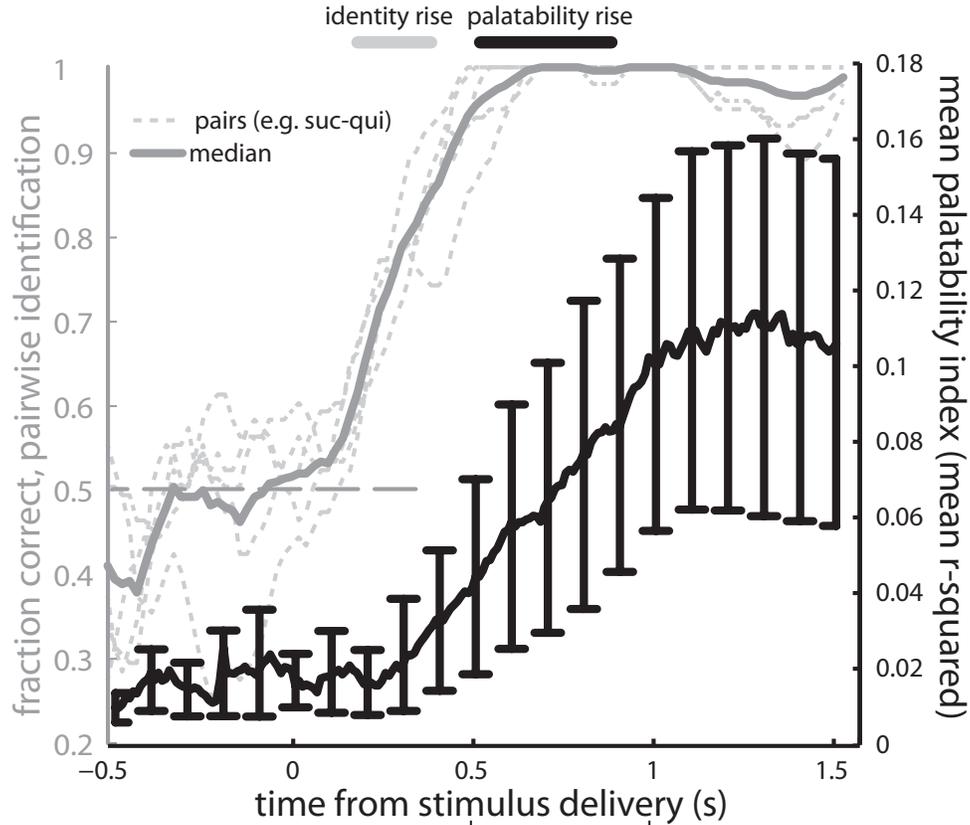


Figure 1.2: **Taste identification precedes taste evaluation.** The ramp of choice-related activity evident in single-neurons (e.g. Figure 1.1) is consistent across the entire neural sample (black line,  $\pm 2x$  standard error of the mean,  $n = 68$ ), though lower than for individual exemplars because of the inclusion of non-responsive neurons in the analysis. This palatability-related firing ramped upward, achieving significance above baseline at 0.83 seconds post-stimulus ( $p < 0.05$ , Tukey-Kramer test). The onset of evaluative coding occurred only long after tastes could be identified by ensemble activity—individual taste pairs (grey dashed line) were reliably discriminated above chance by 0.4 s following taste delivery, and 100% classification was achieved across all taste pairs (gray solid-line) 0.55 s post-stimulus. The duration of each curve—specifically, the time period spanning between 30% and 70% of the total ramp times for taste identification and palatability coding—are displayed (horizontal lines) above the histogram.

CHAPTER 1. DISCRETE BEHAVIORALLY RELEVANT CORTICAL STATES

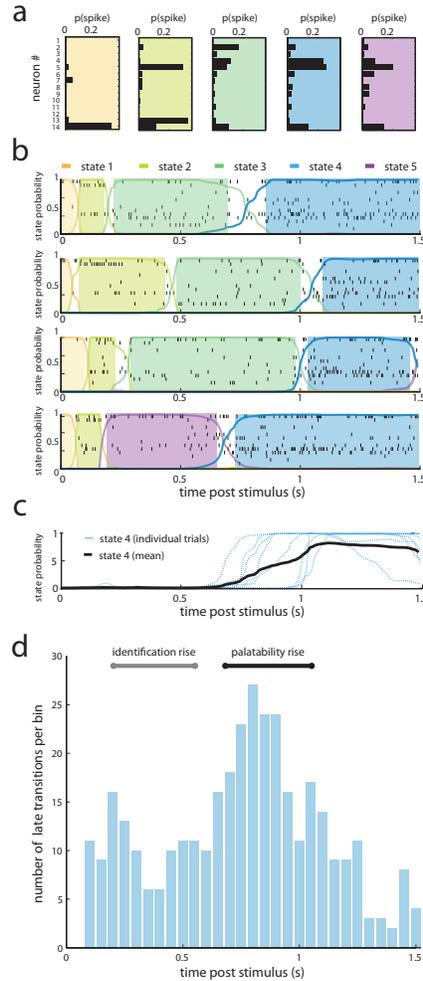


Figure 1.3: **Ensemble cortical responses form reliable sequences of states with coherent, trial-specific state-to-state transition times.** **A:** HMM-determined probability that a set of simultaneously-recorded cortical neurons achieves each firing-rate state (colored curves), plotted together with ensemble spiking activity (each vertical notch represents a spike), for four consecutive trials of one stimulus (sucrose). The same sequence was identified in most trials (here, the first 3 of 4 trials), but the times of state-to-state transitions varied from trial to trial (periods of high state-likelihood [ $>80\%$ ] are highlighted in color). **B:** Each neuron’s spiking probability (per 10ms) is plotted for each of the five states (color-coded to panel **A**) that occurred in the HMM solution derived for sucrose. **C:** The time courses of the HMM-derived “late-state” (i. e., the state dominant after 1 second) probability for all sucrose trials in one session, revealing both the reliability and suddenness of this state’s emergence (it progresses from 0 probability to 1.0 probability across a  $< 100$ ms period in almost every trial) and the considerable variability in the state’s onset latency from trial to trial. The solid black line shows the time-average probability across trials, which forms a gradual ramp with a time-course reminiscent of [Figure 1.1D](#). **D:** The distribution of identified ‘late’ state onsets across all modeled ensembles (42 models total).

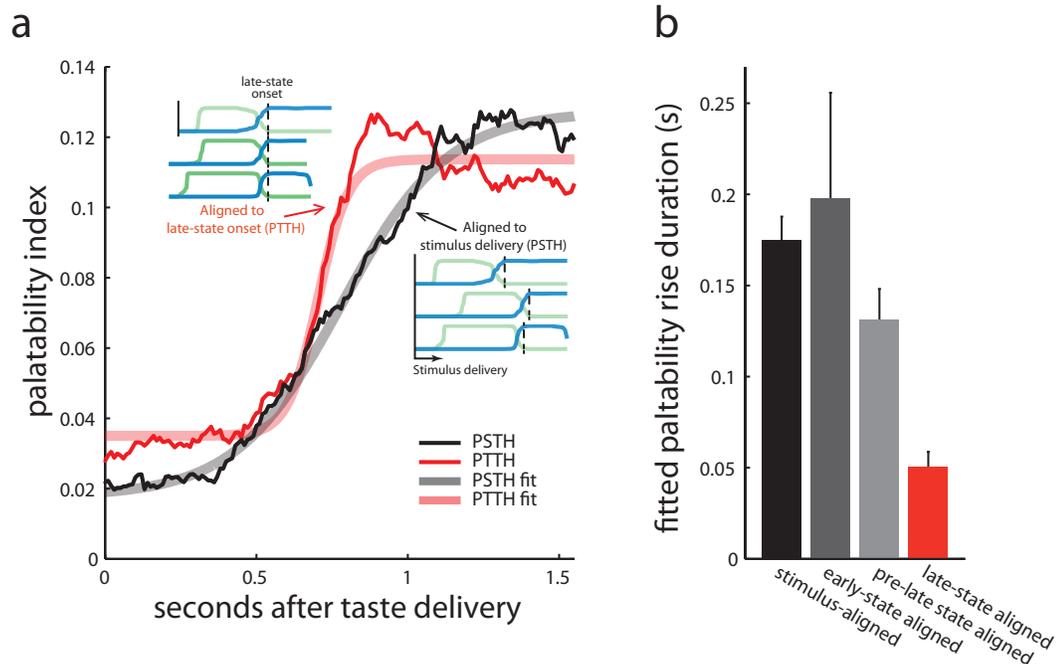


Figure 1.4: **Realignment of cortical ensemble data to the appropriate HM state onsets sharpens the emergence of choice-related firing.** **A:** The emergence of choice-related firing in ensemble activity is sharper following realignment of each trial’s spiking activity to the onset of the state identified as dominant at 1 second after stimulus delivery (PTTH, red dashed line) than when that activity is aligned to stimulus delivery (PSTH, black dashed line). Sigmoidal fits to each time series are overlain (solid lines), and provide estimates for the speed with which activity transitions to choice-relatedness. **B:** The values and 95% confidence intervals for the ‘duration’ of the transition from little choice-related firing to asymptotic choice-related firing—the time across which the slope of the fit sigmoid curve is “high” (parameter  $\frac{1}{\beta}$ ) - for PSTH and PTTH data; data realigned to either the state prior-to the identified late-state or realigned to the state dominant during taste identification show a less substantial drop in transition duration than does data realigned to the state dominant during palatability processing.

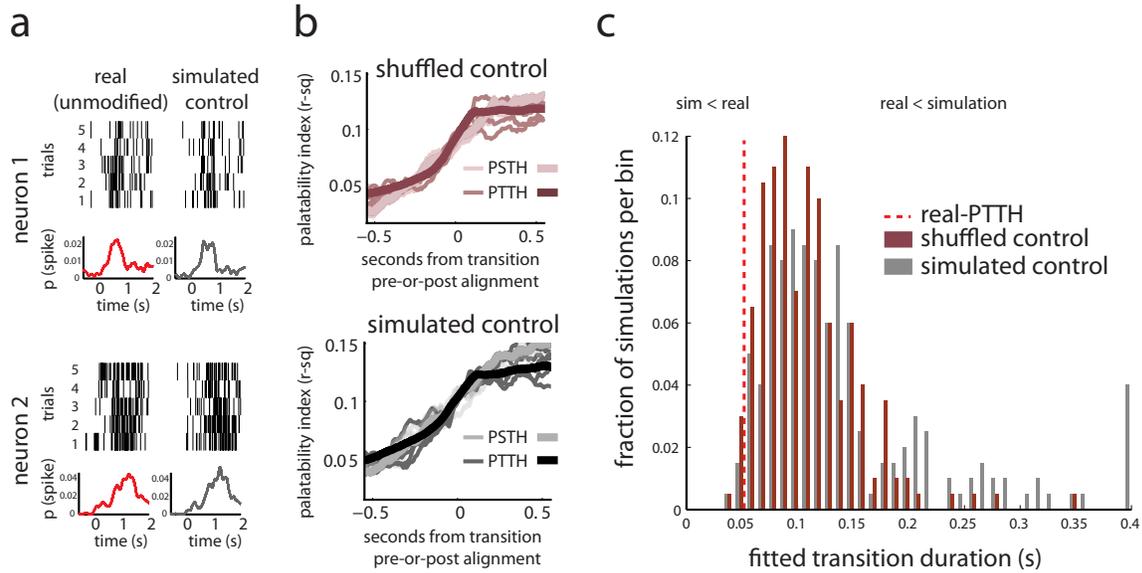


Figure 1.5: **The sharpening of choice-related firing by alignment to HM states is significantly greater than expected by chance.** **A:** Single-trial raster plots (above) and peri-stimulus spiking probabilities (below) for two example cortical single-neuron taste responses (in this case, to quinine, red trace) as compared with a simulation of spiking generated from the mean response of that neuron (grey trace). **B:** To estimate if the sharpening shown in Figure 1.4 was merely an effect of HMM realignment, irrespective of genuine rapid transitions in the neural data, trials from the experimental data were reshuffled (PSTH-shuffled, pink traces), and were modeled using HMM (PTTH-shuffled, maroon traces). Realignment produced a modest sharpening in this control data. For a second control dataset, (as plotted in A) spike-trains were simulated from the PSTHs of real neurons, maintaining the average activity of the ensemble, minus moment-to-moment correlations among neurons (PSTH-simulation, light gray traces). These simulated data were also modeled using HMM (PTTH-simulation, dark gray traces). Again, realignment caused a modest sharpening of choice-related activity. **C:** The results of sharpening 100 simulated datasets and 100 shuffled datasets are compared to the effect of sharpening the experimental data. Only 4 of the 200 simulations, and 7 of the 200 shuffled datasets transitioned into choice-related firing as suddenly as the real data, confirming that the effect shown in Figure 1.3A-B was not merely an effect of HMM-cued realignment ( $p_{\text{simulated}} < 0.05$ ,  $p_{\text{shuffled}} < 0.05$ ; one-tailed binomial test).

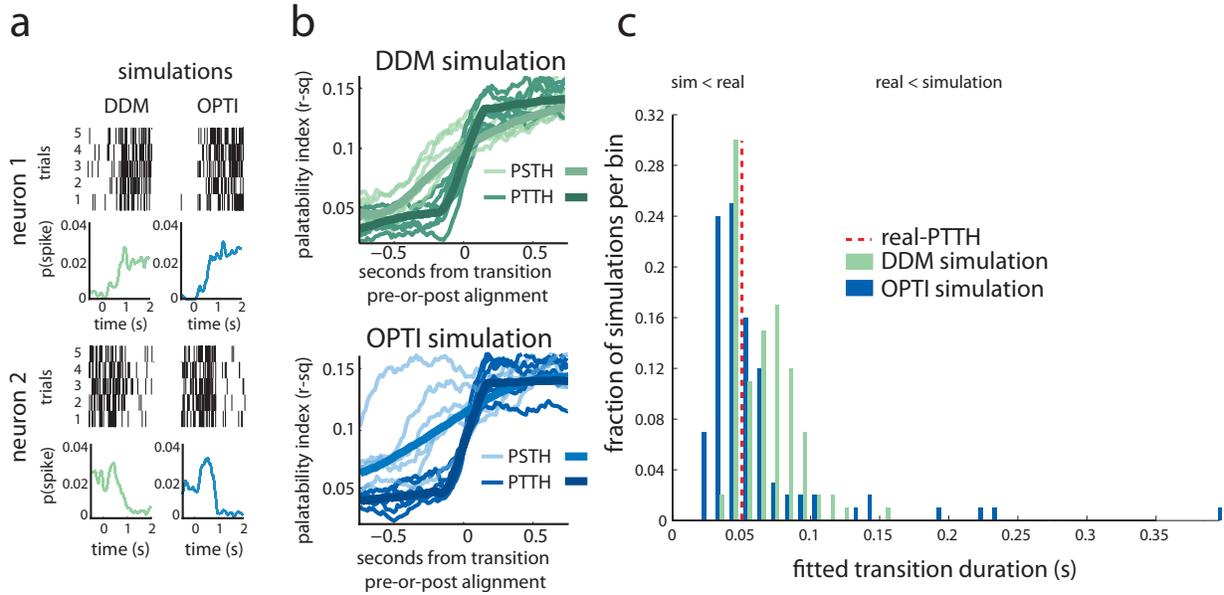
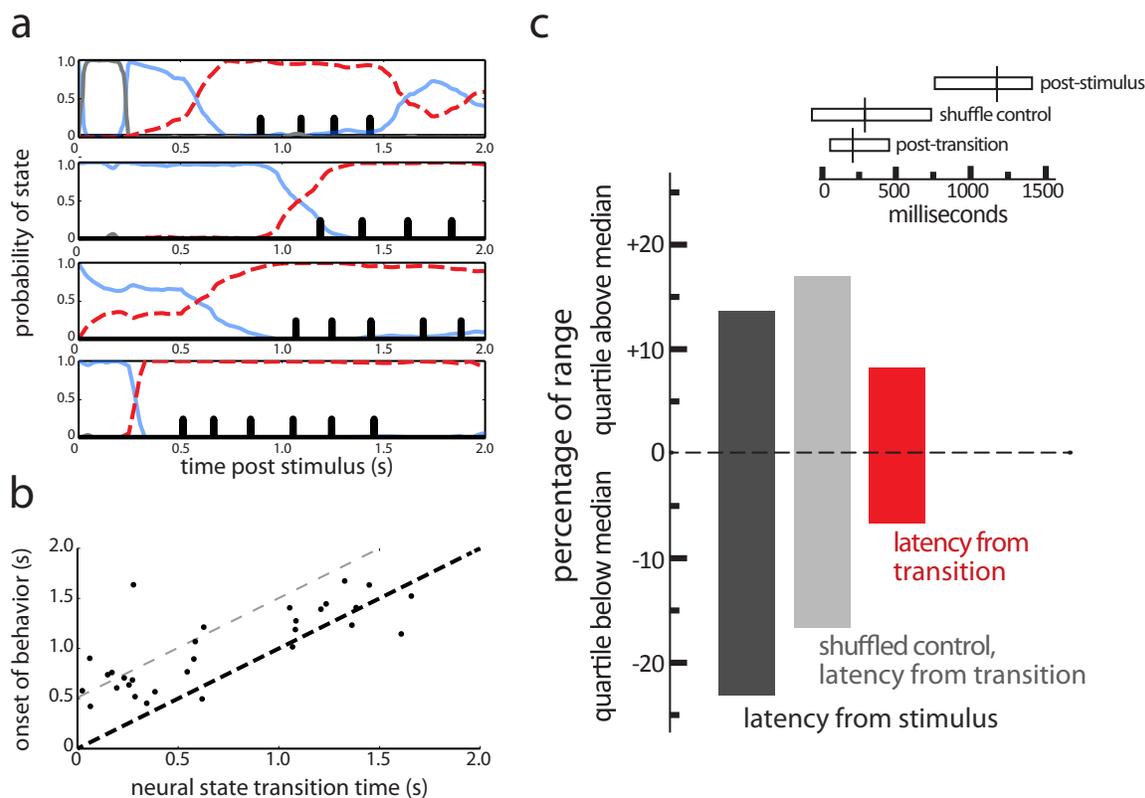


Figure 1.6: **The emergence of choice-related firing in cortical ensembles resembles instantaneous state transitions.** **A:** Single-trial raster plots (above) and peri-stimulus spiking probabilities (below) for two example cortical single-neuron taste responses (to quinine). Shown are the corresponding simulated neural spike-trains and mean responses generated from a ramping drift-diffusion model (DDM, green traces) and from an instantaneous rate-state transition model (OPTI, blue traces) for these two neurons. **B:** The simulated data were modeled in the same manner as the real ensembles and PSTH-derived simulations (Figure 1.3), and choice related activity was calculated pre-and post-HMM state alignment. Emergence of choice-related firing in both DDM- and OPTI-derived simulations (light shading, stimulus aligned; dark shading, state-aligned) substantially benefited from realignment. **C:** When directly compared to the suddenness of choice-related activity onset in real cortical ensembles, the instantaneously-transitioning OPTI data (blue bars) both transition significantly more rapidly than this DDM (green bars) simulation ( $p < 0.01$ , K-S test) and are a better match to the real data.



**Figure 1.7: The onset of the palatability-related state predicts consumption decisions.** **A:** Four consecutive quinine trials (time post stimulus delivery on the x-axis), showing the HMM solution (state probability on the y-axis); the state represented by a dashed red line is the palatability-related state. Overlain on this presentation are the times at which the rat gaped (vertical hash marks) - the time of the first gape is the behavioral readout of the decision to reject the quinine stimulus. **B:** For every trial in the session from which the panels in **A** were culled, the latencies of the palatability-related state (x-axis) are plotted against the latencies of gaping onset (y-axis). The thick diagonal dashed line is unity: above this line, the ensemble transition preceded the making of the decision; below, the opposite is true. In the vast majority of trials, the ensemble transition preceded the behavioral choice with a latency of less than 500ms (thinner diagonal dashed line). **C:** Across 8 sessions (see Methods), the mid-50% of times to first gape were closer to the median, when timing was aligned to the HM transition (red) than to the stimulus onset (dark gray) or to trial-shuffled transitions (light gray). Inset: absolute time. Main figure: times normalized to entire distribution width.

# Chapter 2

## Dynamical structure of cortical taste responses revealed by precisely-timed optogenetic perturbation

### Co-author contributions

The pre-print of this study is available on bioArXiv as [Mukherjee et al. \(2019\)](#) and is currently under peer-review for journal publication.

The following people contributed to this study (and were co-authors on the published paper):

1. **Narendra Mukherjee**: Designed research, performed research, analyzed data, wrote paper, gathered funding.
2. Joseph Wachutka: Designed research, performed research.
3. Donald B. Katz: Designed research, wrote paper, gathered funding.

### 2.1 Abstract

The purpose of perception is driving action. During tasting, for instance, every stimulus must be either swallowed or rejected (the latter via a behavior known as “gaping”). Taste responses in the rodent primary gustatory cortex (GC) span this sensorimotor divide, progressing through a series of firing epochs that culminate in the emergence of action-related firing.

Population analyses reveal this emergence to be a sudden, coherent ensemble transition that, despite varying in latency between trials, precedes gaping onset by 0.2-0.3s. Here, we tested whether this transition drives gaping, delivering 0.5s GC perturbations at various time-points in tasting trials. Perturbations significantly delayed gaping, but only when they preceded the action-related transition - thus, the same perturbation might have an impact or not, depending on the transition latency in that particular trial. Our results suggest a distributed attractor network model of taste processing, and a dynamical role for cortex in driving motor behavior.

## 2.2 Introduction

One of the primary purposes of sensory processing is to drive action, such that the sources of sensory information can be either acquired or avoided (Prinz (1997), Wolpert and Kawato (1998), Wolpert and Ghahramani (2000)). To the extent that this is true, sensory and motor processing should be tightly coupled (Wolpert et al. (1995), Huston and Jayaraman (2011)). The gustatory system is an ideal model to study this proposed coupling, because animals necessarily respond to tastes with discriminative behaviors - specifically, they must decide to either swallow or expel the sensory stimulus in the mouth (Grill and Norgren (1978a), Katz and Sadacca (2011), Li et al. (2016)).

Sensory-motor coupling is visible in the temporal response patterns of rodent gustatory cortical (GC) neurons to taste stimulus administration. GC neurons respond to taste presentation with a sequence of firing-rate “epochs”, two of which are taste-specific: neural firing first carries information regarding the physio-chemical identity of the taste stimulus, and then correlates with palatability, a variable intimately linked with the animal’s decision to ingest or expel the taste (Katz et al. (2001), Fontanini and Katz (2006), Grossman et al. (2008), Piette et al. (2012b), Sadacca et al. (2012), Maffei et al. (2012), Jezzini et al. (2013); see also Crouzet et al. (2015)). Ensemble analyses further reveal that the transition between these two epochs happens suddenly and coherently within neural ensembles (Jones et al. (2007), Sadacca et al. (2016)). This ensemble transition to palatability coding, though highly variable in latency (between 0.5 and 1.5s post stimulus, depending on the trial), is a strong predictor of the onset of the animal’s consumption-related orofacial behavior (Sadacca et al. (2016)), even when the timing of this behavior is manipulated by learning (Moran and

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

Katz (2014b)) or cueing (Li et al. (2016)). That is, GC neural ensembles appear to “hop” from one attractor state to another during taste processing (Miller and Katz (2010), Miller (2016)), with the hop representing the reaching of a consumption decision - and (potentially) the emission of a motor signal to brainstem circuits that generate orofacial behavior.

A direct prediction of this temporally dynamic model of gustatory sensorimotor processing, and most specifically of the suggestion that the transition into the later firing-rate epoch represents the emission of a motor command, is that well-timed perturbations of GC activity should affect the time course of a rat’s taste-reactive ingestion-egestion behavior. This prediction recently received indirect support when it was shown that optogenetic inhibition of the entire GC taste response (Li et al. (2016)) modestly changes the probability of rejection behaviors in response to aversive tastes (“gapes”, Grill and Norgren (1978a), Li et al. (2016)).

However, such gross perturbations of gustatory processing are an inadequate test of this very specific prediction: for one thing, multi-second inactivations likely have secondary effects that confound interpretation, particularly regarding an outcome variable (ability to gape) that is known to depend on an interconnected network of brain regions (including GC; see Smith and St John (1999), Riley and King (2013), Samuelsen and Fontanini (2016)); in addition, it is impossible to disambiguate any epoch- or moment-specific effects on consumption behavior using whole-response perturbations. A much more definitive test would involve using optogenetics to inhibit GC taste responses for short periods of time as awake rats process and respond to a range of tastes.

Here we report the results of this precise experiment, performed in awake, tasting rats. We recorded the activity of GC ensembles while simultaneously inhibiting the firing of these neurons using an optogenetic silencer (specifically, the proton-pump ArchT) for brief (0.5s) periods before, during or after the “hop” to the palatability- (i.e., decision-) related state. Our results provide strong support for the hypothesized importance of the transition time itself, and in addition suggest that important pre-transition taste processing is performed within GC. Furthermore, our data provide a glimpse into the attractor-like dynamics underlying the neural processing of taste, demonstrating that GC is one participatory node in a larger network with attractor dynamics: the fact that GC perturbations only delay the system settling into the decision-related “stable” state suggests that this stable state is a function of activity spread across multiple regions; in addition, the fact that post-decision perturbations have no impact suggests that behavioral control shifts to brainstem circuits once this stable state has been reached.

## 2.3 Materials and Methods

### 2.3.1 Experimental design

#### Subjects

Adult, female Long-Evans rats (n=5; 275-300g at time of virus injection; 300-350g at time of electrode implantation) served as subjects in our study (in our hands, female Long-Evans rats have proven more docile than males, but we have observed no sex differences in the basic cortical dynamics of taste responding). The rats were housed in individual cages in a temperature and humidity controlled environment under a 12:12h light:dark cycle. All rats were given *ad libitum* access to food and water before experiments started. Rats were weighed daily and observed to never drop below 80% of their pre-surgery weight. All experimental methods were in compliance with National Institutes of Health guidelines and were approved in advance by the Brandeis University Institutional Animal Care and Use Committee.

We also performed a set of control analyses on data taken from 10 adult, female Long-Evans rats, previously published in [Sadacca et al. \(2016\)](#) and [Li et al. \(2016\)](#).

#### Virus injections

We injected adeno-associated virus (AAV9) coding for ArchT and green fluorescent protein (AAV9-CAG-ArchT-GFP,  $2.5 \times 10^{11}$  particles per mL) into GC. This AAV serotype has been shown to effectively spread to and infect all cell types ([Aschauer et al. \(2013\)](#)) in regions including GC ([Maier et al. \(2015\)](#), [Li et al. \(2016\)](#)).

Rats were first anesthetized using a ketamine/xylazine mixture (1mL ketamine, 0.05 mL xylazine/kg body weight) delivered *via* an intra-peritoneal injection. Supplemental anesthetic injections were given as needed. The head was shaved, cleaned with an iodine solution and 70% ethanol, and positioned into the stereotax. We then excised the scalp and cleaned and leveled the top of the skull. Small craniotomies were drilled bilaterally over the location of GC (anteroposterior +1.4mm from bregma, mediolateral  $\pm 5$ mm from bregma; [Paxinos and Watson \(2007\)](#)), the meningeal tissues were gently excised, and virus was infused.

We lowered a glass micro-pipette (tip diameter: 10-20 $\mu$ m) filled with the infusate - virus particles suspended in a solution of phosphate-buffered saline (PBS) and Oregon Green 488 (Invitrogen) - into the centers of the craniotomies, and performed a sequence of 3 injections

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

bilaterally into GC: at 4.9, 4.7 and 4.5mm ventral to dura, virus was injected in discrete pulses (44 pulses/location, with 25nL per pulse, 7s between consecutive pulses = 1.1 $\mu$ L total volume injected per depth) controlled by a Nanoject III microinjector (Drummond Scientific). Following each unilateral set of injections, the micropipette remained in place for 5 min, after which it was smoothly removed over the course of 1 minute so that fluid would not spread back up the micro-pipette track. Craniotomies were then sealed with silicone (Kwik-Sil, WPI), the scalp was sutured, and the rat was given analgesic (meloxicam 0.04mg/kg), saline and antibiotic (Pro-Pen-G 150,000U/kg) injections. Similar antibiotic and analgesic injections were delivered 24 and 48 hours later.

Rats were allowed to recover for 4-6 weeks from this procedure, in order to ensure adequate infection and subsequent expression of optical channels (ArchT) and GFP.

### **Opto-trode, intra-oral cannula and EMG electrode implantation**

After recovery from virus infusion surgery, rats were again anesthetized, and implanted with bilateral GC opto-trode bundles. Each bundle consisted of either 30 or 32 recording microwires (0.0015inch formvar-coated nichrome wire; AM Systems) and 1 optical fiber (0.22 numerical aperture, 200 $\mu$ m core, inserted through a 2.5mm multimode stainless-steel ferrule; Thorlabs). The microwire bundle was glued to a custom-made electrode-interface board (San Francisco Circuits) and connected to a 32 channel Omnetics connector. In the case of the 30 microwire bundles, the final two pins were connected to 2 electromyography (EMG) electrodes (PFA-coated stainless steel wire; AM Systems) implanted into the digastric muscle under the jaw. Finally, the microwires and optical fiber were connected to a custom-built 3D printed microdrive that allowed the entire assembly to be moved ventrally after implantation. The microwire tips were located 0.5mm ventral to the tip of the optical fiber - this maximized the likelihood that the electrodes recorded the activity of neurons that were illuminated by the laser. For more information on the implanted apparatus and associated electronics, see [Katz et al. \(2001\)](#), [Sadacca et al. \(2016\)](#) and [Li et al. \(2016\)](#), as well as the [Katz Lab webpage](#).

Rats were anesthetized, after which we shaved and cleaned the scalp and situated the head in the stereotax. After excising the scalp and leveling the skull, we drilled 5 self-tapping screws into the skull for supporting and grounding the opto-trode bundles. The silicone seal was removed from the craniotomies, as were any tissues that had grown in since the prior surgery. We then slowly (over 5-10 minutes) lowered the opto-trode bundles to a depth of

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

4.3mm from the dura mater (0.2mm above the most dorsal location of virus injection). The ground wires were wound tightly around the skull screws and the bundles were cemented in place with dental acrylic. The optical fiber was looped so that the ferrule could be cemented away from the microdrive - this configuration reduced the stress on the microdrive when the animal was later plugged in to the experimental apparatus.

Once the opto-trode assembly was cemented in place, the rat was removed from the stereotax and implanted with a single (right-side) intra-oral cannula (IOC) for controlled delivery of tastants on the tongue. IOCs were made with thin polyethylene tubing and inserted in the space between the first maxillary molar and the lip, through the masseter muscle and inside the zygomatic arch, and out through the opening in the scalp (Phillips and Norgren (1970), Katz et al. (2001)) The IOC was topped with a plastic connector that could be attached to the taste delivery apparatus, and cemented in place with dental acrylic.

The EMG electrodes were channeled down the left side of the face (opposite from the IOC); after the overlying skin had been teased away from the belly of the digastric muscle, one end of each EMG electrode was tied to a suture needle, which was then inserted into the muscle, such that the electrode could be pulled into the desired position (for more details, see Loeb and Gans (1986); Travers and Norgren (1986); Dinardo and Travers (1994); Li et al. (2016)). The electrode wires were trimmed and held in place with vetbond tissue adhesive (3M) and the skin covering the anterior digastric was sutured back into place. Finally, a modified falcon tube was glued to the front of the headcap as a protective cap, and bacitracin ointment was applied all around the base of the headcap and over the wound under the jaw.

Rats were postoperatively injected with analgesic (Buprenorphine 0.05mg/kg), saline, and antibiotic (Pro-Pen-G 150,000U/kg). Similar antibiotic, saline and analgesic injections were delivered 24, 48 and 72 hours later, and bacitracin ointment was reapplied. The rats were handled every day and allowed to recover to 90% of their pre-surgery weight (at least 7 days after surgery) before being introduced to the experimental apparatus.

### Habituation

Following recovery from the opto-trode implantation surgery, we habituated rats to passive water deliveries for 3 days before beginning data collection. In these daily habituation sessions, we attached the rats to the electrophysiology acquisition system, laser patch cables and taste delivery apparatus, and infused 100 pulses of distilled water ( $\sim 40\mu\text{L}$  per pulse;

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

15s inter-pulse interval) into the animal’s oral cavity through the IOC. Starting with the second habituation day, we also placed rats on a mild water restriction schedule - 20mL of water (not including the 4mL delivered during habituation sessions themselves) per day. This water restriction schedule was maintained for the duration of the study (~7 days per animal).

Opto-trode bundles were driven deeper after each habituation session using the microdrive built into the assembly; by the end of the habituation period, the distance traveled was 0.2mm, such that the tips of the electrodes lay within the region of GC infected with the virus.

### Passive taste administration and laser stimulus delivery

We used 2 concentrations of palatable sucrose (30mM: Dilute Sucrose (Dil Suc), 300mM: Concentrated Sucrose (Conc Suc)) and of aversive quinine-HCl (0.1mM: Dilute Quinine-HCl (Dil Qui), 1mM: Concentrated Quinine-HCl (Conc Qui)) dissolved in distilled water as the stimuli in our experiments. Concentrated sucrose and quinine are rich in palatability-related valence and evoke strong orofacial responses; the dilute stimuli are of similar but far less extreme palatability – a fact that aided in the analysis of palatability-related neural firing (Li et al. (2016); see also below). The taste delivery apparatus consisted of gently pressurized tubes containing taste solutions; the tubes converged upon a manifold of finer polyamide tubes that could be inserted into (to 0.5 mm past the end of) the IOC, thus eliminating any chance of mixing. The manifold could be locked securely into the dental acrylic cap. The tastes were then delivered under slight nitrogen pressure - this taste delivery protocol has been consistently shown to ensure reliable tongue coverage at short latencies (Katz et al. (2001), Sadacca et al. (2016), Li et al. (2016)).

Data were collected during 2 types of optogenetic perturbation sessions: 1) sessions made up of “long” perturbation trials in which the laser was turned on for the period of 0-2.5s post taste delivery; and 2) sessions made up of “short” perturbation trials in which the laser was turned on for 0.5s at either 0.0, 0.7, or 1.4s post taste delivery. One experimental session was run per day. Some rats received only the latter (short-perturbation) session; for those that received both, we counterbalanced session type, such that a rat that experienced 2.5s perturbations in one session got 0.5s perturbations the following day, and *vice versa* (see below).

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

Sessions with 2.5s perturbations consisted of 8 sets of trials (2 sets per taste - one with the lasers on and one with no laser). Each set included 15 trials, for a total of 120 trials per session. Similarly, sessions with 0.5s perturbations included 16 sets of trials (4 sets per taste - one with lasers on from 0.0-0.5s, one with lasers on from 0.7-1.2s, one with lasers on from 1.4-1.9s, and one with no lasers). To keep the total number of trials per session from ballooning (a basic concern in taste research is the awake animal’s finite appetite), each set included only 8 trials (total, 128 trials per session). Again, we moved the opto-trode bundle 0.075mm ventrally (deeper into GC) prior to each session, to ensure that we obtained fresh units in every session. Trials were delivered in pseudo-random order and each involved delivery of  $\sim 40\mu\text{L}$  of fluid through the IOC, for a total volume of 5mL per session.

We used a 532nm, DPSS laser (Laserglow Technologies), connected to the implanted ferrules using standard FC/PC patch cables (Thorlabs), for all optogenetic perturbations. Taste and laser delivery were controlled through a Raspberry Pi computer. The strength of the laser input was calibrated, prior to opto-trode implantation, to yield an illumination power of 40mW at the tip of the optical fiber. This output power perturbs all ArchT infected neurons in a  $1\text{mm}^3$  sphere below the tip of the fiber *in vivo* (Han et al. (2011), Yizhar et al. (2011b)) - a sphere that encompasses about 33% of GC in the caudal-rostral axis (Kosar et al. (1986), Maier et al. (2015), Li et al. (2016)). These parameters have previously been shown to reduce the activity of ArchT+ cortical neurons with minimal latency and damage (Maier et al. (2015), Li et al. (2016), Flores et al. (2018)).

### Acquisition of electrophysiological data

We collected 30k voltage samples per second from each implanted neural and EMG electrode, using a 32-channel analog-to-digital converter chip (RHD2132) from Intan Technologies. These chips are capable of recording voltage signals over a wide range of frequencies (0.1Hz-20kHz) and amplitudes (microvolts to millivolts), thereby enabling us to record neural and EMG signals through the same hardware system. The experimental chamber was ensconced in a Faraday cage that shielded recordings from external electrostatic and electromagnetic influences.

### Histology and evaluation of GFP expression

In preparation for histology, rats were deeply anesthetized with an overdose of the ketamine/xylazine mixture, after which DC current ( $7\mu\text{A}$  for 7s) was passed through selected microwires, marking the area below the electrode tips. We perfused the rats through the heart with 0.9% saline followed by 10% formalin and harvested the brain. The brain tissue was incubated in a fixing mixture of 30% sucrose and 10% formalin for 7 days before GC was sectioned into  $50\mu\text{m}$  coronal slices.

We rinsed the slices 3 times with 1X-PBS over 15 minutes and permeabilized them in a 0.3% Triton X-100+1% normal Donkey serum+1X-PBS blocking solution for 2 hours at room temperature. We replaced the blocking solution with primary antibody solution (1:500 anti-GFP-rabbit IgG; Life Technologies) for 12 hours at  $4^{\circ}\text{C}$ . After incubation with the primary antibody, the slices were rinsed with 1X-PBS 3 times over 15 minutes followed by incubation with the secondary antibody incubation of (1:200 Alexa Flour 488 donkey anti-rabbit IgG (H+L); Life Technologies) for 12 hours at  $4^{\circ}\text{C}$ . After a final set of rinses with 1X-PBS (3 times over 15 minutes), we mounted the slices on charged glass slides and cover-slipped them with Fluoromount Aqueous Mounting Medium. Slices were imaged with a Keyence fluorescence microscope to confirm successful virus infection and opto-trode location for each animal.

The spread of AAV in GC was evaluated *via* the expression of GFP, as has been done previously (Maier et al. (2015), Li et al. (2016), Flores et al. (2018)).

### 2.3.2 Data analysis

Most statistical analyses in this paper were performed using Bayesian methods implemented in the PyMC3 probabilistic programming package (Salvatier et al. (2016)). Although the far more common practice in the literature is to implement analyses similar to ours in a frequentist/maximum likelihood estimation (MLE) paradigm, the Bayesian approach offers several advantages. For one, Bayesian statistics provides a natural way to infer the entire joint posterior distribution of the model parameters in the light of the data at hand. This allows the Bayesian methodology to make robust inferences without being constrained by the sampling-related assumptions of parametric frequentist statistics or the lack of statistical power of non-parametric frequentist techniques. Relatedly, the flexibility of the Bayesian framework allows the construction of statistical models appropriate for the data-generating

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

process that can include non-standard (such as multi-modal) parameter distributions. Such models (of which we use several in this study) often cannot be accommodated by frequentist approaches at all, even if they are “true” descriptions of the underlying generative process. Finally, despite working with highly flexible models, Bayesian approaches provide the added advantage of using model priors to regularize parameter estimates - we use “weakly informative” priors in our analyses that are known to reduce the susceptibility of the inference process to noise by penalizing model flexibility (unless supported by the observed data).<sup>1</sup> We will describe the properties of each statistical model used in our analyses, and our specific prediction(s) for each such model, in the sub-sections below.

Recent advances in statistical computing have made it possible to circumvent the analytical challenges that have historically plagued the application of Bayesian techniques to many practical problems. In particular, new Markov Chain Monte Carlo (MCMC) techniques have been developed to facilitate arriving at an approximation to the posterior distribution of the model parameters by drawing samples from it. We performed inference in our Bayesian probabilistic models using the No-U-Turn-Sampler (NUTS; [Hoffman and Gelman \(2014\)](#)), a state-of-the-art, self-tuning Hamiltonian MCMC algorithm that efficiently draws samples from the posterior distribution described by the data at hand. The performance of the sampler can be evaluated by running several independent sampling chains - a properly tuned sampler that explores the parameter space in an unbiased manner and draws samples from the correct posterior distribution will result in all the chains “converging” to the same distribution. Statistically, this is evaluated by computing the Gelman-Rubin  $\hat{R}$  statistic ([Gelman et al. \(2011\)](#)) across all the sampling chains.  $\hat{R}$  close to 1 indicates that the sampling runs have converged and produced samples from the same posterior distribution (we allow values from 0.99 to 1.01). Each analysis finally reports the uncertainty for the inferred parameters as 95% credible intervals - essentially the interval that covers 95% of the probability mass under the posterior distribution of the parameters. Credible intervals inherently serve as significance tests in this setting - for instance, if the 95% credible interval for an estimated parameter does not overlap 0, we can conclude that this parameter is different from 0 at the 5% level of significance.

---

<sup>1</sup>For a detailed comparison of frequentist and Bayesian estimation in statistics and a discussion of weakly informative priors, please refer to [Gelman et al. \(2013\)](#) and [McElreath \(2015\)](#).

### Single unit isolation

We followed a semi-supervised spike sorting strategy: intra-cranial voltage data was filtered between 300-3000Hz, and a Gaussian Mixture Model (GMM) identified potential clusters which were refined manually. For more details on our spike sorting methods and its efficacy in isolating single units, please consult [Mukherjee et al. \(2017\)](#). Our spike sorting code is freely available at [blech\\_clust](#).

### Impact of optogenetics on neural firing

We built a hierarchical Poisson generalized linear model (GLM) for the spiking of a single neuron to evaluate the impact of optogenetic perturbations on firing. Hierarchical GLMs provide precise estimates of condition-specific model parameters, especially when they are expected to vary around condition-agnostic means. In our case, the model parameters are the mean firing rates for every taste and optogenetic condition, that are in turn composed of taste- and optogenetic-specific effects ("random effects") and means across tastes and optogenetic conditions ("fixed effects"). Coupled with the Poisson distribution's suitability for count (here spikes) data, this model can accurately estimate the change in neural firing induced by optogenetic perturbations.

For each neuron  $n$  in our dataset, we aggregated the spikes produced on trial  $i$  of taste  $T$  in optogenetic condition  $O$ . There were 4 levels for  $T$  corresponding to the tastes in our dataset: Dil Suc, Conc Suc, Dil Qui and Conc Qui. The number of levels for  $O$  depended on the type of optogenetic perturbation being delivered in the session: in the 2.5s perturbation sessions,  $O$  had two levels, corresponding to the laser off (control) and on trials respectively; the 0.5s perturbation sessions had 3 types of perturbation trials - starting at 0s, 0.7s or 1.4s after taste delivery - and therefore had 6 levels for  $O$  (a "laser off-laser on" pair for each of the 3 types of perturbations). Our model posits that the aggregate number of spikes  $S_{n,i,T,O}$  is Poisson-distributed with a mean ( $firing_{n,T,O}$ ) that depends on the taste ( $\mu_T$ ), optogenetic condition ( $\mu_O$ ) and an interaction between the taste and optogenetic condition ( $\mu_{T,O}$ ). As described above, owing to the hierarchical structure of the model, each of these effects is further composed of a fixed effect and a random effect. Using weakly informative Gaussian and Half-Cauchy priors for the mean and variance parameters respectively, our model formally says:

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

$$\begin{aligned}
 \text{Fixed effects: } & F_1, F_2, F_3 \sim \mathcal{N}(0, 10) \\
 \text{Variances: } & \sigma_1, \sigma_2, \sigma_3 \sim \text{HalfCauchy}(1) \\
 \text{Taste-specific means: } & \mu_T \sim \mathcal{N}(F_1, \sigma_1) \\
 \text{Optogenetics-specific means: } & \mu_O \sim \mathcal{N}(F_2, \sigma_2) \\
 \text{Taste-and-optogenetics-specific means: } & \mu_{T,O} \sim \mathcal{N}(F_3, \sigma_3) \\
 \text{Mean firing rate (with log link): } & \log(\text{firing}_{n,T,O}) = \mu_T + \mu_O + \mu_{T,O} \\
 \text{Observed number of spikes: } & S_{n,i,T,O} \sim \text{Poisson}(\text{firing}_{n,T,O})
 \end{aligned} \tag{2.1}$$

As explained in the introduction to the data analysis section, we used MCMC (specifically the NUTS sampler) to sample the posterior distribution of  $\text{firing}_{n,T,O}$  for every taste and optogenetics condition. We performed this analysis for every neuron in our dataset and finally calculated the impact of optogenetics on firing as the difference in  $\text{firing}_{n,T,O}$  between laser off (control) and their corresponding laser on trials. If the 95% Bayesian credible interval for these differences in  $\text{firing}_{n,T,O}$  for a neuron did not overlap 0, we concluded that the optogenetics significantly impacted the firing of this neuron (see the introduction to the data analysis section for a discussion of how Bayesian credible intervals inherently serve as significance tests).

### Regression of single neuron firing with palatability ranks

We analyzed, as we have done previously ([Sadacca et al. \(2016\)](#)), the time course of palatability-related information in the activity of single neurons by regressing their firing rates on the palatability ranks of the tastes (Dil Suc: 3, Conc Suc:4, Dil Qui: 2, Conc Qui: 1; higher is more palatable). In order to estimate the firing rates of neurons, we aggregated the spikes of each neuron, on a trial-by-trial basis, in 250ms bins moved by 25ms steps. We divided the aggregate number of spikes by the width of the bins (250ms) to obtain the near-instantaneous firing rate of each neuron across time on individual trials.

These firing rates, of course, vary widely between neurons. Furthermore, correlations between firing rate and palatability ranks may be significantly positive or significantly negative. We therefore needed to perform a 2-stage transform on neural firing before we could analyze all neurons as a group in our regression analysis. The first step was standardization - we

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

transformed the firing rate of each neuron in each time bin by subtracting the trial-averaged firing rate of the neuron in that time bin and scaling by its standard deviation across trials (to get z-scores), ensuring that the firing rates of all neurons were on a comparable scale. Next, we multiplied the standardized firing rate of each neuron by the sign of the time-averaged Spearman correlation coefficient between its firing and the palatability ranks. This ensured that the sign of the relationship of neural firing with palatability was the same for all neurons in our dataset, but left the magnitude of that relationship unaffected.

Our statistical model treats the standardized firing rate  $firing_{t,P,i}$  of a neuron at time bin  $t$  on trial  $i$  of a taste with palatability rank  $P$  as Gaussian-distributed with a mean  $\mu_{t,P}$  that depends linearly on  $P$ . We defined the palatability index in time bin  $t$ ,  $\beta_{Palatability,t}$ , as the change in  $\mu_{t,P}$  induced by a unit change in  $P$ .  $\beta_{Palatability,t}$  is, therefore, the slope of the line that explains  $\mu_{t,P}$  in terms of  $P$ , an estimate of the strength of the firing-palatability relationship. Using weakly informative Gaussian and Half-Cauchy priors for the mean and variance parameters respectively, our model formally says:

$$\begin{aligned}
 \text{Prior on palatability index: } & \beta_{Palatability,t} \sim \mathcal{N}(0, 1) \\
 \text{Prior on observation noise: } & \sigma \sim \text{HalfCauchy}(1) \\
 \text{Mean firing rate: } & \mu_{t,P} = \beta_{Palatability,t} \times P \\
 \text{Firing rate: } & firing_{t,P,i} \sim \mathcal{N}(\mu_{t,P}, \sigma)
 \end{aligned} \tag{2.2}$$

We used MCMC to infer the posterior distribution of  $\beta_{Palatability,t}$  across all neurons in our dataset (again, see above). The firing rate transformations defined previously put the activity of all neurons on the same scale and allowed us to infer a single posterior distribution of  $\beta_{Palatability,t}$  across all the neurons in our dataset. We repeated this regression for each time bin  $t$  from 0.25s before to 1.5s after taste delivery, obtaining posterior estimates of  $\beta_{Palatability,t}$  specific to each time bin. Finally, we normalized  $\beta_{Palatability,t}$  by subtracting its average baseline value (from 0.25 to 0s before tastes). We report the baseline-normalized  $\beta_{Palatability,t}$  as the palatability index  $\beta_{Palatability}$ .

### Characterizing the time course of the palatability index

In a manner similar to our previous work (Sadacca et al. (2016)), we modeled the time course of the posterior mean of the single neuron palatability firing index,  $\bar{\beta}_{Palatability}$ , with a logistic sigmoid. The difference between the lower and upper asymptotes of the S-shaped

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

logistic function fits the total rise in  $\bar{\beta}_{Palatability}$  across time, while its slope describes the rate of this rise. As  $\beta_{Palatability}$  was already normalized to its average pre-stimulus value, we set the lower asymptote of the logistic function to 0. With weakly informative Gaussian priors (restricted to positive values) on the upper asymptote ( $L$ ), slope ( $k$ ) and inflection time ( $t_0$ , ms post taste delivery) of the logistic sigmoid, our model is as follows:

$$\begin{aligned}
 \text{Prior on upper asymptote: } L &\sim \begin{cases} \mathcal{N}(0, 0.1) & L > 0 \\ 0 & \text{otherwise} \end{cases} \\
 \text{Prior on slope: } k &\sim \begin{cases} \mathcal{N}(1, 1.0) & k > 0 \\ 0 & \text{otherwise} \end{cases} \\
 \text{Prior on inflection time: } t_0 &\sim \begin{cases} \mathcal{N}(675\text{ms}, 75\text{ms}) & t_0 > 0 \\ 0 & \text{otherwise} \end{cases} \\
 \text{Prior on observation noise: } \sigma &\sim \text{HalfCauchy}(1) \\
 \text{Mean palatability index: } \bar{\beta}_{Palatability}(t) &\sim \mathcal{N}\left(\frac{L}{1 + e^{-k(t-t_0)}}, \sigma\right)
 \end{aligned} \tag{2.3}$$

We defined the peak of the palatability firing index,  $t_{peak}$ , as the time (post taste delivery) when  $\bar{\beta}_{Palatability}$  reached 95% of its maximum value,  $L$ . We transformed the posterior distributions of  $L$ ,  $k$  and  $t_0$  to get  $t_{peak}$  (inferred using MCMC) as follows:

$$t_{peak} = \frac{\ln \frac{95}{5}}{k} + t_0 = \frac{\ln 19}{k} + t_0 \tag{2.4}$$

### Modeling and change-point identification in ensemble firing data

As described in the Introduction (and Discussion), previous analyses reveal that rat GC population activity in response to a taste consists of a sequence of 3 coherent, abruptly-appearing ensemble states (Katz et al. (2001), Jones et al. (2007), Sadacca et al. (2012), Sadacca et al. (2016), Li et al. (2016)) in which firing rates “code”, in turn, taste presence, taste identity, and taste palatability; the transition into this last state has particular relevance for the prediction of palatability-related behavior in single trials, and is the subject of this study. While identifying these sequences typically requires several forward and backward passes through a dataset made up of many identical (i.e., unperturbed) trials, the work already published on the nature of these state sequences (see also Jones et al. (2007) and

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

Moran and Katz (2014b)) renders it possible (for the purposes of the current study) to more concretely define this process as involving ensemble firing change points between states having the following properties (also see Figure 2.6):

1. **Detection state:** a single distribution of population activity for all the tastes, indicating taste presence on the tongue.
2. **Identity state:** 2 distinct distributions of population activity, for the 2 taste identities in our experiments (Suc and Qui).
3. **Palatability state:** 4 distinct distributions of population activity, for the 4 taste palatabilities in our experiments (Dil Suc, Conc Suc, Dil Qui and Conc Qui).

With this characterization we were able to design a relatively simple change-point model that allowed us to detect these coherent transitions in population activity in individual trials. We first prepared the data for the change-point model by aggregating the spikes of each neuron in each trial into 10ms non-overlapping bins, indexing each neuron recorded in a session with a scalar  $i$  running from 0 to the number of neurons in the session  $N$ . We then converted the aggregate spiking data to a categorical format by marking each time bin by the index  $S$  of the neuron that spiked in that bin, with  $S = 0$  corresponding to no spikes from any neuron. If more than one neuron spiked in a time bin - a highly uncommon occurrence, given the relatively low firing rates of GC neurons and the small (10 ms) bins being used - we randomly selected one of the spiking neurons for assignment to that bin (Jones et al. (2007); Sadacca et al. (2016)).

With the (processed) categorical spiking data in hand, we now designed the change-point model to describe the ensemble firing in each of the 3 states (listed above) as categorical distributions with  $N + 1$  emissions, with 1, 2 and 4 such distributions corresponding to the detection, identity and palatability states respectively. Note that the results of this analysis are unchanged if we relax the parameters slightly to allow for 4 “state 2” distributions— that is, if we allow the Identity State to differ for the different concentrations of Sucrose and Quinine; this is probably because while many neurons may code different NaCl concentrations distinctly (Sadacca et al. (2012)), for other tastes the vast majority of neurons appear to code quality rather than concentration (see, for instance, Fonseca et al. (2018)).

We analyzed 1.5s of ensemble activity post taste delivery from each of the 4 optogenetic conditions in the 0.5s perturbation sessions. For the control (laser off) trials, this corre-

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

sponded to 0-1.5s of firing after taste delivery. On the perturbed trials, we ignored the 0.5s of activity when the lasers were on - for example, we analyzed 0.5-2.0s of firing post tastes when the lasers were on from 0-0.5s. In the resultant 1.5s of activity, we assumed that the change from detection to the identity state,  $C_I$ , happens anywhere in the interval  $[0.2s, 0.6s]$  (except the 0-0.5s perturbation trials, where we allowed the identity state to start earlier from 0.1s, to account for the possibility that some amount of taste processing happens in GC even while the neurons are being perturbed). The second change-point,  $C_P$ , from identity to palatability firing, was assumed to occur anywhere in the interval  $[C_I + 0.2s, 1.3s]$  (except the 0.7-1.2s perturbation trials, where the palatability state can start earlier at  $C_I + 0.1s$  for the same reason). This is equivalent to placing uniform priors over the intervals that define  $C_I$  and  $C_P$ , corresponding to the timing of sudden, coherent firing rate transitions in GC ensembles (Jones et al. (2007), Sadacca et al. (2016)).

$C_I$  and  $C_P$  are therefore latent variables of the change-point model that control the probabilities of the emissions actually observed. The Expectation-Maximization (EM) algorithm is the most widely used approach to perform inference in such models with latent variables; for stability and speed issues, we used a “hard-assignment” version of EM to fit the change-point model (Bishop (2016)). Starting with a randomly chosen set of initial emission probabilities  $\alpha_D$ ,  $\alpha_I$  and  $\alpha_P$  for the categorical emissions that define the detection, identity and palatability states respectively, the EM algorithm for our change-point model repeatedly cycled between 2 steps:

1. **“Hard” E-step:** Pick the combination of the latent variables,  $C_I$  and  $C_P$ , that has maximum posterior probability given the observed categorical spikes  $S$  and the ensemble firing probabilities  $\alpha_D$ ,  $\alpha_I$  and  $\alpha_P$ . We directly pick the mode of the joint posterior distribution of  $C_I$  and  $C_P$  in our hard-assignment version of the E-step instead of taking their expectations/means.
2. **M-step:** Set the categorical firing probabilities for each state to values that maximize the likelihood of the data given the  $(C_I, C_P)$  pair picked in the E-step. This is proportional to the number of emissions of each neuron in that state. For example, with  $S_t$  as the emission observed at time  $t$ , the likelihood-maximizing emission probabilities of

neuron  $n$  can be calculated as:

$$\begin{aligned}
 \text{In detection state: } \alpha_{D,n} &= \frac{\sum_{t=1}^{C_I} \mathbb{1}(S_t = n)}{\sum_{n=1}^N \sum_{t=1}^{C_I} \mathbb{1}(S_t = n)} \\
 \text{In identity state: } \alpha_{I,n} &= \frac{\sum_{t=C_I}^{C_P} \mathbb{1}(S_t = n)}{\sum_{n=1}^N \sum_{t=C_I}^{C_P} \mathbb{1}(S_t = n)} \\
 \text{In palatability state: } \alpha_{P,n} &= \frac{\sum_{t=C_P}^{1.5s} \mathbb{1}(S_t = n)}{\sum_{n=1}^N \sum_{t=C_P}^{1.5s} \mathbb{1}(S_t = n)}
 \end{aligned} \tag{2.5}$$

where  $\mathbb{1}$  is the unit function that is 1 when  $S_t = n$  and 0 otherwise.

In order to deal with the possibility that EM can get stuck at sub-optimal local maxima of log likelihood, we ran the algorithm from 100 different random initializations of the  $\alpha$  parameters. We monitored the log likelihood of the data given the model parameters and ran the algorithm to a convergence threshold of  $10^{-8}$  (or a maximum of 300 iterations). Finally, we picked the run with the maximum log likelihood at convergence and reported the change-points (and their posterior probabilities given  $S$  and  $\alpha$ ) found on this run.

It is worth noting that an inevitable result of performing such analyses on discontinuous data - such as trials in which 0.5s of spiking is missing because of optogenetic inactivation - is a certain number of artifactual change-points identified around the start or end of the inactivation time (the alternative is artifactually few change-points identified). This issue is handled in the Results and Discussion sections.

### Measuring aversive orofacial behaviors (gapes)

Bitter (e.g., Quinine) tastes cause rats to produce an orofacial behavior known as “gaping”, the purpose of which is to maneuver the offending substances to the front of the mouth for egestion. As such, gapes index the fact that the neural processing of the bitter taste has (in a certain sense) reached completion - the rat has “decided” that it does not want to ingest

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

the taste. The occurrence of gapes can be measured in a number of ways, the most common of which is *via* human coding of video recordings - in the best of circumstances, gapes are readily visible as large yawn-like movements.

Of course, the best of circumstances often fail to occur in rats free to move and rear. This fact, and the difficulty of getting precise measures of gape onset time from a visual record, renders video coding of gapes suboptimal for our purposes. Much more objective and less noise-ridden is evaluation of jaw electromyography (EMG), in which individual gapes are recognizable as particularly large-amplitude and large-duration electrical bursts (Figure 2.4A1-A2). We have previously built a quadratic classifier to detect these bursts in ongoing anterior digastric EMG signals, achieving 75% accuracy (Li et al. (2016)).

Even this approach has somewhat troubling limitations, however, as its failure to reach close to 100% accuracy indicates. These limitations stem from the facts that: 1) not all high-amplitude jaw movements are gapes; and 2) gapes vary widely in amplitude, and in fact some are small enough to appear similar in size to many other mouth movements (see Figure 2.4A1-A2). In practice, both types of variability leave the classifier subject to false positives that must be somehow recognized and removed - the former most notably at the beginning of trials (when the taste hits the tongue, causing 1-2 relatively large-amplitude licks).

One solution to these problems involves making simultaneous recordings from multiple jaw muscles, but pilot experiments left us concerned that such drastic infiltration of the jaw can compromise normal movement, which would make interpreting our results difficult. Instead, we decided to take advantage of another, more robust feature of gaping: the fact that gapes occur in 4-6 Hz “bouts” of anterior digastric activity (Travers and Norgren (1986), Li et al. (2016)). While identifying gaping bouts as time periods during which this rhythmicity dominates the EMG signal is also imperfect - it is probabilistic and involves smoothing across time - it largely solves the problems described above.

We instantiated just such a procedure here, applying a Bayesian spectrum analysis that estimates the posterior probability that a 4-6Hz rhythm underlies a short time series of EMG activity (see below for technical details). By this analysis, the probability of gaping to any taste is modestly elevated at trial onset (because of the initial large-amplitude licks), but it quickly drops to effectively zero for Sucrose, which therefore contributes nothing to the overall calculation of when gaping begins. On Quinine trials, in contrast, the probability waxes and wanes appropriately with the occurrence of gape bouts (Figure 2.4B1-B2), rising

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

precipitously and reliably just prior to the first gape (detected in a subset of data with both video recordings and the quadratic classifier, [Figure 2.4D](#)).

In important ways, this analysis is analogous to the method of divining palatability-relatedness of single-neuron firing described above and used in many previous studies ([Fontanini and Katz \(2006\)](#), [Sadacca et al. \(2012\)](#), [Li et al. \(2013\)](#), [Sadacca et al. \(2016\)](#), [Li et al. \(2016\)](#)) - the electrophysiological signal (in this case, the posterior probability of the range of gaping frequency in the EMG signal) varies (i.e., correlates) with the palatability of the proffered taste, and we average these correlations to ascertain the palatability-relatedness of the signal at each time point. Sucrose contributes no information to this signal (because rats do not gape to these sucrose concentrations), so the overall average gaping latency is equivalent to the difference between the time distributions of gaping probability to Dil and Conc Qui (see [Grill and Norgren \(1978a\)](#), [Travers and Norgren \(1986\)](#)), which can be statistically assessed as the Kullback-Leibler (KL) divergence (again, see technical details below). Not only does this procedure reveal the onset of orofacial behaviors reflecting aversion, it pits the two Qui concentrations against each other to get rid of most of the nonspecific gape-like EMG activity (mentioned above) which is of similar magnitude on both Dil and Conc Qui trials and does not contribute to the gape onset calculation.

Unlike previously used methods, in which (usually) trials where gapes could not be reliably detected were removed from further analysis, this algorithm combines EMG data from all the trials available, thereby allowing us to avoid making statistical comparisons between conditions with very different sample sizes. At the cost of being unable to precisely detect the specific timing of later gapes in a bout, this procedure provides an estimate of the average timing of the first gape (both robust and reliable enough for the purposes of the within-session, between-condition analyses performed here).

**Bayesian spectrum analysis (BSA) of EMG recordings:** As detailed previously, we recorded voltage signals from 2 unipolar EMG electrodes implanted in the anterior digastric muscle at 30kSamples/s. We used the difference in the voltage recorded by the 2 electrodes as the EMG signal - this procedure helps to cancel out any large artifacts produced by the animal's movements and is equivalent to using a differential amplifier (as done in [Li et al. \(2016\)](#)). We down-sampled the EMG signal to 1000Hz by averaging the voltage values in sets of 30, and highpass filtered the down-sampled signal above 300Hz ([Travers and Norgren \(1986\)](#); [Li et al. \(2016\)](#)) using a  $2^{nd}$  order Butterworth filter. The absolute value/magnitude

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

of the filtered EMG signal was then lowpass filtered (again using a Butterworth filter of order 2) below 15Hz, effectively capturing the envelope of variation of the EMG signal (plotted as the black curve in [Figure 2.4A1-A2](#)). This cutoff of 15Hz is sufficient for identifying orofacial behaviors, all of which occur at frequencies smaller than 10Hz ([Grill and Norgren \(1978a\)](#); [Li et al. \(2016\)](#)).

We subjected the envelope of the EMG signal to Bayesian spectrum analysis (BSA). BSA involves the construction of a probabilistic model of the generation of periodic signals from the superposition of sinusoids of different frequencies. We divided the signal on each trial into bins of width 300ms, with a step size of 1ms. We assumed that the EMG signal in each bin is produced by a sinusoid of a single frequency (plus noise) - in a probabilistic setting, this assumption implies the same model as a discrete-time Fourier transform. Contrary to the Fourier transform, however, BSA infers the posterior distribution of frequencies given the data. BSA has been shown to provide posterior estimates of frequencies that are an order of magnitude more precise than the Fourier transform ([Bretthorst \(2013\)](#); [Granqvist et al. \(2011\)](#)). We used the BaSAR R package for BSA ([Granqvist et al. \(2012\)](#)) and calculated the posterior probabilities of frequencies from 1Hz to 10Hz in 20 steps for each 300ms wide bin of data.

**Identifying the mean onset of aversive orofacial behavior:** Rats respond to intra-oral deliveries of Qui in the concentration range used in our experiments ( $10^{-4}$  to  $10^{-3}$ M) with an initial set of non-specific investigative licks that are followed by large, jaw-opening mouth movements called gapes ([Grill and Norgren \(1978a\)](#), [Figure 2.4A1-A2](#)). Gapes primarily involve activity of the anterior digastric muscle at 4-6Hz ([Grill and Norgren \(1978a\)](#), [Li et al. \(2016\)](#)) - we, therefore, used the probability of movements at 4-6Hz in the digastric EMG signal (from BSA, see previous section) as the probability of gaping ( $\text{Pr}_{\text{Gape}}$ ). This spectral measure of  $\text{Pr}_{\text{Gape}}$  has a strong correspondence with a previously-defined and above-discussed quadratic classifier (that tags individual mouth movements as gapes ([Li et al. \(2016\)](#))). On individual Qui trials ([Figure 2.4B1-B2](#)),  $\text{Pr}_{\text{Gape}}$  from BSA is high (close to 1.0) when the quadratic classifier tags mouth movements as gapes. In addition, the average probability of gaping ( $\overline{\text{Pr}}_{\text{Gape}}$ ) from BSA ([Figure 2.4C1-C2](#)) is very similar to an across-trial, peri-stimulus average of the gapes picked by the quadratic classifier. In contrast to the quadratic classifier, however, the BSA measure of  $\text{Pr}_{\text{Gape}}$  is based entirely on the spectral content of the EMG signal. It, therefore, does not require the construction of a sufficiently complex classifier

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

function (with a large enough set of experimenter-tagged examples to train the classifier) to pick out individual gapes. This also ensures that BSA considers bouts of movements together while calculating  $\text{Pr}_{\text{Gape}}$ , making it robust against isolated large amplitude movements early in the animal’s orofacial response. These initial movements were often found to be large licks on video and limited the accuracy of the quadratic classifier in Li et al. (2016) to 75%.

The probability of the transition from the rats’ initial investigative licks to gapes depends on the concentration of Qui delivered:  $10^{-3}\text{M}$  (Conc Qui) elicits gapes on more than twice the number of trials as  $10^{-4}\text{M}$  (Dil Qui) (Grill and Norgren (1978a), Li et al. (2016)). Comparison of  $\text{Pr}_{\text{Gape}}$  on Dil and Conc Qui trials, thus, provides a natural way to calculate the mean onset of gaping across all the Qui trials in an experimental condition (again, Suc trials add little to this analysis, as the probability of 4-6Hz activity drops to 0 within 100-200msec of taste delivery). We expect the distribution of  $\text{Pr}_{\text{Gape}}$  on Dil Qui trials to be similar to that on Conc Qui trials in the investigative licking phase. Once gaping starts, however, we expect a large difference in the distributions of  $\text{Pr}_{\text{Gape}}$  on Dil and Conc Qui trials.  $\text{Pr}_{\text{Gape}}$  on Dil Qui trials, therefore, acts like a baseline for  $\text{Pr}_{\text{Gape}}$  on Conc Qui trials: we conclude that gapes have started only when  $\text{Pr}_{\text{Gape}}$  of Conc Qui begins to differ significantly from this baseline.

We used Beta distributions to describe  $\text{Pr}_{\text{Gape}}$  on Dil and Conc Qui trials. The Beta distribution is commonly used to model the probability parameter of a Bernoulli (1/0) process<sup>2</sup>. Gaping being a Bernoulli process, the Beta distribution is an appropriate choice for modeling  $\text{Pr}_{\text{Gape}}$ . We defined one such Beta distribution in each time bin for Dil and Conc Qui separately, parametrized by the number of trials where the animal was gaping ( $\text{Pr}_{\text{Gape}} > 0.5$ ) or not ( $\text{Pr}_{\text{Gape}} < 0.5$ ). The Kullback-Leibler divergence of these Beta distributions ( $D_{\text{KL}}(\text{Conc Qui}||\text{Dil Qui})$ )<sup>3</sup> provides a natural way to quantify the difference between  $\text{Pr}_{\text{Gape}}$  on Dil and Conc Qui trials and shows a sharp jump  $\sim 1\text{s}$  post taste delivery (Figure 2.4E), consistent with the timing of the transition from investigative licks to gapes (Grill and Norgren (1978a), Travers and Norgren (1986), Li et al. (2016)). Finally, we calculated the cumulative sum of  $D_{\text{KL}}(\text{Conc Qui}||\text{Dil Qui})$  across time: the jump corresponding to the

<sup>2</sup>The Beta distribution for the parameter  $p$  of a Bernoulli process is expressed in terms of its concentration parameters,  $\alpha$  and  $\beta$ .  $\alpha$  = observed number of 1s and  $\beta$  = observed number of 0s.

<sup>3</sup>The KL divergence between two Beta distributions with concentration parameters  $(\alpha_1, \beta_1)$  and  $(\alpha_2, \beta_2)$  can be written as:  $D_{\text{KL}} = \log \Gamma(\sum_{j=1}^{j=2} \alpha_j) - \sum_{j=1}^{j=2} \log \Gamma(\alpha_j) - \log \Gamma(\sum_{j=1}^{j=2} \beta_j) + \sum_{j=1}^{j=2} \log \Gamma(\beta_j) + \sum_{j=1}^{j=2} (\alpha_j - \beta_j)(\psi(\alpha_j) - \psi(\sum_{j=1}^{j=2} \alpha_j))$ , where  $\Gamma$  and  $\psi$  are the gamma and digamma functions respectively.

mean onset of gaping is expressed as a change in slope of the cumulative sum. We fit two straight lines to the cumulative sum to capture this change in slope: the intersection of the two lines defines the mean timing of the onset of gaping (Figure 2.4F).

## 2.4 Results

### 2.4.1 Experimental paradigm and data overview

Figure 2.1A depicts the preparation used for our experiments - IOCs for taste delivery, bilateral GC opto-trodes for recording of neural ensemble activity and delivery of laser light, and EMG electrodes in the anterior digastric (jaw) muscle for simultaneous assaying of consumption-related mouth movements. Four weeks prior to the surgery in which we installed these assemblies, we injected AAV carrying the optogenetic silencer ArchT (along with green fluorescent protein - GFP) into GC. The GFP allowed us to confirm (post-mortem) infection of GC neurons by immunohistochemical verification of the GFP tag (Figure 2.1B).

The rats received intra-oral deliveries of 30mM sucrose (Dil Suc), 300mM sucrose (Conc Suc), 0.1mM Quinine-HCl (Dil Qui) and 1mM Quinine-HCl (Conc Qui). One set of sessions involved “brief perturbation” trials: on 75% of the trials in these sessions, we inhibited GC neurons for 0.5s, beginning either at 0s, 0.7s or 1.4s post taste delivery (Figure 2.1C). These three perturbation windows tile the period containing the temporal epochs that characterize GC taste responses (Katz et al. (2001), Sadacca et al. (2012), Sadacca et al. (2016)). More specifically, the earliest (0-0.5s) and latest (1.4-1.9s) inhibitions affect GC neurons before and after the range of likely transition times into the behaviorally-relevant state containing palatability-related firing, which typically occur just before, during, or just after the middle (0.7-1.2s) perturbations (Figure 2.1C, also see Figure 2.1D for a basic schematic of coding across the first 2.0s of GC taste responses). In a separate set of experimental sessions (performed using a subset of the same rats), we inhibited GC across the entire duration of the taste responses (0-2.5s post stimulus) (Figure 2.1C) as a control comparison for the brief 0.5s perturbations.

We recorded the activity of 244 GC single neurons across 10 sessions ( $24.4 \pm 13$  units/session) of 0.5s inhibition, and of an additional 73 GC single neurons in 5 sessions ( $14.6 \pm 4.7$  units/session) of 2.5s inhibition. The two types of experimental sessions were counterbalanced, such that 3 rats received 2.5s inhibition sessions first, and 2 received 0.5s inhibition

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

sessions first. No differences with order were noted.

The AAV-ArchT construct used in this study has been shown to infect neurons of multiple types (e.g., pyramidal neurons and interneurons) in an unbiased manner (Aschauer et al. (2013)). Our optogenetic inhibition protocol, therefore, can be thought of as a general perturbation of the dynamics of GC neurons in response to tastes. Note as well that any such perturbation (including of individual neuron types) would be expected (perhaps paradoxically) to enhance the firing of some neurons through network-level effects (like disinhibition, *via* suppression of the firing of inhibitory neurons, Allen et al. (2015)). This expectation was borne out in the data: the firing of most of the recorded GC units (146/244, 60%, example unit in Figure 2.2A1-A4) was significantly suppressed when the laser was switched on for 0.5s, but the firing of an additional 20% (49/244) was significantly enhanced.

The same pattern of results was observed when the duration of optogenetic inactivation was increased to 2.5s: the firing of 82% of GC neurons (60/73, example unit in Figure 2.2B1-B2) was inhibited, and the activity of 15% (11/73) was enhanced. The fact that 2.5s of laser stimulation appeared to inhibit a larger percentage of neurons is likely an artifact of analysis methods: suppression of the low firing-rates (3-10Hz) that dominate GC taste responses (Katz et al. (2001), Jones et al. (2007), Samuelson et al. (2012), Kusumoto-Yoshida et al. (2015), Mazzucato et al. (2015)) can be difficult to detect, particularly in short time windows; consistent with this, we observed that the highest likelihood of detecting suppression in 0.5s perturbation sessions occurred when that perturbation was delivered in the middle of taste processing (0.7-1.2s, Figure 2.2C) - at the time of peak firing rate modulations. With 2.5s of inactivation, which covered the entirety of GC taste responses, we naturally had the power to detect suppression in a larger fraction of neurons (Figure 2.2D).

Although this specific optogenetic protocol cannot be used to answer cell-type/microcircuit-specific questions, its network-wide effects are ideal for testing the macroscopic dynamical properties of taste processing in GC (the purpose of the current work): GC taste responses evolve through a sequence of temporal epochs (Katz et al. (2001), Maffei et al. (2012), Jezzini et al. (2013)) which have the hallmarks of emergent, quasi-stable states of a system that can be speculatively described, at a high level, as an attractor network (Jones et al. (2007), Miller and Katz (2010), Mazzucato et al. (2015), Sadacca et al. (2016)); our optogenetic protocol brings about a strong perturbation of the network activity characterizing these stable states, and by mapping the state dependence of the effects of these perturbations, we are able to directly test the proposed function of these states (and of the transitions between them).

### 2.4.2 Early perturbations delay single-neuron palatability-related responses while late perturbations do not

We first assessed the impact of optogenetic perturbation on neural activity - that is, on the palatability-related content of GC taste responses that had been smoothed (using 250ms-wide windows moved in 25ms steps) and standardized to be on a uniform scale (see Materials and Methods for details). The set of responses (1 per taste) were regressed against the palatability ranks of the taste stimuli (Conc Suc:4, Dil Suc:3, Dil Qui:2, Conc Qui:1) to obtain a palatability index,  $\beta_{Palatability}$ . Being a Bayesian analysis (consult Materials and Methods for details on model setup and inference), this regression gives access to the entire posterior distribution of  $\beta_{Palatability}$  at every time point. Knowing the spread of the posterior distribution of  $\beta_{Palatability}$  at every time point allows us to more simply perform significance tests: we can conclude that  $\beta_{Palatability}$  is different from 0 at the 5% level of significance if the 95% extent of its posterior distribution (generally known in Bayesian analyses as the “credible interval”) does not overlap 0 (such time points are marked by dots in [Figure 2.3A](#)). We used logistic sigmoid functions to better characterize the time evolution of the posterior mean of  $\beta_{Palatability}$  (shown with dashed lines in [Figure 2.3A](#)), and defined the size and latency (time to attain 95% of maximum size) of the upper asymptote of the logistic fit as the magnitude and latency of the peak of  $\beta_{Palatability}$  respectively.

As expected, perturbation for 2.5s had a devastating impact on palatability-related responses of neurons in the affected GC network ([Figure 2.3A](#)). In control (laser-off) trials, as in previous studies ([Sadacca et al. \(2016\)](#)),  $\beta_{Palatability}$  climbed to an asymptote  $\sim 0.8$ s after taste delivery. However, on trials where the lasers were switched on at the time of taste delivery and left on for 2.5s,  $\beta_{Palatability}$  never rose significantly from 0. Note that the latency to peak palatability firing is comparable in the two conditions (blue bars in [Figure 2.3B](#)), but that the magnitude of the peak is close to 0 when GC neurons are being perturbed (red bars in [Figure 2.3B](#)).

The impact of brief (0.5s) perturbations on the palatability content of single-neuron GC taste responses was smaller in magnitude, but could be quite dramatic with regard to peak timing, depending on when the perturbation occurred ([Figure 2.3C](#)). In these sessions, just as in the 2.5s perturbation sessions,  $\beta_{Palatability}$  peaked  $\sim 0.8$ s after taste delivery when the lasers were left off. Furthermore, neither the timing nor magnitude of this peak was significantly affected by perturbation of GC neurons in the later part of the taste response (1.4-1.9s, after

palatability-related firing had emerged).

In contrast, if activity was perturbed for the first 0.5s of the GC taste response, the palatability content of this response did not reach asymptote until  $\sim 1.3$ s, a lag of almost 0.5s compared to the control condition (laser-off trials). Note that the failure of GC firing to “bounce back” immediately after laser-off (which occurred 300-400ms before the time of peak palatability content in control trials) implicates GC in the processing of palatability itself (see Discussion). Note as well that despite delaying the peak of  $\beta_{Palatability}$ , the early perturbation did not affect its later emergence - if anything, the magnitude of the peak was larger in this condition (red bars in [Figure 2.3C](#)). The 0-0.5s perturbation thus appears to produce a transient shift out of the attractor dynamics responsible for GC taste responses followed by gradual relaxation back into the stable state after the end of the perturbation; variability in this process (which can overshoot the stable point, depending on the speed of relaxation) could explain the apparent increase in the magnitude of the peak palatability index in this condition.

Finally, 0.5s perturbations delivered in the middle of the taste response (0.7-1.2s) also had a powerful impact on GC palatability-related firing: the magnitude of the peak of  $\beta_{Palatability}$  was significantly lower in this condition (red bars in [Figure 2.3C](#)); the latency of this peak, meanwhile, was (like that produced by earlier perturbations) about 0.5s later than no-laser trials. The former effect was unsurprising, as this particular perturbation overlaps the heart of palatability-related activity in GC neurons ([Katz et al. \(2001\)](#), [Jezzini et al. \(2013\)](#), [Sadacca et al. \(2016\)](#)).

### 2.4.3 GC perturbation delays the onset of aversive orofacial behavior

We monitored our rats’ mouth movements *via* electromyography (EMG). Specifically, we implanted EMG electrodes in the anterior digastric muscle; as a jaw moving muscle, the anterior digastric plays a major role in the production of “gapes”, the rhythmic orofacial behavior that serves to move aversive tastants to the front of the mouth in preparation for expelling. Far less accessible tongue muscles underlie mouth movements that support behaviors (such as “lateral tongue protrusions”) that help the rat prepare to ingest appetitive tastants ([Grill and Norgren \(1978a\)](#), [Travers and Norgren \(1986\)](#), [Li et al. \(2016\)](#)). For that reason, we focus solely on gapes in this study (but see Discussion).

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

Individual mouth movements can be recognized as bursts of anterior digastric EMG activity (Figure 2.4A1-A2). However, the variability in the amplitudes and durations of these EMG bursts reduces our ability to separate gapes from other large mouth movements. We, therefore, made use of a more robustly distinctive feature of gaping – the fact that gapes occur in 4-6Hz bouts (Travers and Norgren (1986), Li et al. (2016)). We analyzed the spectral content of the envelope of the EMG signal using Bayesian spectrum analysis (BSA; see Materials and Methods for a detailed discussion) and measured the probability of gaping as the total posterior probability of 4-6Hz movements.

While easier to calculate and less subject to error, this estimate of the probability of gaping has strong correspondence with gaping bouts identified by a classifier trained on individual bursts of EMG activity (Li et al. (2016), see Figure 2.4B1-B2); the trial-averaged probability of gaping calculated by BSA and more classic techniques are also similar, for both trial types in which gaping occurred (Dil and Conc Qui trials, Figure 2.4C1-C2). Finally, the fact that the probability of gaping jumps precipitously just before the first gape as identified on video (Figure 2.4D) confirms this algorithm’s reliability in identifying periods of gaping in the EMG signal (see Materials and Methods for more details).

With this information in hand, we were able to investigate the effects that perturbations of GC activity have on the animals’ rejection of aversive Qui. On average, gaping begins  $\sim 0.9$  sec after Qui delivery in control trials – that is, when analysis is restricted to trials in which the laser was off (trials in which GC neurons were not perturbed, (Figure 2.5A)). This latency is consonant with that reported in video analysis (Grill and Norgren (1978a)) and classic burst-oriented analysis of EMG (Travers and Norgren (1986)). Furthermore, this estimate matches that observed in control rats (published in Sadacca et al. (2016) and Li et al. (2016)) that received neither laser nor ArchT expression. Thus we can conclude that, at least with regard to the driving of gaping, our preparation leaves the system capable of normal function.

Previous work has shown that while the appearance of palatability-related firing in GC (which arises suddenly and coherently across neurons in single trials) robustly predicts the onset of gaping bouts (see below and Sadacca et al. (2016)), it is unrelated to the mechanics of individual gapes within gaping bouts (Grill and Norgren (1978b), Li et al. (2016)). We therefore predicted that GC perturbations delivered once gaping was already underway would have minimal impact on gaping behavior.

In fact, our data show that rats gaped normally, with gape bouts beginning at approxi-

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

mately the same time as in control (no laser) trials, if perturbations arrived late in the trial (1.4-1.9s, [Figure 2.5B](#)). Furthermore, this late perturbation failed to prematurely end gaping bouts that had already begun. Figures 5C1-C4 show example trials in which the probability of gaping rhythm in the EMG signal went high following Conc Qui delivery, and stayed high despite late (1.4-1.9s) GC inhibition. In fact, the percentage of trials in which gaping was maintained into this period was unchanged by late GC perturbation - 57% (36/63) of control trials vs 55% (26/47) of laser trials. We can thus conclude that GC is of no consequence for the maintenance of ongoing gaping.

In contrast, GC activity plays a clear role in the initiation of gaping. GC perturbations occurring 0-0.5s after taste delivery - that is, before transitions into the palatability-related state of GC activity - delayed gaping onset by approximately 0.25s on average ([Figure 2.5B](#)). This delay cannot be explained in terms of removal of early gaping - gaping latencies as early as 0.5s after taste delivery were rare, and an analysis of control (no laser) trials showed that removing latencies of less than 0.5s had essentially no impact on the mean onset time of gaping. The much more likely explanation is that GC inhibition (which is inevitably partial, see Discussion) perturbs the ongoing process that leads to the release of a “decision to gape” signal visible in GC ([Sadacca et al. \(2016\)](#)).

Similarly, GC perturbations timed to occur squarely around the average time of the palatability / decision-related neural state change (0.7-1.2s; see [Sadacca et al. \(2016\)](#)) delayed the onset of gaping until just before 1.2s after taste administration - approximately 0.25s after gaping on control trials and in control sessions (with no laser or ArchT). That is, brief disruptions of GC activity occurring before or during the “heart” of quinine processing had a strong impact on the latency of aversive orofacial behavior. Not only is the impact of brief optogenetic perturbation significant, it was every bit as large as that observed with whole-trial (i.e., 2.5s) perturbations, which delayed the appearance of gaping by  $\sim 0.2$ s ([Figure 2.5B](#)). These long perturbations are not discussed further, because they had the additional unintended consequence of impacting gaping behavior on control trials (see [Figure 2.5A](#) and Discussion).

#### 2.4.4 GC perturbation impacts orofacial behavior only if delivered before the onset of palatability-related ensemble activity

We have previously demonstrated that the temporal dynamics of GC taste responses are well described as sudden transitions between two stimulus-specific ensemble firing rate “states” (Jones et al. (2007)), the latter of which is laden with information about stimulus palatability and highly predictive of the latency of gaping on single trials (Sadacca et al. (2016)); the trial-to-trial variability of both behavioral and transition latencies is large (the neural transition happens at a range of latencies spanning the approximate interval between 0.4 to 1.5s, and the behavior follows close behind), such that trial averaging smears the changes in firing rates into a more gradual-seeming ramp.

We timed our 0.7-1.2s perturbations to overlap with the transition into this palatability-related ensemble activity state, but due to the above-described variability in timing, there were inevitably a subset of trials in which the ensemble state transition occurred before the perturbation. This fact afforded us an opportunity: we predicted that identical 0.7-1.2s perturbations would impact gaping latency differently depending on whether the transition into the late ensemble activity state had already occurred in that specific trial; this prediction implies that the results in Figure 2.5B, averaged across all trials receiving the perturbation, occlude our ability to see the diversity of that perturbation’s possible effects, and mask a larger impact of the perturbation in one independently identified subset of trials.

While we have previously used Hidden Markov Models (HMMs) to detect ensemble firing rate transitions in GC responses to tastes (Jones et al. (2007), Moran and Katz (2014b), Sadacca et al. (2016)), this analysis is not amenable to the data in our study: a dataset made up of all 4 trial types (early, middle, and late perturbation, plus control) would be complex enough (each trial type would likely involve distinct sets of firing rates, see below) that the HMM would seldom reach stable solutions; divided into individual trial types, meanwhile, the datasets would be too small to allow convergence to even simple stable fits. Instead, we took advantage of the insights gained from our previous publications (Katz et al. (2001), Fontanini and Katz (2006), Jones et al. (2007), Grossman et al. (2008)) and built a constrained change-point model of GC population activity; specifically, the model consisted of 2 activity change-points, the latter of which introduced palatability-related firing. This model constrained the general HMM framework in a way that allowed us to estimate transitions in individual trial types (see Figure 2.6 and Materials and Methods for details).

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

The distributions of putative transition times (identified by the change-point model) into the palatability-related ensemble state are shown in [Figure 2.7A](#) for all Qui trials in which GC firing was perturbed from 0.7s to 1.2s post stimulus. As firing rates were suppressed during the perturbation, we did not attempt to identify change-points when the lasers were on (a fact that inevitably impacted change points that could be identified at the “edges” of the excised time period; see below and Methods). According to this algorithm, the palatability-related state emerged before the lasers were illuminated on 55% of the trials (92/168, but see below); on the remaining 45% of trials (76/168), the palatability change-point could not be identified before laser onset. Regression analysis allowed us to confirm that significant palatability-related information appeared before 0.7s in trial-averaged single neuron firing during trials in which the ensemble state transition occurred prior to laser onset time; this information was notably lacking in trials in which the transition had not occurred ([Figure 2.7B](#)).

On the basis of this analysis, we were able to show that, in line with our expectations, identical 0.7-1.2s perturbations had distinctly different effects on the onset of gaping depending on whether or not the transition into palatability-relatedness appeared to have occurred prior to laser perturbation ([Figure 2.7C](#)). Perturbations that arrived before the ensemble transition delayed gaping by more than  $\sim 0.5$ s - that is, gaping appeared more than 0.2s after the end of GC inhibition in these trials. A comparison with control data confirmed that this effect was not caused by a simple truncation of the distribution of gaping latencies: even when we restricted ourselves to analyzing only the proportion (31%, 52/168) of control trials which lacked any gaping-related EMG activity till 1.2s (which was, in perturbation trials, the laser off time), the average gaping latency was still significantly less than that observed in the (larger) subset (45%) of laser trials in which the ensemble transition failed to precede the 0.7s onset of GC inhibition. Clearly, GC perturbation perturbs consumption behavior, if that perturbation begins prior to the ensemble neural transition into palatability coding.

Gaping occurred significantly earlier in trials in which the ensemble transition to the high-palatability state preceded the onset of GC perturbation at 0.7s ([Figure 2.7C](#)), but contrary to our expectation, gaping was still somewhat delayed compared to the no-laser condition even in these trials. As the ensemble state transition purportedly happens by 0.7s on these trials (i.e., earlier than the average transition time on control trials), we expected that the onset of gaping would be similarly expedited. This was not the result that we obtained. We considered several possible explanations for this result (see Discussion), the most reasonable

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

of which seemed the possibility that some transitions identified as happening just prior to laser onset were artifactual - the inevitable effect of attempting to identify firing rate changes next to a data “edge” (see Methods and Discussion) - and thus that for a small subset of trials in this group transitions into palatability coding did not in fact precede laser onset. Note that this hypothesis would also explain why the percentage of trials in which pre-0.7s transitions were identified was somewhat larger than expected (see above, and compare to the grey bars showing transition times in the no-laser control trials in [Figure 2.7A](#)).

We tested this hypothesis, and found that the delay in the onset of gaping can indeed be entirely attributed to the trials where the ensemble state transition is calculated to occur between 0.65 and 0.7s. Specifically, when we restricted our analysis to trials in which the ensemble transition happened at 0.65s or earlier, the onset of gaping was found to occur more than 300ms earlier than in control trials. We went on to examine the trials in which transitions were identified to occur between 0.65 and 0.7s, and found that “early-onset” gaping occurred in only a subset (15) of these trials - almost precisely the same number (14) as there were control trials in which the transition occurred in the 0.65-0.7s interval; this result suggests that those true transitions that occurred during this interval likely resulted in gaping that was unaffected by the laser perturbation.

As a whole, our results demonstrate that the impact of brief optogenetic inhibition of GC depends on precisely when that inhibition occurs: laser perturbation of GC that begins after the onset of palatability-/decision-related firing utterly fails to impact the timing of aversive orofacial responses, but GC perturbation that begins before the transition significantly delays those responses. Furthermore, given the trial-to-trial variability in the issuance of this decision-related transition, the result of any particular timing of brief GC inhibition will differ in different particular trials, depending on precisely what state the brain has achieved prior to that perturbation. This result provides support for our overarching hypothesis that the onset of palatability-related population activity in GC marks a discrete shift in taste processing - the ensemble transition in taste-related firing that predicts behavior is in fact the emission of the decision to gape.

## 2.5 Discussion

Perception and action are inextricably linked in cortical taste responses. Neurons in gustatory cortex (GC), the primary sensory cortical area for taste, exhibit responses that, across 1.5s of post-stimulus time, shift from first reflecting stimulus identity to predicting a rat’s consumption decision (Katz et al. (2001), Fontanini and Katz (2006), Sadacca et al. (2012), Maier and Katz (2013b)). With ensemble analysis, these otherwise gradual-seeming changes in firing rates are revealed to be swift, coherent transitions between population activity “states” (Jones et al. (2007)) - transitions that vary widely in latency from trial to trial, and that are therefore effectively blurred out in stimulus-aligned averages. Despite (in fact, because of) their highly variable latencies, these ensemble firing states reliably precede the onset of ingestion-egestion mouth movements by  $\sim 0.2-0.3$ s (Sadacca et al. (2016), Li et al. (2016)), predicting not only the nature but the latency of these movements in single trials.

Here we show that GC neural ensemble dynamics described above are not merely “efferent copy” reflections of processes occurring elsewhere, but are instead an indication of processing that is (to at least some extent, see below) intrinsic to GC. Brief (0.5s) optogenetic perturbations of GC neurons impact the timing of the animal’s decision to expel a bitter taste in the mouth, but only if those perturbations begin before the neural ensemble has shifted to palatability-related firing. Thus, a unique moment in time (the shift of population activity to reflect stimulus palatability), despite being enormously variable in latency from trial-to-trial, reflects a tipping point in taste processing; cortical disruptions have no impact beyond this tipping point, as the control of the ongoing movements themselves shifts elsewhere (presumably to brainstem pattern generators that control the ingestion-egestion mouth movements themselves in real time, see Travers et al. (1997), Travers et al. (2000)).

A massively interconnected network of forebrain regions underlies or reflects taste processing - in addition to GC, this network includes the central and basolateral nuclei of the amygdala (CeA and BLA, Nishijo et al. (1998), Grossman et al. (2008), Fontanini et al. (2009), Sadacca et al. (2012)), hippocampus (Ho et al. (2011)), lateral hypothalamus (LH, Yamamoto et al. (1989), Li et al. (2013)), the bed nucleus of the stria terminalis (BNST, Norgren (1976), Li and Cho (2006)), the parabrachial nuclei of the pons (Baez-Santiago et al. (2016)), and the nucleus of the solitary tract (NTS, Di Lorenzo and Lemon (2000)). Several of these brain regions have been shown to integrate sensory and motor aspects of taste stimuli in their responses (Sadacca et al. (2016), Baez-Santiago et al. (2016), Denman

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

et al. (2018)). Furthermore, multiple forebrain regions send direct descending feedback to the primary brainstem taste regions, influencing both their activity (Di Lorenzo (2000), Cho et al. (2003), Li et al. (2005)) and generation of orofacial movements (Zhang and Sasamoto (1990), Berridge and Valenstein (1991), Shammah-Lagnado et al. (1992), Travers et al. (1997)). Given this widely distributed network of processing nodes, it is to be expected that perturbation (or disruption over long periods of time) of one (or a few) of the participatory nodes will initiate homeostatic mechanisms that minimize the resultant degradation of behavior; thus, it is unsurprising that rodents remain able to produce gapes following ablation (King et al. (2015)) or disruption of GC (Li et al. (2016)) - in fact, the basic gaping response to quinine is produced even in decerebrate rats (Grill and Norgren (1978b)). Nonetheless, we find that brief perturbations of GC do significantly alter these behaviors (as do lesions of other areas, such as gustatory thalamus, Grill and Norgren (1978b)), proving that far more than the minimal circuit is involved in triggering them *in situ*.

Longer disruptions of GC activity appear to have lasting effects that can confound the interpretation of their behavioral impact - our 2.5s long optogenetic perturbations delayed the onset of gaping even in control (no laser) trials. Such spillover effects may reflect cellular or network-level processes, but they cannot be attributed to cell death caused by the perturbation: in our case, similar optogenetic protocols have been shown to have no observable impact on cell integrity in GC, even for perturbations much longer than 2.5s (Maier et al. (2015), Flores et al. (2018)); furthermore, the same rats in later sessions produced normally-timed orofacial responses on the control trials. We suggest that, to at least some degree, such effects on behavior reflect the widespread nature of taste processing, and the status of GC as one participatory node.

Despite being just one node of this large network of brain regions, our brief perturbations reveal a temporally-specific role of GC in the driving of orofacial behavior - a role that could not be discerned through wholesale disruption of activity. This conclusion is bolstered by findings showing that: 1) even early - i.e., pre-transition - GC perturbations delay gaping; and 2) palatability-related firing does not immediately return to normal levels following cessation of perturbation (as would be expected if GC was simply an output path reflecting processing performed elsewhere). Our 0.5s perturbations reveal that GC contributes to the instigation of a gaping bout but plays no role in the maintenance of gaping once it begins. These data suggest a dynamic flow of processing control within the larger taste network: modulatory signals propagate out of GC (signals that likely develop under the

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

guidance of basolateral amygdala; [Piette et al. \(2012b\)](#)) to influence the choice of a motor program in brainstem circuits, which is then implemented and controlled locally. At its heart, the proposed role of cortex in this model of taste processing has deep similarities to the role of neuromodulatory systems in the circuits underlying *Aplysia* feeding ([Dacks and Weiss \(2013\)](#)), leech swimming ([Crisp and Mesce \(2004\)](#)), control of gastric rhythms in the lobster and crab ([Marder and Bucher \(2007\)](#)), and rat whisking ([Hattox et al. \(2003\)](#)); in each, temporal aspects of rhythmic motor programs produced autonomously by a pattern generating circuit are influenced by descending signals.

The discreteness, coherence and inter-trial variability of GC ensemble dynamics has several attractor network-like properties ([Hopfield \(1982\)](#), [Amit \(1992\)](#)): 1) attractor networks with multiple quasi-stable states can reproduce the sudden switches of activity seen in GC ensembles ([Miller and Katz \(2010\)](#)); 2) the transition durations and state lifetime statistics of GC population dynamics are more in line with a dynamically switching attractor model than linear models of firing rate evolution ([Jones et al. \(2007\)](#), [Sadacca et al. \(2016\)](#)); and 3) nonlinear attractor-based circuits that exploit the noise inherent in neural processing more optimally perform the decision to ingest or expel a taste, which rats need no training to perform, than do linear integrating circuits ([Miller and Katz \(2013\)](#)). Our optogenetic protocol, with its mix of inhibitory and excitatory effects, presumably introduces a transient disruption in such attractor dynamics; such a perturbation is strong enough to transiently “knock” the network out of stability, but only if it hasn’t already settled into the eventual, decision-related stable state.

The finding that the involvement of GC in the gape instigation process appears to last almost precisely 50ms past the calculated transition times could conceivably be explained in many ways. Firstly, transitions between quasi-stable states of GC processing, however discrete, are certainly not instantaneous - the time constants of neural firing ensure that there is some finite (albeit small) amount of time across which the ensemble makes the “jump” from one state of activity to another. In addition, it is worth noting that both HMMs and change-point analysis techniques provide only a noisy estimate of state transition times, even if the transitions themselves were instantaneous. While both of these explanations have merit, it is also clear that the change-point analysis model, which must deal with a sudden change in firing introduced by the laser, identifies artifactual “change-points” close to the laser onset time on some of the trials, even if palatability firing actually began after the lasers were switched off. Our analysis suggests that some, if not all, of the seeming

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

response delay following change-points occurring between 0.65 and 0.7s may be artifactual, which in turn suggests that GC perturbations may have no impact even scant milliseconds following ensemble transitions. It is worth noting in this context that gaping lags 0.2-0.3s behind the ensemble neural transition (Sadacca et al. (2016)); thus, it appears that GC becomes irrelevant following the emission of a “gape signal”, even before actual gaping has begun.

In this study, we focused exclusively on gapes, the orofacial responses that rats make to expel aversive tastes from the oral cavity. Pilot attempts to implant EMG electrodes in deeper muscles that control the distinctive consumption behaviors that occur in response to palatable tastes resulted in unacceptable levels of distress for the animals. This means that it remains (remotely) possible that gapes and LTPs are produced by separate cortical mechanisms (Peng et al. (2015)), and that therefore our results are informative only about aversion. We consider this possibility highly unlikely, however, for several reasons: 1) GC ensemble firing reflects the palatability of both appetitive and aversive tastes (Figure 2.3, Katz et al. (2001); also see Fonseca et al. (2018)), even if palatability is modified by learning (Moran and Katz (2014b)); 2) the latency and inter-trial variability of the onset of palatability-related ensemble activity is similar for palatable and aversive tastes (Sadacca et al. (2016)); 3) there is considerable overlap in the brainstem circuits that underlie gapes and LTPs (Travers et al. (2000), Chen and Travers (2003), Venugopal et al. (2007), Moore et al. (2014)), resulting in similar latencies in the onset of LTPs and gapes after taste delivery (Travers and Norgren (1986)); and 4) independent analysis has suggested that orofacial behaviors reflecting aversiveness and palatableness lie on a single parametric continuum (Breslin et al. (1992)). These lines of evidence are consistent with the suggestion that cortex plays similar roles in the initiation of LTPs and gapes, which leads us to speculate that the transition of GC population activity to reflect stimulus palatability marks a shift in processing control, irrespective of the palatability of the tastant.

In summary, the balance of our results demonstrate a dynamic role for cortex in the processing of tastes; because this role involves ensemble activity states with variable trial-to-trial latencies, it cannot be discerned using standard analyses that average across trials. They reveal the importance of a unique moment in time that, despite being massively variable in latency from trial to trial, denotes a reliable shift of processing control - a modulatory signal emerging (at least partly) from cortical circuits that is passed (presumably) to a brainstem central pattern generator. These results suggest an attractor-like network of

## *CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES*

activity (although they could also be consistent with networks with thresholds), potentially spread across interconnected brain regions, underlying the animal's decision to ingest or expel the tastant in the mouth - perturbations to this network can disrupt its functioning transiently, but only if it has not yet settled into the final, behaviorally-relevant stable state.

### **2.6 Figures**

CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

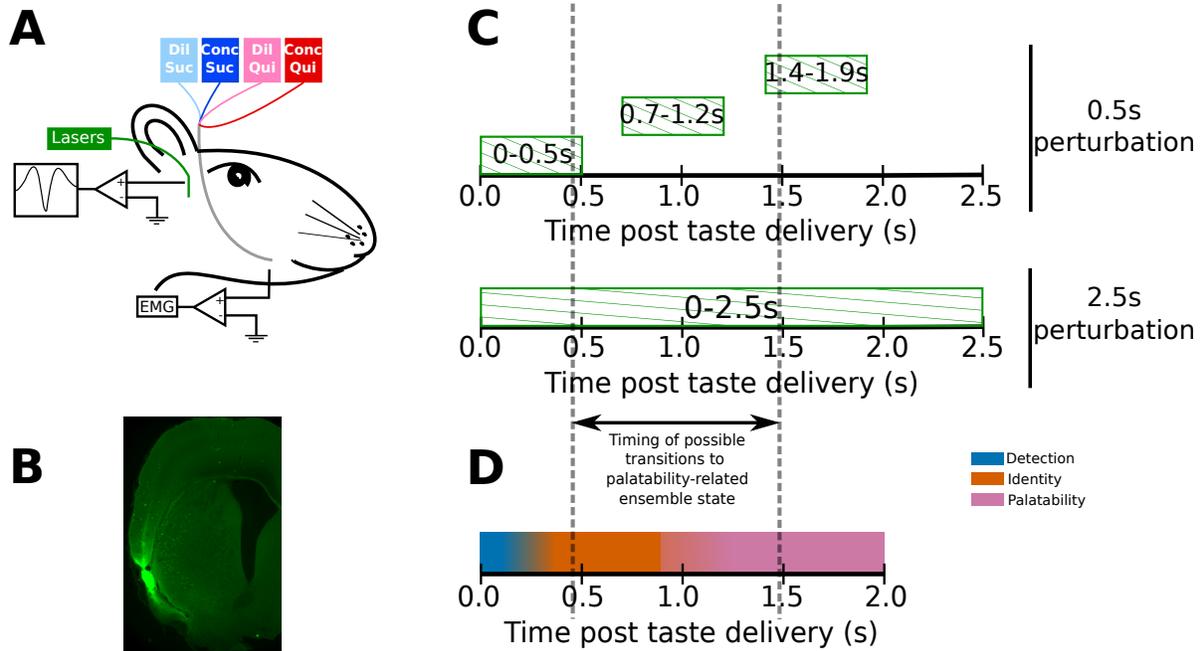


Figure 2.1: Experimental paradigm. **A**: 4-6 weeks after receiving surgeries for virus injections, rats were implanted with opto-trodes and EMG electrodes. Post recovery, they were given intra-oral infusions of Dil Suc (30mM Sucrose), Conc Suc (300mM Sucrose), Dil Qui (0.1mM Quinine-HCl) and Conc Qui (1mM Quinine-HCl), and ArchT-expressing GC neurons were briefly inhibited by green (532nm) laser light. **B**: Coronal slice from a subject, showing ArchT expression (visualized by the GFP tag) localized in gustatory cortex (GC). A small lesion, left by the tip of the opto-trode is visible in the middle of the GFP expressing region, had no general impact on behavior (see below). **C**: Inhibition protocol used in the study: two types of optogenetic perturbations, short (0.5s) or long (2.5s), were delivered in separate experimental sessions; short perturbations were delivered at one of three possible time points on any individual trial. Not shown, but delivered in all sessions, were control trials with no perturbations. Grey dashed lines mark the approximate range of the ensemble transitions to palatability/decision-related firing. **D**: A schematic of the temporal structure of single-neuron coding across the first 2.0s of taste responses in GC. Immediately following taste presentation, responses are nonspecific, indicating only the presence of fluid on the tongue (“detection” epoch). The next two temporal epochs of GC firing are taste specific: the first codes the physio-chemical identity of the stimulus (“identity” epoch); following a transition (that can happen anywhere between 0.5-1.5s post stimulus on individual trials, see grey dashed lines, and on average happens midway through this period) firing rates change to reflect palatability and the upcoming consumption decision (“palatability” epoch).

## CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

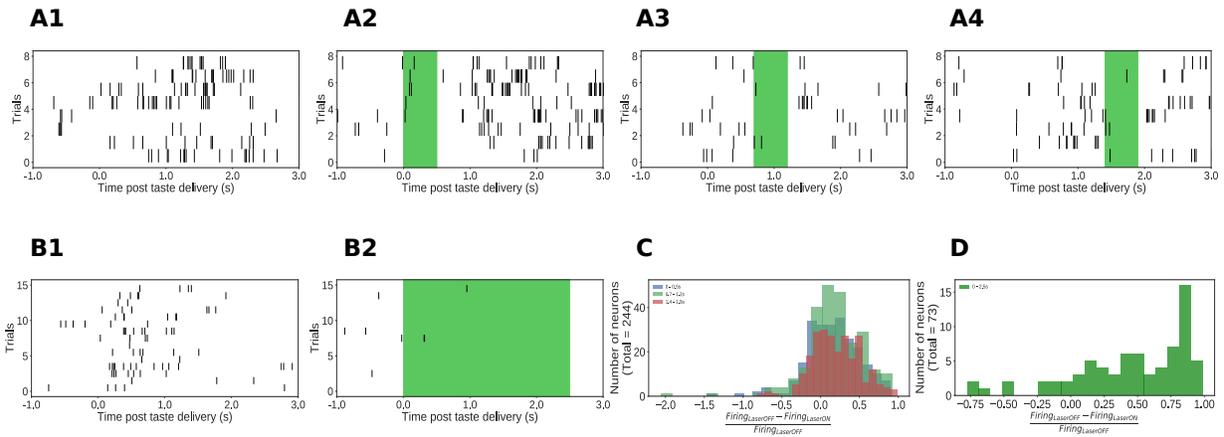


Figure 2.2: Impact of ArchT-mediated inhibition on GC neurons. **A1-A4**: Rasters of spiking in an example single GC neuron in a 0.5s-perturbation session; each hash mark is an action potential. Activity is robustly suppressed during laser stimulation. **B1-B2**: Analogous data from an example single GC neuron in a 2.5s perturbation session, also showing clear inhibition during laser stimulation. **C**: Histogram of changes in firing rates (plotted as a fraction of the firing rate on control trials, x-axis) produced by 0.5s perturbations across the entire sample (y-axis = number of neurons). The majority of neurons show robust firing suppression when perturbed (fraction  $> 0$ ), but a small group of neurons actually increased their firing rates in response to perturbation, presumably due to network-level effects (fraction  $< 0$ ). **D**: Analogous histogram of changes in firing rate produced by 2.5s perturbation. Almost all neurons were affected by the perturbation: the large majority are suppressed, but a small minority show elevated firing rates in response to perturbation.

CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

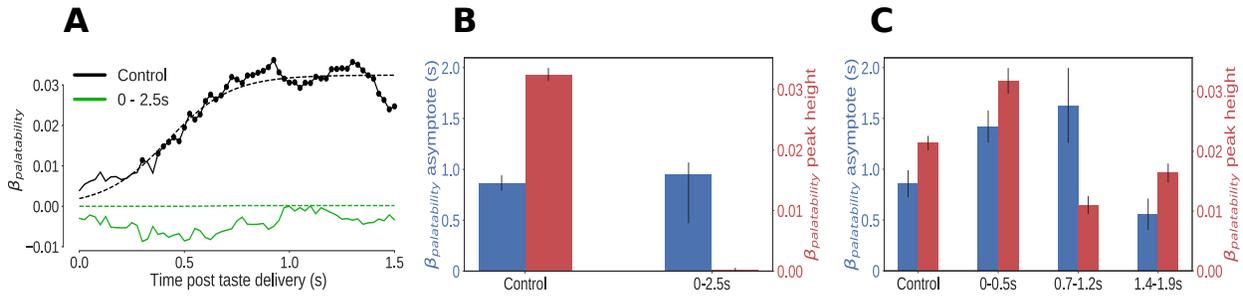


Figure 2.3: Impact of optogenetic perturbations on palatability relatedness of the firing of GC neurons. **A:** Coefficients (palatability-relatedness, y-axis) obtained from the regression of trial-averaged firing rates on palatability ranks of the taste stimuli across time (x-axis). The solid lines depict the mean regression coefficient across time for the entire data sample; coefficients significantly different from 0 at the 5% level are marked by dots. The dashed lines are logistic sigmoid fits for each condition. Disruption of GC firing for 2.5s wipes out the entirety of the palatability response. **B:** The post-stimulus latency (blue bars and y-axis) and magnitude (red bars and y-axis) of the peak (95% of the asymptote) of the sigmoid fits in **A**. Error bars denote 95% Bayesian credible intervals; differences are statistically significant at the 5% level if bars are not overlapping. On control (laser off) trials, GC neurons asymptote to peak palatability firing  $\sim 0.8$ s post stimulus. The 2.5s perturbation, by disrupting the palatability response completely, is fit by a flat sigmoid whose peak magnitude overlaps 0, although the latency to “peak” is similar to that of control trials. **C:** Analogous graph of post-stimulus latency (blue bars and y-axis) and magnitude (red bars and y-axis) of the peak (95% of the asymptote) of the sigmoid fits for each trial type in the 0.5s-perturbation sessions. Error bars denote 95% Bayesian credible intervals; differences are statistically significant at the 5% level if these error bars are not overlapping. On laser off trials, GC representation of palatability peaks  $\sim 0.8$ s after taste delivery, identical to the 2.5s perturbation control trials in **B**. Perturbations early (0-0.5s) and in the middle of the taste response (0.7-1.2s) delay the peak of palatability firing by  $\sim 0.5$ s; the magnitude of this peak, however, is the smallest for the middle perturbation. Perturbations late in the taste trial (1.4-1.9s), after palatability-related firing has mostly subsided, have (as expected) no impact compared to control trials.

CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

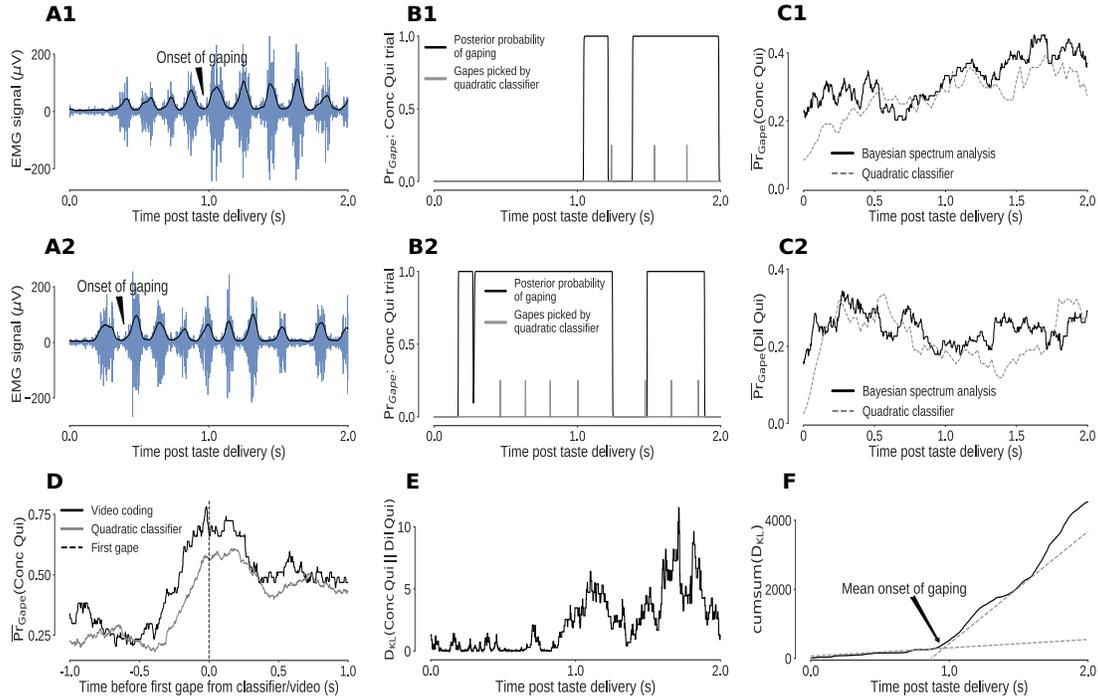


Figure 2.4: Bayesian spectrum analysis (BSA) of anterior digastric EMG recordings - probability of gaping calculated in terms of the total posterior probability of 4-6Hz movements. **A1-A2:** Two representative Conc Qui trials. The animal’s mouth movements can be seen as bursts of higher-amplitude (y-axis) EMG activity (blue) following taste delivery - the onset of gaping, as detected on video, is marked. The time series of the envelope of the EMG signal (black line) are the data subjected to BSA. **B1-B2:** Result of BSA brought to bear on a pair of individual Conc Qui trials. The calculated probability of gaping (y-axis, black lines) matches up with individual gapes (grey vertical hash marks) picked by a previously published quadratic classifier that achieved 75% accuracy; while correlating well with the earlier technique, BSA avoids multiple pitfalls of that technique (and is easier to apply, see Methods). **C1-C2:** BSA (solid line) and the quadratic classifier (dotted line) produce similar estimates of trial-averaged probability of gaping in response to Dil Qui (**C1**) and Conc Qui (**C2**) on a set of control (laser off) trials. **D:** The probability of gaping from BSA rises reliably just before the first gape. Gaping probability was averaged across trials aligned by the time of the first gape, detected either on video (black) or by the quadratic classifier (grey). The black dashed line (0 on the x axis) indicates the occurrence of the first gape. **E:** KL divergence between the probability of gaping to Conc and Dil Qui (higher values indicate larger differences in their gaping distributions, same trials as in **B**). As expected, the distributions of gaping probability on Conc and Dil Qui trials are initially similar (while non-specific investigative licks happen) and diverge out at  $\sim 1$ s post stimulus once gaping begins. **F:** The cumulative sum of the KL divergence in **E** across time. The jump in KL divergence around the mean onset time of gaping is seen as a change in slope of its cumulative sum. We fit two straight lines to the cumulative sum and pick their intersection as the mean onset of gaping across this set of trials.

CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

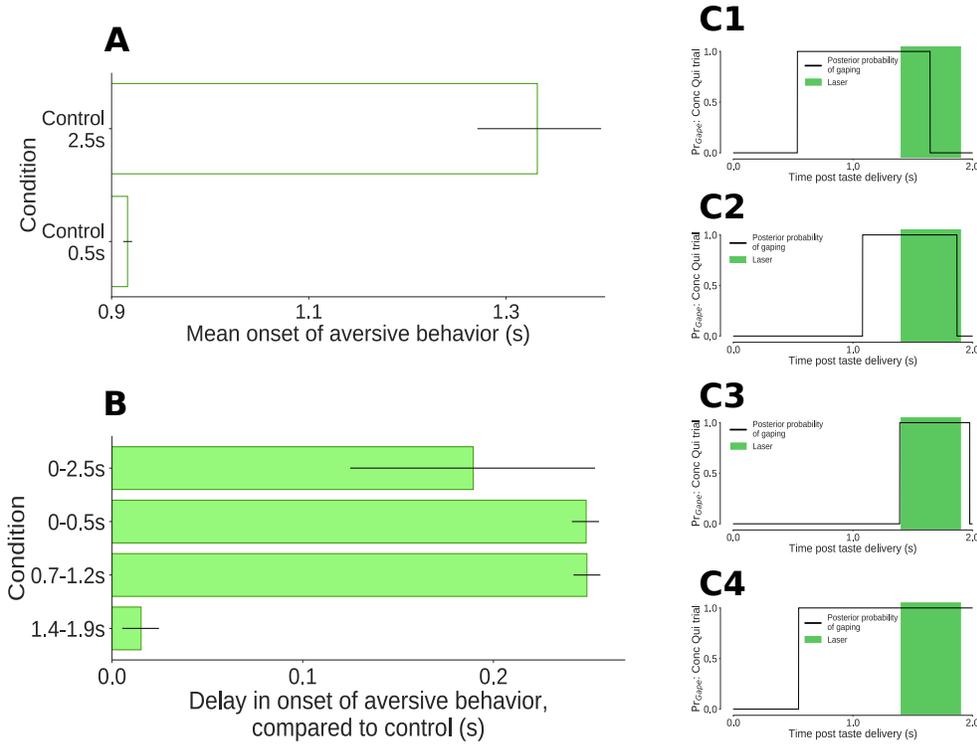


Figure 2.5: Onset times of 4-6Hz aversive orofacial behaviors (gapes) under different conditions. **A**: Onset of aversive orofacial behaviors in control (no laser) trials in 0.5s and 2.5s perturbation sessions. The x-axis presents the mean gape onset times; the extent of their 95% Bayesian credible intervals are shown in the error bars. Non-overlapping error bars depict statistical significance at the 5% level. The 2.5s controls show a delayed onset, likely due to lasting effects of the (relatively) long optogenetic perturbation. **B**: Delay in the onset of aversive orofacial behaviors (compared to control trials) with 2.5s perturbation (top bar), and in the different 0.5s laser trials, with the same conventions as **A**. Early (0-0.5s) and mid-trial (0.7-1.2s) perturbations of the taste response delay the onset of gaping (to the same degree as 2.5s perturbation). The delay in the onset of gaping is insignificant if GC neurons are disrupted late in the trial (1.4-1.9s). **C1-C4**: Four representative Conc Qui trials with optogenetic disruption from 1.4-1.9s post taste delivery. On each of these trials, the probability of 4-6Hz aversive orofacial responses is unaffected by the onset of the laser, confirming that GC perturbation fails to disrupt ongoing bouts of gaping.

CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

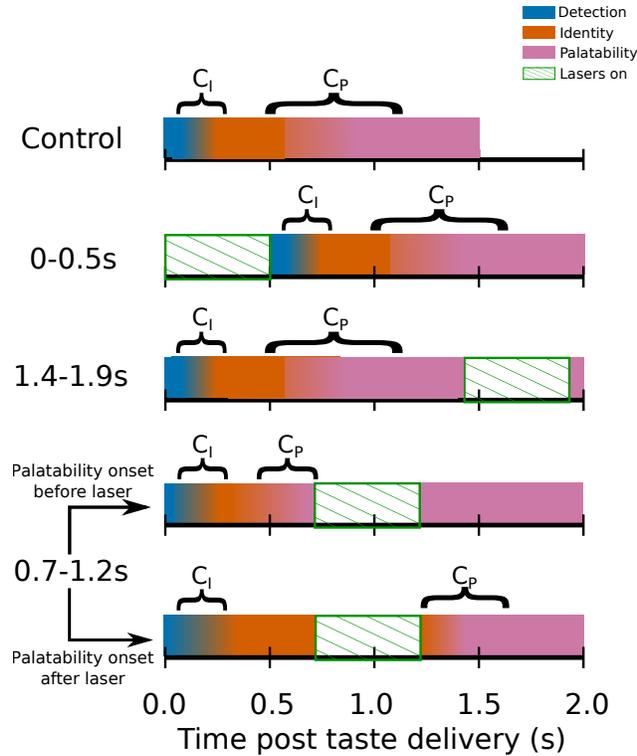


Figure 2.6: Switchpoint model of GC ensemble responses to tastes, which were assumed (on the basis of our previous work) to consist of 3 states as follows: 1) **Detection**: a brief, initial state of nonspecific responses with identical population distributions of activity for each tastant in our battery; 2) **Identity**: responses related to the chemical identity of the taste stimulus with 2 population firing rate distributions, one each for Suc and Qui; 3) **Palatability**: population firing rich in palatability and consumption-decision related information with 4 population distributions of activity, one for each of the 4 tastants in our stimulus battery. The model assumed that the transitions between these states could not occur during the optogenetic perturbation of GC (denoted by periods of green diagonally hatched regions): each row shows how the search for change points is hypothesized to be impacted by GC perturbation; note the two distinct possibilities with regard to 0.7-1.2s trials.

CHAPTER 2. OPTOGENETIC PERTURBATION OF CORTICAL STATES

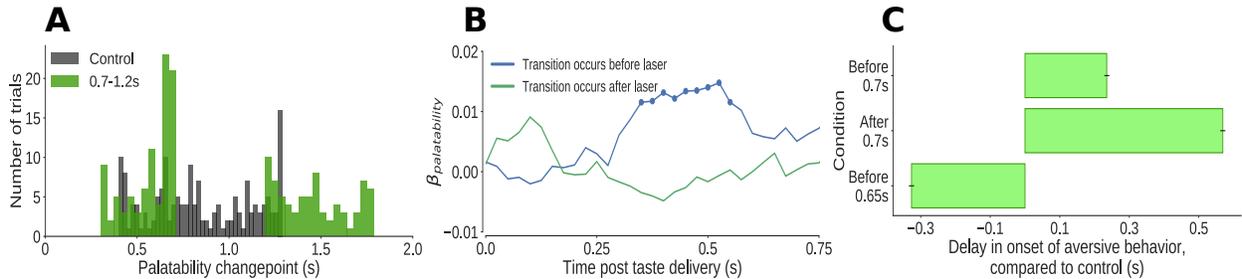


Figure 2.7: The impact of 0.7-1.2s perturbation on GC neural activity and aversive orofacial behavior varies from trial to trial, depending on the progress of taste dynamics. **A**: Distribution of change-points into the palatability-related ensemble state identified in Qui trials (in green). We could not examine the time period of perturbation, from 0.7s to 1.2s, because firing during this period was deeply confounded by laser-induced inhibition: as we concatenated the pre- and post-perturbation periods, an abnormally large number of change points are localized to the time of splicing (compare to change points identified in control trials, in grey). **B**: Correlation (quantified in terms of coefficient of regression) of trial-averaged firing rates of GC neurons with palatability of the taste stimuli in two subsets of trials - those in which the ensemble transition into palatability-related firing was identified to have occurred prior to perturbation (blue line), and those in which it did not (green line). Coefficients significantly different from 0 at the 5% level are marked by dots; these coefficients differ from 0 only within the trials in which the palatability-related ensemble state appeared before the onset of perturbation. **C**: The impact of 0.7-1.2s GC perturbation on the onset of aversive orofacial behavior, quantified in terms of the delay of behavior onset compared to control trials (x-axis). The onset of gaping is delayed significantly more if the perturbation begins before palatability information has appeared in ensemble activity than if it does not - but even on these trials behavior is significantly delayed. When we drop the subset of trials in which transition times into palatability-related firing occur within 50ms of 0.7s (a subset that likely contains artifactually identified transitions), however, gaping on transition-before-perturbation trials is revealed to happen earlier than on control trials - the expected result (see text).

## Chapter 3

# The role of the gustatory cortex in incidental experience-evoked enhancement of later taste learning

### Co-author contributions

The experiments and results from this chapter were published in *Learning & Memory* as [Flores et al. \(2018\)](#).

The following people contributed to this study (and were co-authors on the published paper):

1. Veronica L. Flores: Designed research, performed research, analyzed data, wrote paper, gathered funding.
2. Tamar Parmet: Performed research.
3. **Narendra Mukherjee**: Provided reagents/tools, analyzed data.
4. Sacha B. Nelson: Wrote paper.
5. Donald B. Katz: Designed research, analyzed data, wrote paper.
6. David Levitan: Designed research, performed research, analyzed data, wrote paper.

## 3.1 Abstract

The strength of learned associations between pairs of stimuli is affected by multiple factors, the most extensively studied of which is prior experience with the stimuli themselves. In contrast, little data is available regarding how experience with *incidental* stimuli (independent of any conditioning situation) impacts later learning. This lack of research is striking given the importance of incidental experience to survival. We have recently begun to fill this void using conditioned taste aversion (CTA), wherein an animal learns to avoid a taste that has been associated with malaise. We previously demonstrated that incidental exposure to salty and sour tastes (taste pre-exposure - TPE) enhances aversions learned later to sucrose. Here, we investigate the neurobiology underlying this phenomenon. First, we use immediate early gene (c-Fos) expression to identify gustatory cortex (GC) as a site at which TPE specifically increases the neural activation caused by taste-malaise pairing (i.e., TPE did not change c-Fos induced by either stimulus in isolation). Next, we use site-specific infection with the optical silencer Archaelhodopsin-T to show that GC inactivation during TPE inhibits the expected enhancements of both learning and CTA-related c-Fos expression, a full day later. Thus, we conclude that GC is almost certainly a vital part of the circuit that integrates incidental experience into later associative learning.

## 3.2 Introduction

Consistent pairing of specific taste stimuli with strong reinforcement leads animals to adapt their future responses to those stimuli, thereby making them more successful at consuming nutrients and avoiding toxins. In the lab, the most well-known variety of this adaptive process is called conditioned taste aversion (CTA), wherein animals learn to avoid a taste conditioned stimulus (CS) that has been paired with malaise-inducing unconditioned stimulus (US). While complex, CTA is known to involve: 1) changes in a brainstem-amygdalar-cortical circuit (Bures et al. (1998); Grossman et al. (2008)); and 2) synaptic plasticity governed by the degree of the association between the CS and US (Garcia et al. (1966b); Revusky (1968); Nachman (1970); Ahlers and Best (1971); Balsam et al. (2002); Frankland et al. (2004); Molet and Miller (2014); Adaikkan and Rosenblum (2015)).

Of course, reliable pairings of stimulus and reward are quite rare in the ongoing stream of experience. Most taste stimuli are seldom experienced with strong reinforcement - they are

### CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

“innocuous”, meaning they occur incidentally. Nonetheless, these experiences are important for survival, in that ostensibly innocuous stimuli can have a measurable impact on behavioral adaptations caused by learning - that is, on associative memory strength (Walters and Byrne (1983); Fanselow and Poulos (2005); Johansen et al. (2011); Kandel et al. (2014)). An extensive body of research has shown, for instance, that CTA memory strength is decreased by unreinforced pre-exposure to the CS or US, which renders the stimuli familiar and less salient (e.g., Lubow and Moore (1959); Lovibond et al. (1984)), and pinpoints possible neural loci of these effects (Weiner (2010)).

This work leaves unstudied, however, the potential impact of the most common stimuli - those other than the CS and US in some eventually-experienced learning paradigm. There are at least two reasons why it is reasonable to ask whether even totally “incidental” experience with a set of tastes might in fact have an impact on learning about a new taste: 1) general environmental “enrichment” has been shown to affect both neural development and specific sensory responses (Liu and Urban (2017)); and 2) “innocuous” stimuli have been suggested to enhance sub-threshold learning experiences (Ballarini et al. (2009)) and latent inhibition (Merhav and Rosenblum (2008)). Still, virtually no work has explicitly investigated how incidental taste experience might change the function of CTA learning circuits in the brain. This noticeable gap in the literature is a potentially significant limiting factor on our ability to generalize the results of lab experiments to the human condition - incidental taste experience is omnipresent in the natural world, a fact that stands in stark contrast to the laboratory, in which learning experiments are performed on animals that have never tasted anything but (mild, nearly tasteless) chow.

We have recently begun an inquiry into this topic (Flores et al. (2016)), showing that experience with salty and sour tastes (hereafter “taste pre-exposure” or TPE) enhances aversions toward a novel taste; experience with both tastes is more effective than experience with either alone, and three sessions of TPE is more effective than two. These results, which contrast with both classic interference effects that reduce conditioning strength (Bouton (1993); Kwok et al. (2012)) and the above-mentioned effects that occur across an entirely different time scale (Riege (1971); Donato et al. (2013); Leger et al. (2014)), demonstrate that benign experience with one set of tastes over two to three days can impact the strength of learning established in a later associative conditioning paradigm using a different taste.

While much remains to be learned about this behavioral phenomenon, the above results do enable us to formulate basic hypotheses regarding how sensory taste information acquired

during TPE is processed in the brain and integrated into future learning. An appropriate place to start in this regard is with primary gustatory cortex (GC), which resides in the anterior insula: GC has been amply shown, in electrophysiological studies (Katz et al. (2001); Grossman et al. (2008); Moran and Katz (2014b); Sadacca et al. (2016)), immediate early gene imaging (Desmedt et al. (2003); Contreras et al. (2007); Inberg et al. (2016); Uematsu et al. (2014)), and loss of function experiments (Berman (2001); Stehberg et al. (2011); Levitan et al. (2016); Li et al. (2016)) to play a role in mediating taste behavior and CTA learning, as well as taste novelty/familiarity (Gallo et al. (1992); Rosenblum et al. (1993); Koh et al. (2003); Bahar et al. (2004); Koh and Bernstein (2005); Roman and Reilly (2007); Merhav and Rosenblum (2008); Lin et al. (2012); Inberg et al. (2013); Hadamitzky et al. (2015)). Here, we test the hypothesis that GC is vital for representing incidental taste experience and for driving the impact of that experience on later CTA learning.

Specifically, the current study: 1) directly characterizes CTA-training-related c-Fos expression—an immunohistological marker for neural activation—in GC following TPE; 2) replicates the previously-shown TPE phenomenon in virus-infected rats; and 3) tests the impact of optogenetic silencing of GC activity during TPE on later learning. Our results support the hypothesis that GC plays an important role in integrating incidental taste experience with future learning (without changing processing of the taste experience alone), thereby mediating later memory formation - both TPE itself, and manipulation of GC specifically during TPE, impact a future CTA (and learning-related neural activation) towards a novel taste.

## 3.3 Materials and Methods

### 3.3.1 Subjects

127 naïve adult (6-8 weeks, 225-250g at time of surgery) female Long Evans rats acquired from Charles River Laboratories (Wilmington, MA) served as subjects for all experiments. Females were chosen for their behavioral temperance, and because published evidence does not show major sex differences in CTA trainability (Randall-Thompson and Riley (2003); RINKER et al. (2008); Dalla and Shors (2009)). Rats were housed individually in humidity- and temperature-controlled cages (Innovive), kept on a 12-hour light-dark cycle, and given ad libitum access to food and water (prior to experiments) which was replaced twice a week. At

## CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

least 10 days post arrival to the facility, all animals were randomly assigned to experimental groups.

All procedures were conducted in accordance with the guidelines established by the Brandeis University Institutional Animal Care and Use Committee (IACUC).

### 3.3.2 Stimuli

Taste solutions used for taste pre-exposure sessions (TPE) consisted of 0.01M sodium chloride (N) and 0.02M citric acid (C), as well as distilled water (W). Both concentrations are comfortably above detection thresholds for rats (Kolodiy et al. (1993); SCALERA (2004); Li et al. (2012); Sadacca et al. (2012)). These specific stimuli were used to ensure experience with both palatable and less palatable tastes - this concentration of N is modestly palatable, and this concentration of C is mildly aversive. A novel, innately palatable 0.2M sucrose (S) solution was always used as the conditioned stimulus in all experiments to facilitate detection of aversion learning. Note that while this means that we did not counterbalance completely, for instance making N the conditioned stimulus and S a TPE stimulus for some rats, in our previous work we demonstrated some basic generalizability, showing that N and C were similarly effective as single TPE stimuli (Flores et al. (2016)).

### 3.3.3 Experimental apparatus

Experiments were conducted in the morning, following a 21-hr water deprivation period. All sessions occurred in a Plexiglass experimental chamber (8.5 x 9.5 x 11.5 in) that was distinct from the rats' home cages. The experimental chamber and bottles were rinsed and sterilized before and after each use.

### 3.3.4 Experiment 1

#### Surgery

Rats were anesthetized using a ketamine/xylazine mixture (1ml ketamine, 0.05 ml xylazine/kg body weight) delivered via intraperitoneal injection. The head was shaved and positioned into a stereotaxic device, after which the scalp was exposed, leveled, and cleaned. Four self-tapping support screws were implanted into the skull. Rats were then removed

## CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

from the stereotax and laid prone for bilateral implantation of an intra oral cannula (IOC) - flexible hollow plastic tubes inserted parallel to the masseter muscle into the mouth posterolateral to the first maxillary molar (Phillips and Norgren (1970)). A stable, rigid dental acrylic head cap was formed around the IOCs and skull screws.

Following surgery, rats were given analgesic (meloxicam 0.04 mg/kg), saline, and antibiotic (Pro-Pen-G 150,000 U/kg) injections. Additional antibiotic and analgesic injections were delivered 24 and 48 hours later. The weight of each animal was recorded each day; any rat displaying lethargy, lack of grooming or weight loss greater than 15% of pre-surgery weight were removed from the study. All rats were given 7 days of recovery post-surgery before any experimentation.

### **Adaptation sessions**

Following recovery, rats were given 2 days of access to distilled water through a bottle in the experimental chamber for approximately 30 min. This ensured familiarization with the testing environment.

### **TPE/WPE sessions**

TPE and water pre-exposure (WPE) sessions, which followed adaptation, were identical to those used in our previous experiments (Flores et al. (2016)). Each rat ( $n = 66$ ) received 1 such session per day - 100 aliquots of either tastes (pseudo-random ordering) or water alone, delivered to the oral cavity via IOC (brief opening of a solenoid valve caused 50 $\mu$ l of fluid to be delivered) at 15-second inter-trial intervals, for a total of 5 mL of fluid. This protocol continued for 3 consecutive days (Figure 3.1). The IOC was used for delivery of TPE because it ensured experimenter control, such that TPE rats consumed equal volumes of unpalatable C and palatable N and WPE rats received the same volume of W.

### **Conditioning sessions**

A single conditioning session took place the day after completion of TPE/WPE, in the same experimental chamber. Sucrose CS was delivered via bottle (5mL available for 5 min) and then IOC (60 deliveries for a total of 3mL; Figure 3.1). This procedure allowed us both to: 1) take advantage of the literature indicating that c-Fos expression is stronger with bottle than IOC conditioning (Wilkins and Bernstein (2006)) and 2) ensure substantial and

### CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

reliable sucrose consumption. Immediately after sucrose consumption animals in the aversion conditioning groups (TPE,  $n = 15$  and WPE,  $n = 15$ ; [Figure 3.1](#)) received subcutaneous injections of lithium chloride (LiCl, 0.3 M, 0.5% of current body weight) to induce the malaise US. Use of this concentration of LiCl, which is lower than that typically employed to induce CTA, ensured that CTA learning would be sub-maximal, thereby allowing us to observe enhancements of learning ([Nachman and Ashe \(1973\)](#); [Stone \(2005\)](#); [Levitan et al. \(2016\)](#)).

In addition to the above-mentioned WPE controls, additional control rats were given TPE followed by either: 1) S paired with subcutaneous injections of harmless saline - essentially S alone ( $n = 9$ ); or 2) administration of LiCl alone, without pairing of S ( $n = 8$ ; [Figure 3.1](#)). Finally, two further groups of controls received WPE followed by S + saline ( $n = 10$ ) or LiCl-alone conditioning sessions ( $n = 9$ ; [Figure 3.1](#)). LiCl/saline Injections were administered on the experimenter's lap (a location distinct from both testing chamber and home cage), to ensure that malaise would not be associated with a context that could confound the results (or cause a great deal of potentially novelty-related c-Fos). Rats were then briefly (see below) returned to their home cages with access to ad lib food and water.

#### **Brief access task**

An additional set of animals ( $n = 8$ ) took part in a brief access task (BAT, aka the "Davis rig," Med Associates Inc.) to provide a separate, consumption test of the possibility that TPE changes sucrose processing itself.

The procedure was similar to previously used protocols to investigate neophobia ([Lin et al. \(2012\)](#); [Monk et al. \(2013\)](#)). Following IOC surgeries, recovery and adaption (see above), rats underwent 2 days of water habituation in the BAT rig, during which they learned to drink from the single lick spout in the chamber. Water deprivation for the 21 hours before each session ensured motivated drinking. At the start of each BAT trial, a mechanical shutter was raised to allow access to the lick spout for 15 mins (15mL of W), after which the shutter descended. The experimental chamber and bottles were rinsed and sterilized before use for each animal and session.

Following the two days of water habituation, animals underwent TPE ( $n = 4$ ) or WPE ( $n = 4$ ) sessions as described above. 24 hours after the last of these 3 pre-exposure sessions, they were returned to the BAT Rig, and this time given access to novel sucrose (15mL available

### CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

for 15 mins). Consumption, lick rate and interlick intervals were recorded for each rat. Rats were then returned to their home cages with access to ad lib food and water.

None of the rats in this experimental group ( $n = 8$ ) were given CTAs toward sucrose or perfused for immunohistochemistry analyses. Although animals were presented with a contextual change in experimental location between the TPE and BAT sessions, usage of the BAT was necessary to test neophobic reactions, which are only evident via voluntary consumption (see Results).

#### **c-Fos immunohistochemistry**

To capture peak c-Fos expression levels, rats were deeply anesthetized with an overdose of the same ketamine/xylazine mix used for surgery ninety mins after the conditioning session and perfused transcardially with isotonic phosphate buffered solution (1X PBS) followed by 100ml of ice-cold 4% paraformaldehyde. Brains were then extracted and post fixed in 4% paraformaldehyde for three days, after which coronal brain slices (60  $\mu\text{m}$ ) containing the region of interest were sectioned on a vibratome. Sections were chosen based on anatomical landmarks for Gustatory cortex (GC (+2.5 mm for anterior GC, +1.4 mm for middle GC, +0.36 mm for posterior GC relative to bregma; see [Paxinos and Watson \(2007\)](#))) and on published reports demonstrating the regions importance to CTA or taste processing ([Katz et al. \(2001\)](#); [Fontanini and Katz \(2006\)](#); [Piette et al. \(2012a\)](#); [Sadacca et al. \(2012\)](#); [Maier and Katz \(2013b\)](#); [Schier et al. \(2014\)](#); [Schier et al. \(2015\)](#)), the driving of taste-responsive behavior ([Li et al. \(2016\)](#), [Sadacca et al. \(2016\)](#)), and the coding of taste learning ([Stone \(2005\)](#); [Fontanini and Katz \(2006\)](#); [Grossman et al. \(2008\)](#); [Moran and Katz \(2014a\)](#); [Levitan et al. \(2016\)](#)).

The c-Fos antibody protocol used was adapted from the manufacturer's recommendation ([Santa Cruz Biotech](#)). Slices were rinsed with 1X PBS and incubated in a blocking solution (1XPBS/.3% TritonX-100/3% Bovine serum albumin) for 12 hours at 4°C. Blocking solution was removed and replaced with the primary antibody solution which consists of 1:100 c-Fos polyclonal rabbit IgG (SC-52G; Santa Cruz Biotechnology) for 12 hours at 4°C. After incubation, slices were rinsed using a 1XPBS/.3% Triton X-100 solution followed by the secondary antibody incubation of 1:500 c-Fos Alexa Flour 546 Goat-Anti-Rabbit IgG (H+L) (Life Technologies) and 5% natural goat serum for 12 hours at 4°C. Sections were then rinsed 5-6 times over 90 mins (1XPBS/.3% Triton X-100) and mounted on charged glass slides and

cover slipped with antifade mounting medium with DAPI (Vectashield) to verify that c-Fos expression was specific to the nucleus of GC cells.

To monitor the expression of c-Fos, bilateral GC sections (limited to ventral gustatory cortex (Figure 3.2A and Figure 3.4A)) were viewed by confocal fluorescence microscopy with a Leica Sp5 Spectral confocal microscope/Resonant Scanner with 405 lasers equipped with x/y/z movement stage. Imaging and quantification were performed blind - the experimenter was unaware of the experimental group from which tissue was collected at the time of analysis.

### 3.3.5 Experiment 2

#### Virus Injection Surgery

Rats were anesthetized and prepped as for Experiment 1. With the skull exposed, cleaned, and leveled, bilateral craniotomies were then made at stereotactic coordinates above the part of GC (AP = 1.4 mm, ML = 5 mm from bregma; see Paxinos and Watson (2007)) previously shown to contain neurons that respond distinctly to tastes in awake rats (e.g., Katz et al. (2001); Katz et al. (2002); Maier and Katz (2013a); Sadacca et al. (2016)).

We infused either adeno-associated virus (AAV serotype 9,  $n = 23$ ) coding for ArchT and green fluorescent protein (AAV9-CAG-ArchT-GFP,  $2.5 \times 10^{11}$  particles/mL) or control virus coding for green fluorescent protein alone (AAV9-CAG-GFP,  $n = 13$ ,  $2.5 \times 10^{11}$  particles/mL, UNC vector core) into GC. Thus, we had “ArchT+” and “ArchT-” rats. ArchT has been shown to have better light sensitivity ( $1-10 \text{ mW/mm}^2$ ) than other optogenetic AAV constructs, and thus to be particularly useful for investigations of the neural mechanisms of behavior (Zhang et al. (2010); Yizhar et al. (2011a)). This AAV serotype has been shown to effectively spread and infect all cell types (Aschauer et al. (2013)), and to be effective within GC (Maier et al. (2015); Li et al. (2016)).

Viral particles were suspended in a phosphate buffered solution containing Oregon Green 488 (Invitrogen). To infect GC, micropipettes (tip diameter  $10-20\mu\text{m}$ ) carrying this solution were lowered to a sequence of three depths (4.9, 4.7, and 4.5 mm from dura), at each of which virus was delivered in discrete pulses (50 nl/pulse, 7 seconds between each pulse) controlled by an automatic Nanoject III Microinjector (Drummond Scientific). Following each unilateral set of injections, micropipettes remained in place for 5 min, and were then smoothly removed over the course of one minute so that fluid would not spread back up the cannula track.

## CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

A total volume of 1.25  $\mu\text{L}$  of virus was delivered in 25 pulses per each injection depth. Following bilateral injections wounds were sutured and rats entered post-operative care. To ensure high expression, all viral injections detailed here were made 4 weeks prior to any further procedures. During these 4 weeks rats remained in their home cages, with no experimental manipulation.

### Optical Fiber and Intra-Oral Cannulation surgery

Four weeks after viral injection surgeries, rats were again anesthetized and prepped for surgery. Following the insertion of four self-tapping support screws, the bilateral GC craniotomies created during the previous surgery were re-opened. Custom-built optical fiber assemblies (multimode fiber, 0.22 numerical aperture, 200  $\mu\text{m}$  core, inserted through a 2.5mm multimode stainless-steel ferrule; THORLABS) were lowered to 4.7 mm ventral to the reflected dura mater (targeting the center of virus expression).

Placement of the fibers was stabilized with dental acrylic. Once the dental acrylic was dry, rats were removed from the stereotax and laid prone, and IOCs were implanted bilaterally (see above). The rigid dental acrylic head cap included the optical fibers, IOCs and skull screws. Rats were given post-operative care as detailed above, as well as 7 days of recovery prior to the start of experimentation.

### Adaptation sessions

Following recovery, rats were given 2 days of adaptation, as *per* Experiment 1.

### TPE/WPE sessions

Following adaptation sessions, rats were given either TPE or WPE sessions (1/day) for 3 days, as in Experiment 1. For Experiment 2, however, fluid deliveries were accompanied by optical illumination using 532nm laser light (Shanghai Dream Lasers), coupled to optical fibers (multimode fiber, 200  $\mu\text{m}$  diameter, 0.22 NA) using customized FC/PC patch cables (THORLABS). For every taste delivery, the laser was turned on 850ms before a solenoid valve opened to release a taste onto the tongue and turned off 2500ms later (see [Figure 3.5](#); [Mukherjee et al. \(2017\)](#)). Precisely this same protocol was run for 4 groups of both ArchT+ and ArchT- rats (see below), ensuring that: 1) GC activity was perturbed during the period in which we know taste processing to occur post-delivery ([Katz et al. \(2001\)](#); [Katz et al.](#)

(2002)); 2) GC was intact during the 15 second interval between all taste deliveries, as well as before/after each session; and 3) even control rats received GC virus infection and taste-coincident laser illumination.

The power of illumination was adjusted, before implantation, to be 40mW at the tip of the fiber using an Optical Power and Energy Meter console (THORLABS); this intensity has been calculated to inactivate neurons in vivo within an approximate 1mm sphere around the tip of the optic fiber (Han et al. (2011); Yizhar et al. (2011b)) - a sphere that encompasses about 33% of GC in the caudal-rostral axis (Kosar et al. (1986); Maier et al. (2015); Li et al. (2016)). The same parameters have previously been shown to reduce the activity of ArchT+ single cortical neurons with minimal latency and damage (Maier et al. (2015); Li et al. (2016)); pilot experiments (Figure 3.5) revealed that the same is true in our rats.

In all, 5 groups were run ( $n = 53$ ). Three of these were ArchT- groups—two TPE, and one WPE (allowing us to replicate the TPE effect in infected, ArchT- rats, and to evaluate consumption levels in sham-trained TPE rats). The remaining two were ArchT+ groups: 1) a group (the most important group) that allowed us to test the impact of GC inactivation during TPE; and 2) a group that received unperturbed TPE, and then, six hours after the taste session, received “delayed” but otherwise identical protocol of GC inhibition (allowing us to test the impact of GC inactivation lacking taste stimuli).

### Conditioning session

Conditioning sessions for Experiment 2A involved presentation of sucrose CS *via* IOC only (100 deliveries, 5mL total); this control ensure substantial and reliable sucrose consumption, such that otherwise-possible between-rat differences in amounts taste consumption could not confound the results (Figure 3.6). We have previously shown that both the Experiment 1 and 2 conditioning methods (i.e., Bottle + IOC; IOC-only) do cause taste aversions, although these aversions are somewhat stronger (as expected) with bottle delivery (Flores et al. (2016)). Sucrose consumption was immediately followed by subcutaneous injections of lithium chloride (LiCl, 0.3 M, 0.5% of current body weight) to induce the malaise unconditioned stimulus. One group received sham training (sucrose + saline), and all others received CTA-inducing sucrose + LiCl. In Experiment 2B, which explicitly examined the impact of cortical inactivation TPE on learning-related c-Fos expression, conditioning sessions were identical to those described for Experiment 1. No laser illumination was delivered during

### CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

any conditioning session in Experiment 2A or B - that is, these experiments examined the impact of inhibiting GC activity during TPE sessions on aspects of conditioning delivered 1 full day after the end of the neural inhibition itself.

#### **Testing Session (Experiment 2A)**

Testing sessions were, like conditioning sessions, given without optical illumination. Rats were first presented with a bottle containing 5 mL sucrose (S) for 5 minutes; following a 5-minute pause in which no bottle was available, rats were then presented with a different bottle containing 5 mL water (W), again for 5 minutes.

Following test day, rats were deeply anesthetized with an overdose of the Ketamine/Xylazine mix and perfused as above. Brain tissue was harvested for localization of optical fiber placement.

#### **c-Fos immunohistochemistry (Experiment 2B)**

Rats used for Experiment 2B did not receive testing sessions. Instead, they were perfused for c-Fos analysis 90 minutes after the US administration in the conditioning session. Collection and analysis of immunohistochemistry proceeded as for Experiment 1.

#### **GFP expression**

To monitor expression of the AAV infection via visualization of GFP, slices were stained using previously developed protocol ([Maier et al. \(2015\)](#); [Li et al. \(2016\)](#)). All slices were rinsed 3 times with 1XPBS over 15 mins. Slices were then permeabilized in a 0.3% Triton X-100/1% normal Donkey serum/1XPBS blocking solution for 2 hours at room temperature. Blocking solution was removed and replaced with primary antibody solution which consists of 1:500 anti-green fluorescent protein – rabbit IgG fraction (Life Technologies) for 12 hours at 4°C. After incubation, slices were rinsed 3 times over 15 mins using a 1XPBS followed by the secondary antibody incubation of 1:200 Alexa Fluor 488 donkey anti-rabbit IgG (H+L) (Life Technologies) for 12 hours at 4°C. Sections were then rinsed 3 times over 15 mins using 1XPBS. Slices were then mounted on charged glass slides and cover slipped with Fluoromount Aqueous Mounting Medium.

Slices were imaged with a Keyence fluorescent microscope to confirm successful virus infection and optical fiber location after each animal. Although spread of virus does not

itself define the area of perturbation, fluorescence confirms that the region of interest was infected and more importantly, that the optical fiber tip was localized within that infected region. Rats ( $n = 5$ ) that did not have confirmed bi-lateral virus infection and/or incorrect position of the optical fiber were excluded from the study.

### 3.3.6 Quantification and Statistical analysis

All results were analyzed using SPSS and MATLAB. Significance across all experiments was defined as  $p < 0.05$ .

#### Experiment 1

**c-Fos quantification and analysis** Measures of c-Fos were collected for 6 independent groups of animals (TPE: Sucrose + LiCl, Sucrose + saline, LiCl alone, WPE: Sucrose + LiCl, Sucrose + saline, LiCl alone). To minimize systematic bias, c-Fos counts were performed blind and semi-automatically, using FiJi (University of Wisconsin-Madison) software ([Schindelin et al. \(2012\)](#)). Using the Analyze Particles function (after first rejecting particles outside the size range 10-infinity  $\mu m^2$  and circularity of  $0_{\text{elongated polygon}} - 1.00_{\text{perfect circle}}$  as non-neural), a manual threshold (0-0.50%) was applied across all samples to differentiate between background and somae. Un-normalized soma counts were averaged across hemispheres and anterior, middle and posterior regions of GC resulting in one soma count per slice/rat. The mean soma counts for each of the 6 groups were used in all statistical analyses.

For Experiment 1, we tested the hypothesis that TPE specifically enhances learning-related c-Fos expression in GC. Analysis began with a one-way ANOVA but centered on planned comparisons between subsets of conditions. This analysis was deemed more appropriate than a two-way ANOVA, because the planned comparisons cut across traditional main and interaction affects: our hypothesis specifically turned on an evaluation of whether CTA-related c-Fos activity when preceded by TPE was greater than that when: 1) CTA was preceded by WPE; 2) presentation of S alone was preceded by TPE; or 3) presentation of LiCl alone was preceded by TPE. Data passed assumptions of homogeneity of variances (Levene's test,  $p = 0.118$ ) and is presented as mean  $\pm$  standard error.

The enhancement of CTA-related c-Fos expression caused by TPE was subject to further analysis, including a two-way repeated measures ANOVA that allowed us to evaluate any potential anatomical specificity of the phenomenon within sub-regions of GC. Again, data

## CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

passed assumptions of homogeneity of variances (Levene's test,  $p > 0.05$  for all groups) and is presented as mean  $\pm$  standard error.

**Consumption analysis** To test the potential impact of TPE on sucrose consumption during conditioning sessions, we performed several analyses (using independent samples t-tests to test for significance of differences) on BAT data—comparing raw consumption of sucrose between TPE and WPE groups, as well as lick rate and initial (1st 3 minutes) lick rate.

### Experiment 2

**Analysis of CTA learning** Sucrose and water consumption (mL) on testing day were used to evaluate learning (less sucrose consumption = stronger learning). Consumption for different groups was compared using a one-way ANOVA followed by Fisher's LSD post hoc analyses. Data passed assumptions of homogeneity of variances (Levene's test,  $p = 0.123$ ) and is presented as mean  $\pm$  standard error.

**Localization of Optical Fiber** We also performed a Pearson's bivariate correlation to determine any relationship between placement of optical fibers (i.e., depth) and test day sucrose consumption: each fiber was given a score of 1-4 (ranging from 1: granular, 2: dysgranular, 3: agranular dorsal and 4: agranular ventral as defined by [Paxinos and Watson \(2007\)](#)) for each hemisphere; these scores were then averaged across each animal for a single Fiber Depth Index where 1 was most dorsal insular cortex and 4 was most ventral insular cortex.

**c-Fos Quantification and Analysis** Slices in Experiment 2B were analyzed for c-Fos expression as in Experiment 1 (see above).

## 3.4 Results

### 3.4.1 Experiment 1: TPE increases CTA-related c-Fos activity in GC

Rats were subjected to the TPE protocol—3 days of pre-exposure to sodium chloride (N, NaCl, salty taste), citric acid (C, sour taste), and distilled water (W) via IOC—and then given 1 day of aversion conditioning to a novel sucrose CS (Figure 3.1); TPE in this protocol was previously shown to enhance the strength of CTA (Flores et al. (2016); see Experiment 2 for a replication and extension of the behavioral phenomenon).

Brains were harvested for immunohistochemistry 90 min after the training trial and processed for c-Fos. The 90-min waiting period allowed c-Fos activity to approach its peak consistent with previous studies (Koh and Bernstein (2005); Uematsu et al. (2014)), and the decision to look at c-Fos caused by the training trial (the only session that was identical for all trained rats) allowed us to specifically compare the number of GC neurons activated by the associative learning situation in TPE and non-TPE rats. c-Fos differences observed after testing sessions would have been difficult to interpret, as they would reflect some combination of: 1) learning differences; 2) retrieval differences; and 3) differences in consumption caused by differential learning (see Methods).

Given that CTA learning has repeatedly been associated with increased levels of c-Fos expression in GC (Koh and Bernstein (2005); Hadamitzky et al. (2015); Soto et al. (2017)), and that TPE enhances learning (Flores et al. (2016)), we hypothesized that TPE would enhance learning-related c-Fos expression in GC. We therefore compared c-Fos in CTA-conditioned TPE rats to that of five controls: 1) rats that underwent the exact same CTA – conditioning procedure, but were pre-exposed only to water (i.e., WPE rats); 2-3) rats that received TPE or WPE, but for which the “training trial” was actually a “sham-training” trial in which sucrose was delivered was paired with harmless saline injection; and 4-5) rats that received TPE or WPE followed by “training” that consisted of LiCl alone (Figure 3.1). These experiments allowed to test, most centrally, whether CTA training causes more c-Fos activation when preceded by TPE than when preceded by lack of TPE, and whether CTA training causes more c-Fos activation than either sham-training or LiCl alone when each was preceded by TPE. In addition, examining “conditioning-trial” c-Fos in this set of conditions allowed us to test whether sham-training or LiCl alone causes more c-Fos activation when

preceded by TPE than when not.

Figure 3.2B shows, at two magnifications, representative examples of c-Fos expression in GC (Figure 3.2A, from left to right) for TPE-CTA, WPE-CTA, TPE-Sham and TPE-LiCl (the groups involved in the primary pair of tests above). In the main panels of Figure 3.2B, black spots represent the somae of c-Fos positive neurons revealed by our immunostaining and imaging procedures (the insets confirm that the signal comes from cell bodies rather than noise). It can easily be seen that TPE led to higher levels of CTA-related (i.e., following Sucrose + LiCl) activation than did WPE, and that even fewer neurons were activated by the taste of sucrose or LiCl alone, despite their being preceded by TPE.

Analysis of the group data, which are shown in Figure 3.3A, supports the conclusions suggested by visual scrutiny of the representative data shown in Figure 3.2B. An ANOVA (see methods) revealed significant differences between the groups ( $F(5, 65) = 5.092, p = 0.001$ ); subsequent Fisher's LSD post hoc analyses reveal a specific increase in CTA-related activation of GC caused by TPE—c-Fos expression for this group ( $102 \pm 8$  somae) was higher than the CTA group preceded by WPE ( $78 \pm 9$  somae,  $p = 0.025$ ). In fact, CTA preceded by TPE was found to induce significantly higher c-Fos counts than any of the control groups ( $p$ 's all  $< 0.05$ ). This supports our central hypothesis that TPE enhanced the processing of sucrose-LiCl pairing.

Of course, as noted in the Introduction, CTA would also be strengthened if TPE was to enhance responsiveness to either the taste or LiCl itself (Logue (1979); Franchina and Slank (1988); Flores et al. (2016)), perhaps by enhancing novelty or neophobia. The c-Fos evidence suggests that TPE does not enhance processing of either stimulus presented alone, however: levels of c-Fos expression for TPE-sucrose (+ saline) and TPE-LiCl groups proved indistinguishable from those observed for WPE-sucrose (Fisher's LSD,  $p = 0.596$ ) and WPE-LiCl (Fisher's LSD,  $p = 0.234$ ) groups (Figure 3.3); if TPE had enhanced stimulus novelty, it would have increased c-Fos expression following exposure to sucrose or LiCl (Koh et al. (2003); Lin et al. (2012)).

As an independent test of this last finding, we performed an analysis of licking to novel sucrose in a brief access task (BAT, Med Associates Inc.). The BAT provides a set of particularly rich, sensitive assays of taste responsiveness using individual licks as the basic units of measure (see Methods and DAVIS (1989)); notably, differential lick rates across even small sets of trials reliably reveal differences in novelty (i.e., neophobia; Monk et al. (2013)), including differences caused by IOC pre-exposure as used here (Neath et al. (2010)). Our

### CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

comparison of sucrose licking in TPE and WPE rats revealed no such differences in 15-min raw consumption ( $t(6) = -0.112, p = 0.914$ ) and session lick averages ( $t(6) = -1.098, p = 0.314$ ), or even in lick averages during the initial three minutes of sucrose trials (independent samples t-test,  $t(6) = 1.671, p = 0.146$ , [Figure 3.3B](#)). These results confirm a lack of enhanced neophobia to sucrose caused by TPE. In fact, by the most sensitive lick measure across the first three minutes, TPE trends (insignificantly) toward the less novel (a result consistent with previous suggestions about the impact of innocuous taste experience, [CAPRETTA et al. \(1975\)](#); [Braveman \(1978\)](#); [Miller and Holzman \(1981\)](#); [Franchina and Gilley \(1986\)](#); [Pliner et al. \(1993\)](#); [Morón and Gallo \(2007\)](#)); of course, this lack of significant effect is unsurprising, given that we found no evidence of neophobia to sucrose even in naïve rats (data not shown, but consistent with previous studies observing notably weak neophobia to sucrose, see [Franchina and Dyer \(1985\)](#); [Franchina and Slank \(1988\)](#)). Thus, our c-Fos and behavioral results suggest that TPE changes learning itself, rather than enhancing sucrose novelty.

Further examination of the [Figure 3.3A](#) c-Fos data largely conformed to expectation. A significant difference between training-session c-Fos observed in WPE-CTA and WPE-sham rats (Fisher’s LSD,  $p = 0.044$ ) replicated previous studies ([Navarro et al. \(2000\)](#); [Koh and Bernstein \(2005\)](#); [Wilkins and Bernstein \(2006\)](#); [Andre et al. \(2007\)](#); [Hadamitzky et al. \(2015\)](#)). The difference between the WPE-CTA and WPE-LiCl groups (Fisher’s LSD,  $p = 0.549$ ) failed to achieve significance here, but by far the simplest explanation for this lack of significance, which is of little import in the current study, is the fact that the current protocol was specifically designed to minimize it: the need to avoid a ceiling effect in learning (which would have obscured the enhancement of learning caused by TPE) required that we reduce by half the concentration of LiCl used (see Methods), which in turn ensured that WPE-CTA caused relatively mild conditioning (see Experiment 2 and [Nachman and Ashe \(1973\)](#); [Navarro et al. \(2000\)](#); [Hadamitzky et al. \(2015\)](#); [Flores et al. \(2016\)](#)) that was necessarily difficult to differentiate from the c-Fos induced by a powerful emetic stimulus. Regardless, this result only serves to emphasize the fact that TPE “primes” cortical activation in response to a CTA training trial. This enhancement of neural activation could reasonably be expected to enhance CTA strength.

Finally, we asked one further question with these c-Fos data, examining whether the proposed neural substrate of TPE’s impact on learning might be localized to a specific sub-region of GC. A recent study has demonstrated that lesions large enough to include

posterior GC have a larger impact on CTA learning than those limited to anterior GC (Schier et al. (2015)); to test whether this anatomical subdivision also offered a more precise characterization of our TPE effect, we re-analyzed c-Fos data from TPE and WPE rats, making three separate measurements for each rat - one in an anterior GC slice, one in a middle GC slice, and one in a posterior GC slice (Figure 3.4A; see Methods).

The lack of significance in the interaction term of a repeated measures ANOVA (for pre-exposure condition and sub-region) allows us to reject this ancillary hypothesis, revealing that the effect of TPE on CTA-related c-Fos expression did not differ across sub-regions of GC ( $F(2, 52) = 0.168, p = 0.846$ ). Nor did the main effect for amount of c-Fos in different GC regions reach statistical significance ( $F(2, 52) = 2.380, p = 0.103$ ); only the expected main effect showing that TPE enhances c-Fos was borne out ( $F(1, 26) = 6.783, p = 0.015$ ). TPE influenced CTA-related activation across the entirety of GC similarly (Figure 3.4B).

### 3.4.2 Experiment 2: Perturbation of GC activity mitigates the impact of TPE on CTA

The above experiments demonstrate that TPE enhances GC neural responsiveness to the later association of a novel taste with illness. These results suggest, but do not prove, that GC is a part of the circuit responsible for the enhancement of learning caused by TPE (Flores et al. (2016)); similarly, they suggest but stop short of conclusively demonstrating a link between the observed enhancement of c-Fos and enhanced learning (although this link has been proposed previously, see Koh and Bernstein (2005); Hadamitzky et al. (2015)).

Both hypotheses above can be tested with optogenetic silencing of GC. Note, however, that these tests are quite risky and novel. Specifically, if it is true that GC activity during TPE is vital for the subsequent enhancement of learning and learning-related c-Fos, then inhibiting GC activity during TPE should eliminate these enhancements - enhancements observed a full 24 hours after the last session of inhibition. We are unaware of another experiment in which neural inhibition was predicted to have an impact on either learning or neural activation produced by a procedure administered that much later. Here, we test for each of these possible impacts in turn.

**Experiment 2A (impact of TPE and GCx on learning-related consumption)**

To perform a direct test of our hypothesis that GC neural activity plays an important role in the TPE-induced enhancement of learning, we performed a set of experiments in which we optogenetically perturbed GC neural activity (GCx) during TPE sessions using the light-gated optical silencer ArchT (Archaeorhodopsin-T, a light activated H<sup>+</sup> pump, see [Han et al. \(2011\)](#); [Yizhar et al. \(2011b\)](#)). The construct was delivered in an adeno-associated virus (AAV) also expressing green fluorescent protein fused to ArchT (AAV9-CAG-ArchT-GFP), allowing us to visualize infection sites ([Figure 3.5A](#)).

Every rat run in Experiment 2A and B ( $n = 53$ ; regardless of group) was infected with virus and given fiber optic implants into GC ([Figure 3.5B](#)). Recovery was followed (again, for all rats) by 3 pre-exposure sessions in which GC was illuminated by green (532nm) laser light from 0.85 sec before until 2.5 sec after each fluid delivery, which in GCx rats successfully inhibited GC neuron firing ([Figure 3.5C](#)). The final TPE/WPE day was in turn followed by CTA training and testing sessions ([Figure 3.6](#)).

Rats in three of the 5 groups were injected with a “control virus” carrying only GFP (“ArchT-” groups—GC was illuminated by the laser, but this did not cause GC neurons to be inhibited): 1) an ArchT- group that received TPE followed by CTA training ( $n = 20$ ); 2) an ArchT- group that received WPE followed by CTA training ( $n = 16$ ); and 3) an ArchT- group that received WPE followed by sham conditioning ( $n = 4$ ; [Figure 3.6](#)).

Running these control groups allowed us to replicate the basic, previously-reported ([Flores et al. \(2016\)](#)) effect of TPE on learning in a new context: CTA in surgically-prepared, laser-illuminated ArchT- rats is enhanced (that is, consumption of newly-aversive sucrose in the testing session is reduced) by TPE (left WPE and TPE bars, [Figure 3.7A](#)).

The two additional groups were infected with the AAV containing ArchT (GCx rats, [Figure 3.6](#)) such that laser illumination of GC inhibited neural firing. For one of these groups, GC neural inhibition ( $n = 10$ ) was targeted to the period from 0.85 sec before until 2.5 secs following each intraoral taste infusion of each TPE session (i.e., the same parameters used with ArchT- rats).

These data allowed us to confirm our most central Experiment 2A hypothesis ([Figure 3.7A](#)), specifically that GCx during TPE effectively blocks the expected enhancement of CTA. That is, the strength of conditioning when GC neural activity had been disrupted during TPE ([Figure 3.7A](#), 3rd bar) was less than the strength of conditioning when GC was

### CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

intact during TPE (Figure 3.7A, 2nd bar), and was similar in strength to learning of WPE rats (Figure 3.7A, 1st bar).

While the absolute degree and duration of GCx was quite small (GC activity was suppressed for a total of 5.6 min in each of three 25-min sessions, and for only 3.35 sec at a time), we ran an additional control for the remote confound that GCx itself (as opposed to inhibition of GC during the presence of tastes) might reduce later learning. In this control condition, we induced GCx ( $n = 3$ ; ArchT+) using precisely the same parameters as those described above (3.35 sec GCx every 15 sec, for a total of 5.6 min), but doing so hours after the end of each actual TPE session (see the Methods and Discussions sections for logic explaining why this specific protocol was used, rather than one in which GC was inhibited during TPE sessions but between taste deliveries).

This protocol utterly failed to diminish the TPE-related enhancement of learning (Figure 3.7A, rightmost bar), thus demonstrating that it is specifically GC activity during the TPE session that is important for the learning enhancement.

Statistical analysis of the data shown in Figure 3.7A confirmed each finding described above. A one-way ANOVA revealed that test day sucrose consumption was different across the 5 optogenetic groups ( $F(4, 31) = 13.867, p = 0.000$ ). Fisher's LSD post hoc analyses confirmed replication of the original enhancement of learning caused by TPE, in that learning was significantly enhanced (i.e., test day sucrose consumption was reduced) when CTA was preceded by TPE ( $0.879 \pm 0.212$  mL) compared to WPE ( $1.890 \pm 0.450$  mL; Fisher's LSD  $p = 0.032$ ) for ArchT- rats receiving laser illumination during pre-exposure trials.

More centrally regarding the current inquiry, this TPE-driven learning enhancement was blocked by GCx during TPE: learning was significantly stronger (i.e., test day sucrose consumption was lower) in ArchT- rats in which CTA had been preceded by TPE ( $0.879 \pm 0.212$  mLs) than in identically treated ArchT+ rats ( $2.056 \pm 0.374$  mL, Fisher's LSD  $p = 0.009$ ). Inhibition of GC during TPE specifically impacted later learning, reducing it almost precisely to that observed with WPE (compare the 1st and 3rd bars of Figure 3.7A).

Lastly, our results clearly rule out the confounding possibility that inhibition of GC activity at a time in which tastes are not being presented impacts future processing: test day sucrose consumption was lower in ArchT+ rats for which GCx occurred 6 hours post TPE sessions ( $0.423 \pm 0.084$  mL) than in rats for which GCx occurred during TPE ( $2.056 \pm 0.374$  mL, Fisher's LSD  $p = 0.015$ , Figure 3.7A); it was if anything, even lower than consumption in unperturbed TPE rats (although it is important to note that this difference

failed to reach significance, and therefore is not discussed further).

Rats in all 4 trained groups consumed far less than rats receiving TPE followed by a sham-training (horizontal line in [Figure 3.7A](#),  $4.808 \pm 0.042$  mL; Fisher's LSD  $p = < 0.05$  for all comparisons), proof that our experimental paradigm was able to successfully implement aversions in AAV animals. Together these results strongly imply that GC processing of tastes during TPE is vital for the enhancement of learning normally observed after TPE.

Ancillary analysis failed to reveal anatomical specificity of the GCx effect. While fiber placements were too consistent to allow for an analysis of differences in the anterior-posterior plane, Pearson's correlations (these data were normally distributed) showed no reliable dependency between average depth of fiber placement (see Methods) and raw test day sucrose consumption, within the TPE ( $R = 0.156, p = 0.648$ ), WPE ( $R = 0.405, p = 0.191$ ) or TPE-GCx ( $R = 0.393, p = 0.333$ ) groups ([Figure 3.5B](#); for other groups, N was too small to analyze).

In summary, the above results strongly support the view that GC activity during TPE changes how the brain handles the paired presentation of taste and LiCl in later CTA training. More specifically, they suggest that the elevations in CTA-related cortical c-Fos expression, observed in rats that had received TPE in the preceding days, may truly reflect an enhancement in the expression of learning caused by TPE.

### Experiment 2B (impact of GCx on TPE-induced c-Fos)

Perhaps the best evidence for or against the theory supported above can be had from the testing of one additional, further prediction: if the above logic is correct, then GC perturbation during TPE, which inhibits the normal TPE-induced enhancement of learning, should also reduce the TPE-related enhancement of c-Fos expression.

We directly tested this prediction by comparing CTA-related c-Fos expression in rats that had undergone GCx during TPE to that observed in rats in which GC was unperturbed during TPE. ArchT+ ( $n = 9$ ) and ArchT- ( $n = 8$ ) rats underwent identical surgeries, identical TPE (with laser illumination synced with taste infusions), identical training protocols (pairings of taste and LiCl), and identical harvesting of brains 90 mins after the pairing ([Figure 3.7B](#)). The sole difference between groups was the nature of the virus injected prior to the onset of the protocol, and, thus, the impact of laser illumination ([Figure 3.6](#)).

[Figure 3.7B](#) presents the results of this experiment, which support our prediction. While

overall c-Fos levels were lower in this experiment than in Experiment 1 (a result that likely reflects the trauma of surgery, virus infusion, and ferrule implantation), GCx during TPE reliably reduced levels of c-Fos expressed in response to later CTA training, compared to rats for which TPE was offered with GC intact ( $t(15) = -2.238, p = 0.041$ ). In fact, the reduction in c-Fos expression caused by GCx here is similar to the difference between TPE and WPE in Experiment 1.

Consistent with the findings in Experiment 1, anatomical analysis using two-way repeated measures ANOVA revealed absolutely no dependency of AP zone on this effect — GCx during TPE reduced training-induced c-Fos similarly for anterior, middle and posterior regions of GC. These results take our experiments full circle, demonstrating that GC plays a role in both incidental taste experience processing and taste association learning. Activity in GC during TPE processes the taste experience, allowing later enhancement of GC processing of CTA training, and thereby strengthening learning itself.

### 3.5 Discussion

In the current report, we have addressed the question of how non-reinforced incidental taste experience influences the neural processes underlying later learning. Such investigations are potentially of great importance, both because incidental taste experience is omnipresent in the natural world (and, thus, constantly exerting a noticeable impact on brain and behavior) and because the work reveals limits of the generalizability of data collected from laboratory animals lacking such experience.

Our data reveal that innocuous taste experience specifically affects processing of the association between taste and malaise in gustatory cortex (GC). TPE enhances conditioning-session induction of the immediate-early gene c-Fos (expression that both reflects activity levels post aversion learning and has been shown to be pivotal for the formation of CTA enduring memory; see [Swank and Bernstein \(1994\)](#); [Navarro et al. \(2000\)](#); [Spray et al. \(2000\)](#); [Koh and Bernstein \(2005\)](#); [Wilkins and Bernstein \(2006\)](#); [Yasoshima et al. \(2006\)](#); [Andre et al. \(2007\)](#); [Doron and Rosenblum \(2010\)](#); [Uematsu et al. \(2014\)](#); [Mayford and Reijmers \(2015\)](#)), while having no impact on the processing of a new taste (sucrose, see below) or nausea alone. We also found no significant differences in the expression of c-Fos across the anterior to posterior axis of the structure, suggesting that, while distinct sub-

divisions of GC may contribute differently to CTA memory itself (Schier et al. (2015)), TPE affects learning-related activity uniformly across the entirety of GC.

We went on to show that GC neural activity during TPE plays an indispensable role in enhancing CTA memory. GCx during TPE completely suppressed the TPE-induced enhancement of learning (whereas an identical GCx protocol delivered 6 hours after TPE sessions had no such impact), as well as suppressing TPE-induced enhancements of learning-related c-Fos. These results suggest that GC network activity elicited by innocuous taste experience may impact later learning by specifically priming the association between the novel taste and malaise.

### 3.5.1 Learning-related activity in GC does not suggest TPE-induced changes in sucrose processing

In associative learning (of which CTA is a popular model), memory strength is influenced by the salience (or novelty) of both the CS and US—in the case of CTA, these are the taste (Lubow and Moore (1959); Ahlers and Best (1971); Logue (1979); Franchina and Slank (1988); Siddle and Bond (1988); Rosenblum et al. (1993); Delacasa and Lubow (1995); Merhav and Rosenblum (2008); Clark and Bernstein (2009); Lin et al. (2012)) and malaise (Revusky (1968); Nachman and Ashe (1973); Flores et al. (2016); Levitan et al. (2016))—as well as by the degree of their association (Garcia et al. (1966a); Revusky (1968); Nachman (1970); Ahlers and Best (1971); Kalat and Rozin (1971); Adaikkan and Rosenblum (2015)). Expression of c-Fos in GC is similarly dependent on stimulus salience (Koh et al. (2003)). The impact of TPE on learning, and on c-Fos expression in GC, could therefore conceivably reflect an increase in salience of either stimulus: if TPE increased the novelty of novel sucrose, this change in sensory coding could drive the observed CTA changes, thereby implicating the cholinergic system, manipulations of which are known to be linked to taste novelty (Miranda et al. (2000); Miranda (2003)) and novelty-related CTA changes (Gutiérrez et al. (2003); Clark and Bernstein (2009); Neseliler et al. (2011)), in TPE phenomenology.

Our analysis of c-Fos, consumption, and licking reveal this explanation to be highly unlikely, however. It is well established that administration of novel tastes (including those that cause neophobic reactions) causes higher c-Fos induction than familiar tastes (Koh et al. (2003); Lin et al. (2012)); the fact that neither sucrose nor LiCl exposure-alone caused stronger c-Fos induction in TPE rats than in water-exposed rats suggests that TPE did not

### CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

increase sucrose novelty or LiCl salience. In a separate brief access task experiment known to be particularly sensitive to novelty, TPE rats did not consume more sucrose than WPE rats—analysis of neither a 15-min exposure nor even of the initial 3 min of the 15-min session revealed any evidence of enhanced novelty/neophobia in consumption or licking of TPE rats (or of any sucrose neophobia at all in either group).

Incidentally, while we do not provide direct evidence that TPE is sufficient to attenuate any neophobic response to NaCl and citric acid (this issue does not impact the results presented here—whether these stimuli are novel or familiar, TPE enhances learning of a novel sucrose aversion), we are confident that 3 sessions of TPE are sufficient to attenuate neophobia to salty and sour tastes: in our previous work we observed a weaker CTA in TPE animals who were then conditioned to NaCl, citric acid, and sucrose (Flores et al. (2016)) - evidence suggesting latent inhibition, which is to say familiarity with the TPE tastes; attenuation of neophobia begins within minutes of the first presentation (Monk et al. (2013)) and asymptotes within 6 hours (Green and Parker (1975)), such that many studies investigating the neural underpinnings of neophobia limit testing to only 2 sessions (Gutiérrez et al. (2003); Figueroa-Guzmán and Reilly (2008); Pedroza-Llinás et al. (2009)) — fewer than our 3 sessions of TPE. In fact, the only study that to our knowledge familiarized rats to tastes via IOC reports that far less consumption (1mL for 5 days) than offered here is sufficient to attenuate neophobia (Neath et al. (2010)).

We conclude that TPE specifically affects cortical activity related to the central process involved in associative learning—i.e., the association between the stimuli—rather than the processing of taste or malaise alone. Note that we are not saying that GC is not involved in the processing of sucrose (it almost certainly is); we are saying only that our c-Fos and behavioral data suggest that TPE's impact is largely limited to enhancing the processing of the taste-malaise association—the learning effect is not secondary to a significant enhancement of taste novelty/neophobia. Our work does not therefore provide specific evidence suggesting involvement of the cholinergic system (known to be involved in taste novelty, see Miranda et al. (2000); Miranda (2003)) in TPE.

### 3.5.2 GC plays a specific role in mediating the impact of TPE on CTA learning

Of course, the above results do not prove a specific role for GC activity during TPE, either. The c-Fos data do not directly test whether TPE-induced changes in GC activity are truly important for CTA learning; other CTA-relevant brain regions that are connected to the GC, notably including the basolateral amygdala (Allen et al. (1991); Grossman et al. (2008)), could potentially mediate the effect on c-Fos expression during CTA. To more directly test the importance of GC activity during TPE, we made use of the optical silencer Archaelhodopsin-T (ArchT, see Yizhar et al. (2011b); Maier et al. (2015)), and found that temporally controlled perturbation of GC during TPE reduces aversions to a strength similar to that observed in rats pre-exposed only to water (see Flores et al. (2016)). Neither laser illumination alone nor GCx administered after TPE sessions hindered the TPE-related enhancement of CTA strength. These results strongly suggest that GC plays a role in the integration of TPE into future taste learning.

The fact that GC activity during TPE is important to CTA learning suggests a causal relationship between TPE activity and training-related c-Fos induction. We performed one further set of experiments to test this implication, showing that GC activity during TPE is important for not only TPE-induced enhancement of CTA, but also for the attendant enhancement of GC activation—GCx during TPE resulted in lower GC CTA-related c-Fos expression. Taken in context with the above, these results suggest that GC activity during “innocuous” taste experience promotes CTA learning by specifically priming the degree of association between taste and malaise. These novel findings are important since they implicate GC activity in the processing of incidental taste experience into future associative learning.

We cannot yet say what cortical process is being interrupted by our optogenetic inhibition. The fact that identical amounts of inhibition delivered a few hours after each TPE session had no deleterious impact on learning confirms that the effect likely reflects perturbation of the processing of the tastes themselves and not perturbation of a process lasting hours (Kandel et al. (2014); Levitan et al. (2016)); it remains possible, however, that the seconds to minutes after taste administration are important for a consolidation process that must occur following a taste experience. If so, perhaps inhibition induced immediately after TPE sessions (or a more molecular perturbation of plasticity such as protein synthesis inhibition,

e.g. [Ferreira et al. \(2005\)](#); [Garcia-DeLaTorre et al. \(2009\)](#); [Inberg et al. \(2013\)](#)) or even between taste deliveries would forestall the CTA enhancement observed here. It would be fascinating and valuable to test the time-course of this phenomenon across hours, as has been done in recent work on CTA consolidation (e.g., [Levitan et al. \(2016\)](#)).

These experiments were not done here, however, because their results would be difficult to interpret without a great deal of additional experimentation: if between-trial perturbation proved to have an impact on later learning, for instance, it would be necessary to ask whether that impact reflects a simple need for GC potency during TPE sessions, or whether taste processing itself, which is known to extend well beyond the period of taste delivery ([Yamamoto et al. \(1985\)](#); [Katz et al. \(2001\)](#); [Katz et al. \(2002\)](#)), continues throughout a 30-sec interval. Similarly, an impact of inhibition during water pre-exposure could reflect functioning of a broader process, or could reflect the fact that water is itself a taste ([Rosen et al. 2010](#)), and thus that water pre-exposure is simply a weak, 1-taste TPE ([Flores et al. \(2016\)](#)). Thus, we leave these necessarily large sets of experiments for later investigation.

Although also beyond the scope of our current work, we can begin to speculate regarding possible sources of the TPE effect. It is likely that innocuous taste exposure becomes associated with safe outcome ([Lubow and Moore \(1959\)](#); [Ahlers and Best \(1971\)](#); [Lubow \(1973\)](#); [Lovibond et al. \(1984\)](#); [Delacasa and Lubow \(1995\)](#); [Monk et al. \(2013\)](#)). The multiple sessions of multi-taste pre-exposure that make up TPE would strengthen this association, and perhaps even leave rats with a general belief that “tastes are safe.” The subsequent pairing of a novel taste with an aversive outcome (during CTA) would saliently deviate from this expectation, a fact that would result in enhanced learning. This speculation is consistent with the fact that CTA learning enhancement following TPE is correlated with number of tastes to which the rat is safely pre-exposed ([Flores et al. \(2016\)](#)).

### 3.5.3 Possible cellular mechanisms

At a more mechanistic level, the fact that TPE influences learning that occurs days later suggests a metaplastic mechanism—activity-dependent priming of future cellular plasticity ([Parsons \(2018\)](#)). One study examining this phenomenon showed that experience with one learning paradigm (olfactory learning) enhanced CA1 excitability in a manner that later correlated with enhanced learning of a different learning task (water maze; [Zelcer et al. \(2005\)](#)). It is possible that TPE induces similar metaplastic changes—perhaps at synapses

### CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

linking basolateral amygdala to gustatory cortex (BLA→GC), which are associated with CTA learning and have been shown to undergo metaplastic changes (Grossman et al. (2008); Rodríguez-Durán et al. (2011)). Perhaps GC activity during TPE reduces the threshold for the induction of (BLA→GC) synaptic plasticity in future CTA conditioning sessions, thus enabling the enhancement of CTA learning (most likely via NMDA-dependent mechanisms, see Escobar and Bermúdez-Rattoni (2000); Ferreira et al. (2005)).

Particularly relevant guidance for a future investigation of molecular mechanism comes from work examining varieties of pre-experience and learning in several behavioral tasks including CTA (Ballarini et al. (2009); Moncada et al. (2011)). It has been shown, for instance, that protocols that normally fail to produce lasting CTA memories may be bumped “above threshold” by previous (or, intriguingly, later) exposure to a different taste stimulus that can provide the necessary plasticity related proteins (Merhav and Rosenblum (2008); Ballarini et al. (2009)); while this phenomenon, dubbed “behavioral tagging”, differs in several important ways from the phenomenon that we report here [in addition to using a sub-threshold learning protocol, in “behavioral tagging” the effective pre-exposure is a novel stimulus presented as little as 1 hr prior to training, whereas the TPE effect is stronger following multiple presentations in the days prior to training; see (Flores et al. (2016) for more discussion)] it is intriguing to think that TPE stimuli could induce plasticity related genes (Inberg et al. (2013); Inberg et al. (2016)) that are similar to those identified in this earlier work, could act in the same way as behavioral tagging.

Protein synthesis in GC itself has also been identified as vital for a related phenomenon, whereby latent inhibition learning (the ability of pre-exposure to the eventual CS to reduce CTA strength) is itself strengthened by “pre-pre-exposure” to a different, novel taste (Merhav and Rosenblum (2008)). Again, this phenomenon differs from ours in many ways, but it would be unsurprising if in both paradigms protein synthesis driven by presentation of putative “innocuous” tastes was enhancing the next learning opportunity to arise. Future experiments will pursue the nature—and, just as importantly, the time course—of such protein synthesis.

Regardless of the ultimate underlying mechanism, our findings suggest a causal relationship between TPE activity in GC and the strength of later taste aversion learning-related activity. GC activity during “innocuous” taste experience promotes CTA learning by specifically priming the degree of association between a totally novel taste and malaise. These findings are important in that they grow our currently meager understanding of the neu-

## CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

ral substrates underlying the integration of taste experience with future associative taste aversion learning, a subject with great relevance to human research, and do so in a quite surprising way: they demonstrate that primary sensory cortex, far from being just involved in perception or even learning, is involved in integrating recent un-reinforced experience with current stimulus associations. Our results highlight the need for future research into precisely how seemingly innocuous experience affects future learning at both the behavioral and neural levels.

### 3.6 Figures

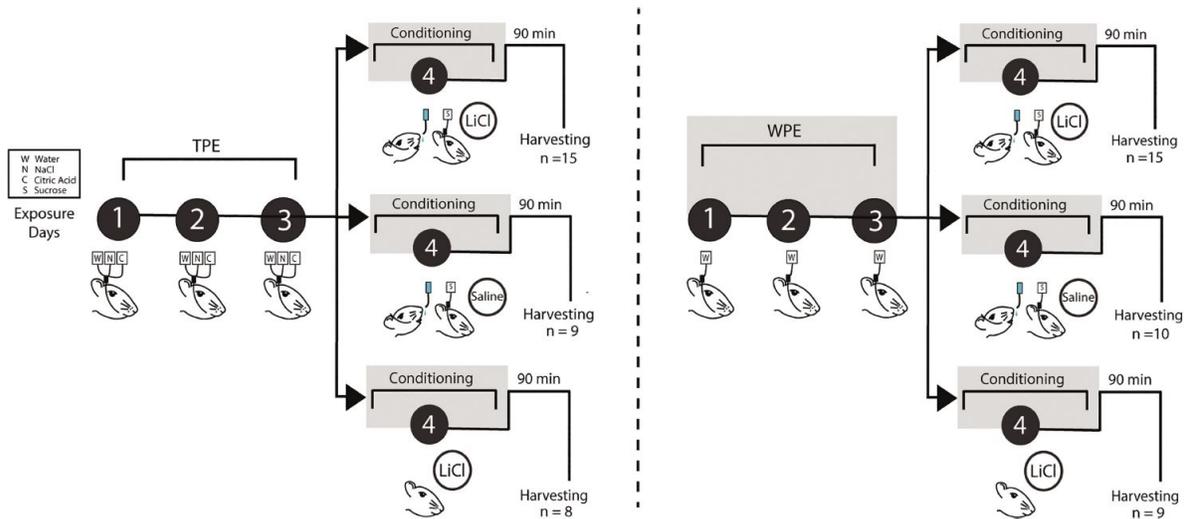


Figure 3.1: **Taste Pre-Exposure paradigm.** A complete timeline of the Taste pre-exposure paradigm showing all groups. Animals were divided into 2 groups: TPE (left) or WPE (right), then further divided into three conditioning conditions: sucrose + LiCl, sucrose + saline, and LiCl alone. Schematic demonstrates the four-day experimental paradigm in which rats receive TPE to water (W), sodium chloride (N) and citric acid (C), via IOC infusions to the tongue for 3d (black circles days 1–3). WPE rats underwent 3 identical days of exposure to water. Aversions were then conditioned on the fourth day, when exposure to sucrose (S) is immediately followed by LiCl injections (0.3 M, 0.5% of current body weight), equal dosages of saline or equal dosages of LiCl without sucrose exposure. Control rats were either given saline injections or LiCl alone. Ninety minutes after the conditioning session, rats were perfused for harvesting of gustatory cortex.

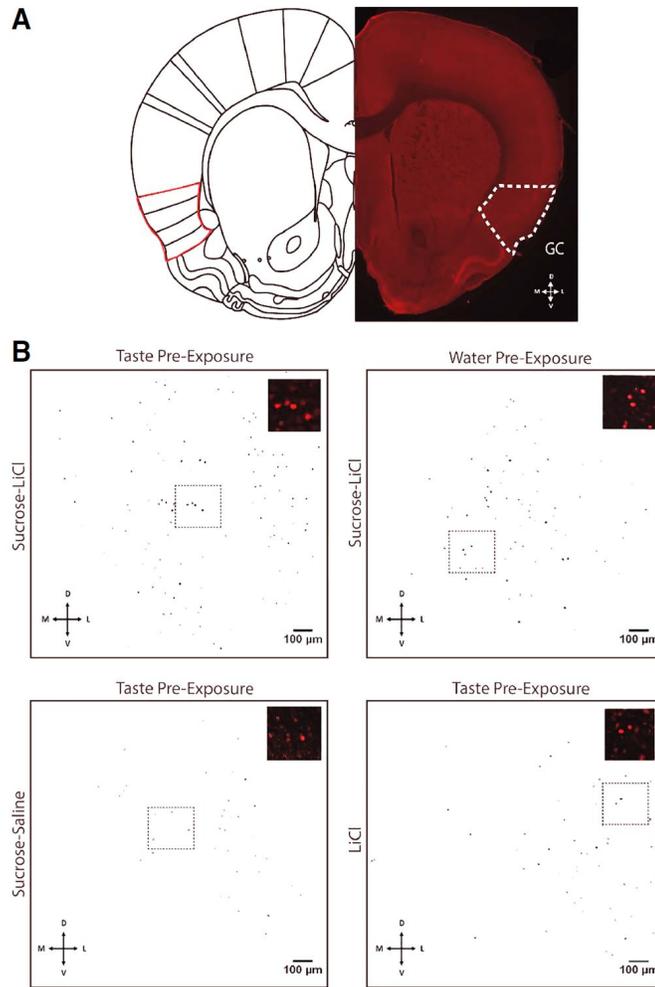


Figure 3.2: **c-Fos positive cells in GC after CTA conditioning to novel sucrose.** **A:** A representative coronal slice indicating the location of gustatory cortex (GC, Left-half hemisphere; Reprinted from Paxinos and Watson (2007) with permission from Elsevier © 2007.) Bottom directional indicates direction of tissue location—dorsal (D), ventral (V), medial (M), and Lateral (L). **B:** Representative images of c-Fos positive somae (masked in black) in GC for the 4 groups most relevant to the central Experiment 1 hypothesis, quantified by the FIJI Analyzing Particles tool. From top left to bottom right: TPE followed by pairing of sucrose and LiCl; WPE followed by pairing of sucrose and LiCl; TPE followed by a pairing of sucrose and saline; and TPE followed by LiCl alone. Insets represent higher-magnification samples of c-Fos positive somae (red) sampled from the region in the dotted black rectangle (note: quantification took place across the entire masked image).

CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

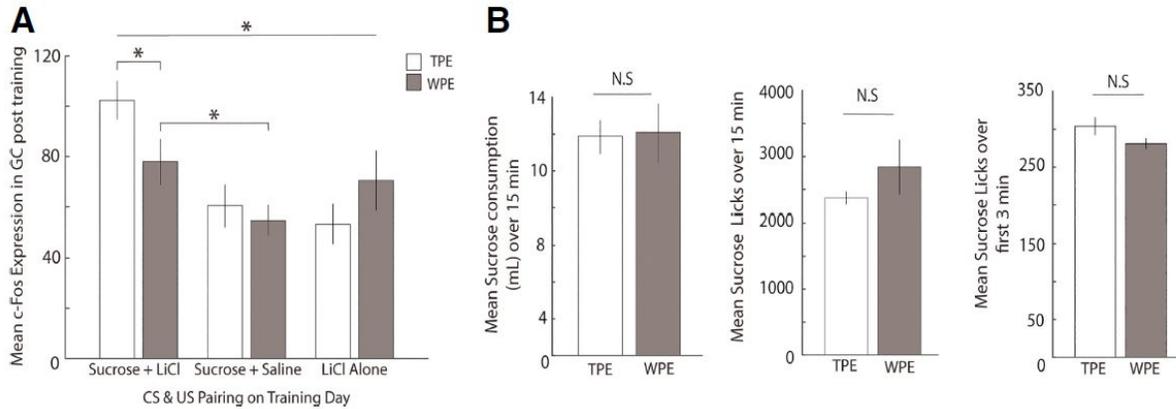


Figure 3.3: **TPE increases CTA-related c-Fos expression in GC** **A:** Pre-exposure to salty and sour tastes (open bars) followed by CTA conditioning resulted in significantly higher c-Fos expression in GC, compared to WPE rats (gray bars); c-Fos in the TPE - CTA conditioning group was also significantly higher than all other groups. These results, and the fact that TPE did not enhance c-Fos in sham and LiCl-alone groups, demonstrate that TPE specifically impacts the pairing of the sucrose with LiCl, and not the processing of either independently. **B: Left:** BAT sucrose consumption (mL) across 15-min access to 15mLs sucrose was similar for TPE and WPE rats. **Middle:** Lick rate for sucrose was similar for TPE and WPE rats during the entire 15min BAT session. **Right:** Initial lick-rate (average licks for the first 3 minutes of the BAT) was similar for TPE and WPE rats. Error bars represent SEM. (\*)  $p < 0.05$ .

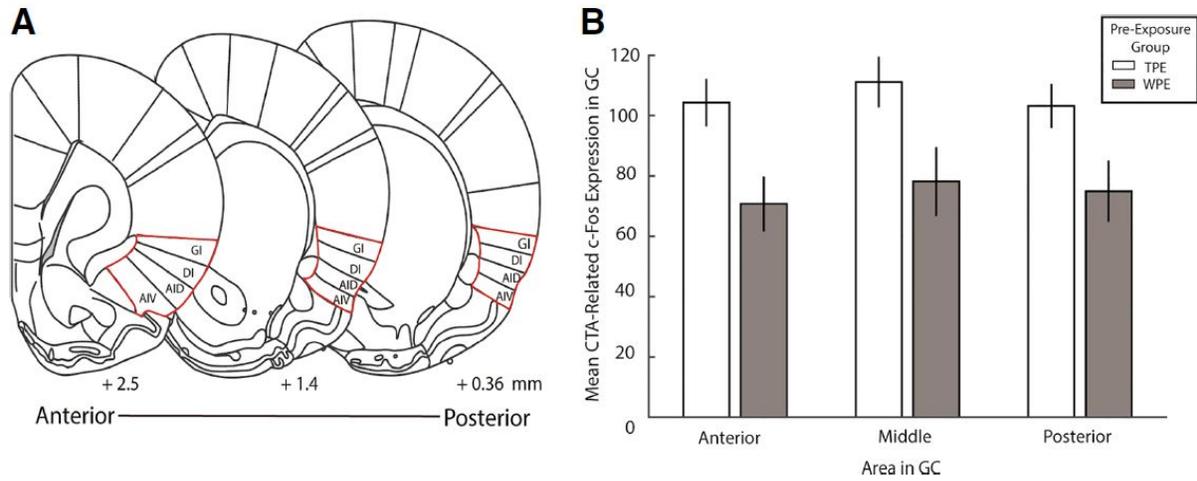


Figure 3.4: **TPE-evoked increases in CTA-related c-Fos expression impacts the entire GC.** **A:** Schematics of coronal sections of anterior, middle and posterior regions of GC (Reprinted from Paxinos and Watson (2007), with permission from Elsevier © 2007). Numbers at the bottom indicate distance from bregma (in mm) for designated anterior, middle and posterior regions of GC. Areas outlined in red indicate the locations of gustatory cortex. **B:** Mean CTA-related c-Fos positive somae in GC corresponding to the regions in Panel A for rats given TPE or WPE. The impact of TPE on CTA-related c-Fos expression was similar for anterior, middle or posterior sub-regions of GC. Error bars represent SEM.

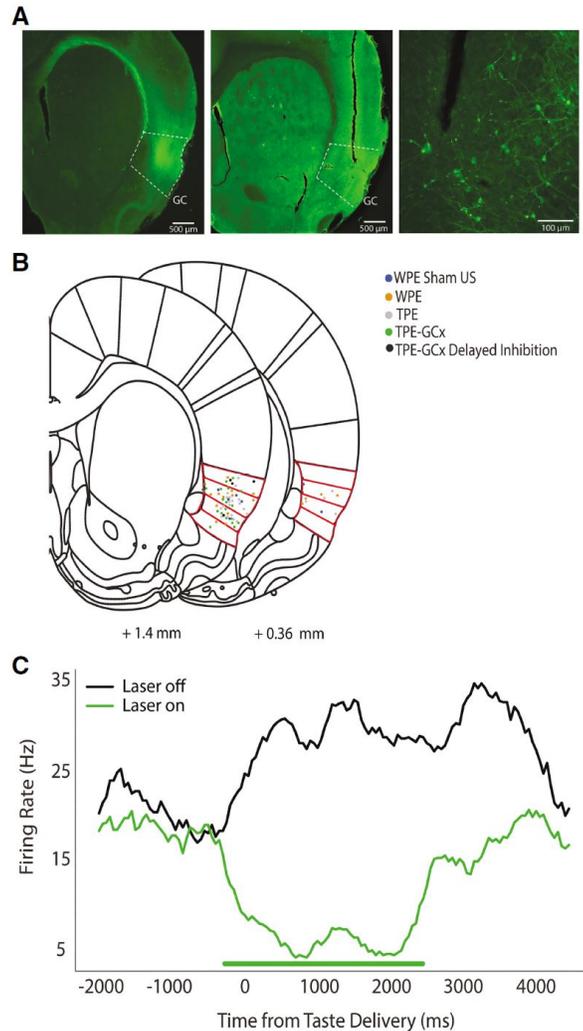


Figure 3.5: **Localization of viral infection and optical fiber placement in GC.** **A:** Representative fluorescent images of gustatory cortex confirming infection of ArchT (left) and placement of optical fiber (middle) in GC. The far-right panel shows, at higher magnification, infected neurons in GC stained with GFP at the tip of the fiber track. **B:** Localization of all fiber tips for all optogenetic groups in Experiment 2A, overlain on schematic coronal slices (Reprinted from Paxinos and Watson (2007), with permission from Elsevier © 2007), demonstrating reliable placement in GC; red outlines are the 4 granular regions in GC. Note that for simplicity and demonstration here, fiber localizations for both hemispheres are overlaid on to one. Each rat received 2 fiber depth scores (one per hemisphere), which were then averaged for a single Fiber Depth index per rat (see Methods). **C.** Peristimulus time histogram (PSTH) of a single GC neuron infected with ArchT demonstrating that our optogenetic inactivation protocol has the desired effect. The firing rate of the infected neuron drops drastically while the laser is on (green line, 0-2500ms post taste delivery); the stimulus response rebounds as soon as the laser is switched off.

CHAPTER 3. EXPERIENCED-EVOKED ENHANCEMENT OF TASTE LEARNING

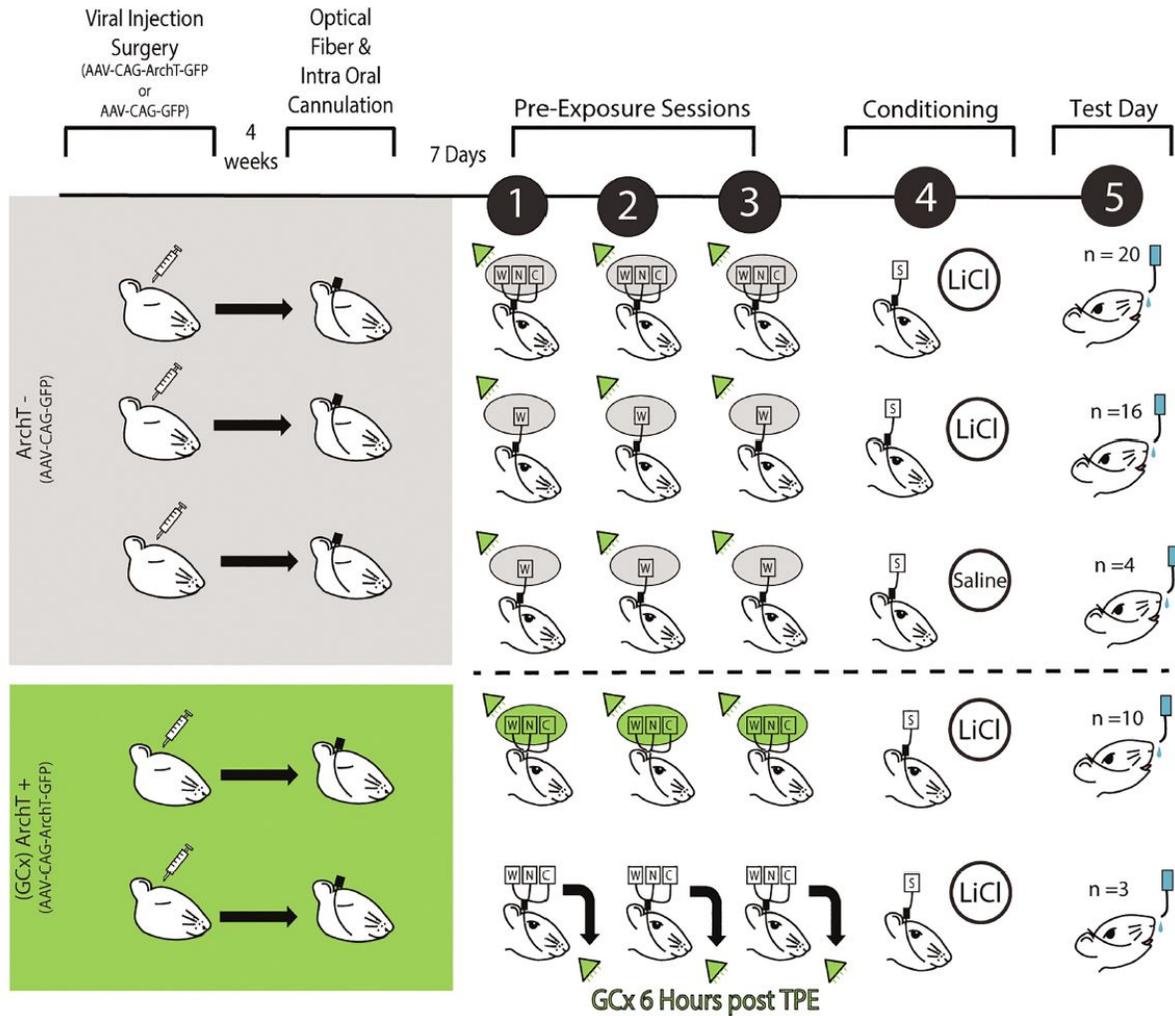


Figure 3.6: **GCx during Taste Pre-Exposure paradigm.** A complete timeline of the optogenetic (Experiment 2A) paradigm showing all groups. Rats first undergo viral injection surgery (either AAV-CAG-GFP or AAV-CAG-ArchT-GFP infused bilaterally into GC). Rats receiving AAV-CAG-ArchT-GFP (ArchT+) are highlighted in green. Rats receiving only AAV-CAG-GFP (ArchT-) are highlighted in gray. To ensure high levels of viral infection, the optical fiber and intra oral cannulation surgery took place 4 weeks after viral injection surgery. Following 7 days of recovery after the optical fiber and intra oral cannulation surgery, all rats encountered 3 TPE (or WPE) sessions, with 532nm laser illumination of GC during each fluid exposure (indicated by green triangle). ArchT- groups (top 3 rows) involved 3 groups: from the top, TPE followed by sucrose + LiCl, WPE followed by sucrose + LiCl and WPE followed by sucrose + saline. ArchT+ groups (bottom 2 rows) involved 2 groups: TPE followed by sucrose + LiCl and TPE followed by sucrose + LiCl in which laser illumination was delayed by 6 hours (see text). On the day following conditioning, aversion strengths are tested via sequential presentation of sucrose and water for all 5 groups.

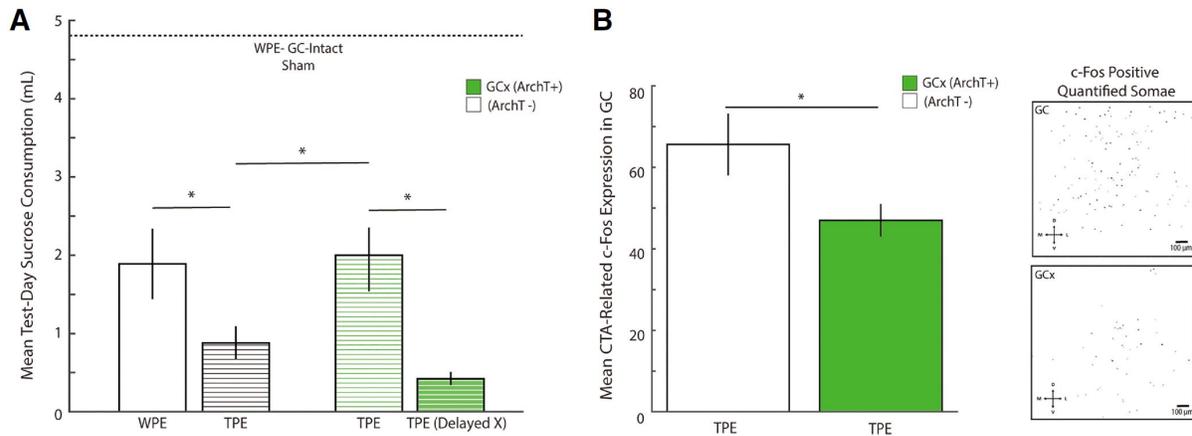


Figure 3.7: **GCx during TPE inhibits CTA learning enhancement and impacts CTA-related c-Fos.** **A:** In ArchT- rats (black and white bars, left side of graph), aversions to novel sucrose were stronger following TPE (black striped bar) than following WPE (open bar) - a replication of the original behavioral result (Flores et al. (2016)). ArchT+ rats receiving GCx during TPE (thin green striped bar) showed significantly weaker CTAs compared to identically run ArchT- rats (compare the middle pair of TPE bars) - demonstration that GC activity during TPE is vital for the behavioral phenomenon. GCx induced 6 hours after each TPE session (thick green striped bar) did not reduce aversion strength. Finally, all conditioned groups showed learning when compared to ArchT- sham-conditioned rats (horizontal dashed line). The x-axis represents average raw sucrose consumption on test day (mL) across all groups. **B: Left:** CTA-related c-Fos expression in GC was significantly stronger in ArchT- rats (open bar), compared to ArchT+ rats (green bar)—GCx reduced CTA-induced c-Fos. **Right:** Representative images of c-Fos positive somae (masked in black) for GC-intact (top) and GCx (bottom) rats, quantified by the FIJI Analyzing Particles tool. Error bars represent SEM. (\*)  $p < 0.05$ .

# Conclusion

The mammalian brain, especially the cerebral cortex, is an incredibly dynamic system. This dynamism is critical for processing the wide variety of sensory inputs from the environment and generating appropriately-timed behavioral responses. Despite the multiple scales of spatio-temporal variability inherent in neural activity, it is commonplace in systems neuroscience to dismiss much of it as “noise”. Standard analysis techniques restrict their focus to the mean response profiles of individual neurons in specific brain regions, leading to the separation of brain regions into processing “modules”. Such trial-averaged analyses imply sensory regions of the cerebral cortex to be “encoding” modules, with computations related to decision-making, learning and behavioral actions happening elsewhere.

In this thesis, we provide several lines of evidence to show that the rodent taste system lacks the spatial separation of processing modules that are thought to characterize the other senses. We show that gustatory cortical ensembles, when analyzed with techniques that do not disregard the temporal dynamics and spatial variability intrinsic to their activity patterns, are revealed to be involved in linking the encoding of taste stimuli to the eventual orofacial behaviors that ingest or expel the taste in the mouth. Furthermore, we find that this connection between sensory encoding, decision-making and behavioral action is dynamic and is moulded by incidental taste experiences which eventually inform future taste-related learning.

## HMMs as models of neural ensemble dynamics

The experiments in [chapter 1](#) and [chapter 2](#) show that HMMs (and related population-level changepoint models) successfully capture the properties of the quasi-stable GC ensemble activity states as well as the suddenness, coherence and speed of the transitions between them. In addition, the control analyses of [Figure 1.5](#) reveal that despite consisting of discrete

## CHAPTER 3. CONCLUSION

activity states by definition, HMMs neither overestimate nor underestimate the sharpness of the sudden transitions between GC ensemble activity states. Using HMMs, we uncover the variability hidden by classical peri-stimulus averaging of neural activity - we find that the seemingly slow (across 1-3s post stimulus) temporal evolution of GC neural firing is actually comprised of relatively long periods of stable ensemble firing punctuated by brief, variably-timed transitions in which the majority of the recorded neurons participate.

Our use of HMMs to model ensemble activity patterns on a trial-by-trial basis is not unique. HMMs have been similarly applied to recordings from frontal areas in monkeys performing Go/No-Go tasks (Abeles et al. (1995), Gat et al. (1997)), spatial location inference from hippocampal place cell assemblies (Maboudi et al. (2018)) and map motor actions from motor cortex ensembles (Kemere et al. (2008)). In each of these cases, the hidden states of the HMMs corresponded to different, behaviorally relevant modes of ensemble firing and could be interpreted as such. We use that strategy to great effect in chapter 1 and find that the GC ensemble activity states discovered by the HMM can be mapped to periods of firing that represent different features of the taste stimulus, culminating in its palatability (on which the eventual decision to ingest or not is based).

From a dynamical systems perspective, despite being easy to interpret, the discrete states of the HMM do not completely elucidate the nature of GC taste responses in state space. In particular, they leave open the possibility that the quasi-stability of GC ensemble states is actually a composition of many (or even a continuous spectrum) of closely similar latent states. As long as latent states within each such group are more similar to each other than across groups, HMMs will tend to cluster them together into common states. Continuous state space modeling techniques (including the Kalman filter, Schiff (2009)) or new variational autoencoder (VAE) based approaches (Pandarinath et al. (2018)) can be used to investigate this possibility.

### Importance of a “moment in time” in sensory cortical processing

The experiments in chapter 2 combine the results from chapter 1 with precisely-timed optogenetic perturbations to show that sensory processing of tastes is casually involved in the appropriate timing of taste-reactive orofacial behaviors. These data show that a unique “moment in time”, internal to the dynamics of sensory ensembles and disconnected from experimentally-imposed time points, is the “tipping point” in the sensory processing of tastes.

### CHAPTER 3. CONCLUSION

On individual trials, perturbations of GC ensembles have a massive impact on the timing of orofacial behavior, but only if delivered before this tipping point.

Internal (to the animal) and behaviorally relevant time points are not uncommon in decision-making tasks that require long periods of pre-training. For instance, it is standard practice to separate out stimulus onset (through a “fixation cue”) and behavioral action (using a “decision cue”) while setting up a behavioral task. Similarly, in spatial exploration tasks (studying, for instance, hippocampal function), it is common to slice and inspect neural firing in relation to behavioral “choice points” - for example, the point in time when the animal makes a decision to move left or right in a spatial maze. Much like the results of [chapter 1](#), trial-averaged analyses reveal neural firing rich in decision-related information when aligned to choice time.

However, the unique moment in time represented by the onset of palatability-related firing in GC ensembles is fundamentally different from these examples - for one, the decision to ingest or expel a taste in the mouth is one that animals need no training to perform. Relatedly, our experiments with naive animals experiencing tastes demonstrate that this internal reference point is a network-level, emergent phenomenon - something that cannot be detected by single neuron analyses. Finally, as [Figure 1.7](#) shows, the onset of taste-reactive behaviors follow the onset of palatability firing by  $\sim 0.3s$ , a much longer brain→behavior lag than that seen in tasks that impose behavioral choice points externally.

The results of [chapter 2](#), in addition to providing support for a dynamical systems-like understanding of sensory processing, argue for a change in the way optogenetic experiments are interpreted. Our optogenetic manipulations affected only  $\sim 60-70\%$  of the neurons in GC, and despite using an activity silencer (ArchT), increased the firing rates of  $10-15\%$  of the recorded neurons (presumably through network-level disinhibition effects, see [Figure 2.2](#)). Despite the variety (and spatial randomness within GC and across rats) of the effects produced by our optogenetic protocol, its behavioral impact was remarkably straightforward - a delay in the shift of sensory processing to discriminatory behavior, contingent upon the state of the brain when the perturbation was delivered. These results point at the need for focusing on time, and not on the exact nature (activation or inhibition), of perturbation when applying optogenetics to large, dynamic networks of neurons. In fact, our data show that the impact of optogenetic perturbations depend not on when they were delivered, but on the internal, network-level activity state of the neurons right when they were perturbed. These results, taken together with our dynamical systems-like interpretation of GC taste

## CHAPTER 3. CONCLUSION

processing, make us hypothesize that the nature of optogenetic perturbation is unimportant in the context of GC - we expect to see similar impacts on the timing of orofacial behaviors if we activated (using channelrhodopsin) instead of inhibiting GC neurons.

### **Implications for taste processing and neocortical function**

The experiments in this thesis reveal a critical role for GC in the entire loop from sensory encoding to decision-making and consumption behaviors. However, as discussed in the Discussion section of [chapter 2](#), GC is just one node of a massively distributed network of taste processing regions in the brain, many of which exhibit temporally dynamic ensemble responses to tastes and have direct descending connections to the brainstem circuits driving orofacial responses. Of these, the interactions between GC and amygdala are known to be particularly important in both the evaluation of taste palatability as well as the establishment of taste aversion learning ([Grossman et al. \(2008\)](#), [Piette et al. \(2012b\)](#), [Sadacca et al. \(2012\)](#)). This makes it likely that the amygdala is one of the key players in the shift of processing control away from GC when palatability-related responses emerge in GC ensembles.

Our results from [chapter 3](#) also reveal that population activity in GC is critical in the modulation of learning by incidental taste experiences. These findings are in line with a relatively new understanding of GC as a focal point for processing the novelty of taste stimuli, and shaping consumption decisions and learning on that basis ([Lin et al. \(2012\)](#)). A different line of work has implicated GC as a key player in multi-sensory integration of tastes (in general ([Vincis and Fontanini \(2016a\)](#)), or with olfactory stimuli in particular([Maier et al. \(2015\)](#), [de Araujo et al. \(2012\)](#))). The results in this thesis, by supporting the dynamic, network-level model of taste processing bring all of these lines of evidence together and paint a unique picture of GC - as a sensory cortex that doesn't just "encode" stimuli. Instead, it maps taste stimuli to behavioral programs, and in particular ropes in contextual information - based on stimulus novelty, other sensory domains and learning/memory - to modify this mapping dynamically.

We think that these findings, while being specific to GC, have lessons to offer for neocortical function in other sensory domains too ([Vincis and Fontanini \(2016b\)](#)). For one, the behavioral relevance of temporal dynamics in ensemble-level taste responses argues strongly against reduction of sensory coding to tuning curves and trial-averaged responses. Secondly, our results question the conclusions about sensory cortical responses to stimuli drawn from

### CHAPTER 3. CONCLUSION

behavioral paradigms that require long periods of training - for instance, decision-related firing is observed even in the visual cortex when untrained subjects are exposed to bistable percepts ([Brascamp et al. \(2015\)](#)). Thirdly, our findings add to the long literature of learning-related plasticity in taste cortex (also see [Berlau and Weinberger \(2008\)](#) and [Lacefield et al. \(2019\)](#) for evidence of learning-induced changes in primary auditory and somatosensory cortices respectively) - classically thought of being limited to sub-cortical/limbic areas in the other senses. Relatedly, the relevance of GC in the processing of incidental taste experiences further supports the idea of stimulus encoding in sensory cortical neurons being shaped by task context. Taken together, our results with GC provide a common computational framework from stimulus encoding to behavioral action that can be applied to and tested in other sensory domains.

# Appendix A

## Python meets systems neuroscience: affordable, scalable and open-source electrophysiology in awake, behaving rodents

### Co-author contributions

The tools and results from this chapter were published in the proceedings of the 2017 Scientific Python (Scipy) conference as [Mukherjee et al. \(2017\)](#).

The following people contributed to this study (and were co-authors on the published paper):

1. **Narendra Mukherjee**: Designed research, performed research, analyzed data, wrote paper, gathered funding.
2. Joseph Wachutka: Designed research, performed research.
3. Donald B. Katz: Wrote paper, gathered funding.

### A.1 Abstract

In-vivo electrophysiology, the recording of neurons in the brains of awake, behaving animals, is currently undergoing paradigm shifts. There is a push towards moving to open-source

technologies that can: 1) be adjusted to specific experiments; 2) be shared with ease; and 3) more affordably record from larger numbers of electrodes simultaneously. Here we describe our construction of a system that satisfies these three desirable properties using the scientific Python stack and Linux. Using a Raspberry Pi to control experimental paradigms, we build a completely open-source, HDF5-based analysis (spike sorting) toolkit in Python. This toolkit can be easily parallelized and scales to incorporate increasing electrode counts and longer recordings. Our rig costs about \$5000, an order of magnitude less than many comparable commercially available electrophysiology systems.

## A.2 Introduction

The process of recording neural activity in awake, behaving animals (in-vivo extracellular electrophysiology, hereafter ‘ephys’) is key in systems neuroscience to understanding how the brain drives complex behaviors. Typically, this process involves voltage recordings from bundles of microwire electrodes (10-20 microns in diameter) surgically implanted into the brain regions of interest. Across hundreds of papers, ephys has increased our understanding of brain systems, function and behavior in a wide range of animal species from invertebrates (locusts and grasshoppers – [Stopfer et al. \(2003\)](#), [Bhavsar et al. \(2015\)](#)) to fishes ([Chou et al. \(2016\)](#)), birds ([Liberti III et al. \(2016\)](#)), rodents ([Jones et al. \(2007\)](#)) and primates ([Ghazanfar et al. \(2005\)](#)). Ephys in awake, behaving animals provides an unprecedented view of the complex and highly variable neural dynamics that underlie accurate behavioral responses. It provides a unique degree of resolution at both the spatial and temporal (sub-millisecond) scales, yielding insights into brain structure and function ranging from the cellular ([Heynen and Bear \(2001\)](#)) to the systems ([Howard et al. \(2014\)](#), [Gurden et al. \(1999\)](#)) levels.

The world of ephys hardware and software has classically been dominated by proprietary and closed-source technologies. These closed-source designs are, by default, not easily modifiable to suit specific experimental circumstances, and, like any closed-source technology, go against the philosophy of open science ([Siegle et al. \(2015\)](#)). It is also harder for other investigators to replicate experimental results obtained by the use of such proprietary software, given that most calculations and operations happen *under-the-hood*, with underlying algorithms either being opaque or not technically accessible to researchers of all skill levels

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

(Ince et al. (2012)). Furthermore, proprietary ephys hardware and software is prohibitively expensive, and poses a high ‘barrier to entry’ for neuroscientists starting to set up their laboratories or working under tight budgetary constraints in smaller schools — particularly those in nations in which research funding is scarce. Finally, the use of closed-source technologies has ensured that ephys hardware and software packages are slow to change. In fact, dominant ephys technology has been virtually unchanged for the last 20 years despite the fact that electronics and computing technology have taken giant strides forward in that time.

With reproducible and affordable science in mind, some ephys laboratories have recently started to explore open source ephys hardware and software (Siegle et al. (2015)). The possible value of this move is manifold: new ephys hardware and software, apart from being open-source, affordable and reproducible, can easily ‘scale’ with growing experiment and data sizes. It is, therefore, much easier with open-source technology to follow the trend in modern ephys towards increasing ‘channel counts’ - recording from hundreds, or even thousands, of electrodes implanted in several different brain regions to better understand the inter-regional coordination that underlies brain function and animal behavior.

In this paper, we describe a completely open-source, Python-based hardware and software setup that we are currently using to study the role of gustatory (taste) cortex in taste-related learning and behavior in rats. We use a Raspberry Pi based system to coordinate the various stimulus control needs of our experiments. This includes the delivery of precise amounts of taste solutions to the animals (Katz et al. (2002)) and the optogenetic perturbation of the firing of neurons in the taste cortex with laser sources (Li et al. (2016), Pastrana (2011)). To handle the ephys signals, we use chips from Intan Technologies and a HDF5 and Python-based software setup for spike sorting (picking out action potentials from individual neurons) (Lewicki (1998)) and analysis.

Starting with a brief description of the hardware we have constructed to control experimental paradigms, we will focus on describing the computations involved at every step of our spike sorting toolchain, highlighting software principles that make such an analysis setup: 1) scale with increased channel counts and longer recordings; and 2) easily parallelized on computing environments. Traditionally, manual approaches, closed-source software and heuristics abound in the electrophysiologist’s spike sorting toolchain - these are time-consuming, error-prone and hard to replicate in a principled manner Wood et al. (2004). We automate several key steps of the spike sorting pipeline with algorithms that have been suggested elsewhere (Quiroga et al. (2004), Fee et al. (1996)) and describe the accessibility and ease-of-use

that the scientific Python stack offers to electrophysiologists. Finally, we demonstrate the use of this system to record and analyze ephys data from 64 electrodes simultaneously in the taste cortex of rodents and point out future directions of improvement keeping the modern ephys experiment in mind.

## A.3 Experimental paradigm and hardware

### A.3.1 Animal care, handling and surgeries

We use adult, female Long-Evans rats (300-325g) and adult mice (15-20g) in our experiments. They are prepared with surgically implanted bundles of microwire electrodes bilaterally in the gustatory (taste) cortex and intra-oral cannulae (IOCs) behind the cheek for delivering taste solutions. All animal care and experiments comply with the Brandeis University Institutional Animal Care and Use Committee (IACUC) guidelines. For more details on experimental protocols, see [Sadacca et al. \(2016\)](#).

### A.3.2 Raspberry Pi based behavior control system

We use a Raspberry Pi running [Ubuntu-MATE](#) to weave together the various behavioral paradigms of our experiments. This includes 1) delivering precise amounts of taste solutions to the animals via pressurized solenoid valves, 2) measuring the animals' licking responses with an analog-to-digital converter (ADC) circuit and 3) controlling laser sources for optogenetic perturbation. Most of these steps involve controlling the digital I/O pins (DIO) of the Pi – the `Rpi.GPIO` package provides convenient functions:

```
1 import RPi.GPIO as GPIO
2 # The BOARD mode allows referring to the GPIO pins
3 # by their number on the board
4 GPIO.setmode(GPIO.BOARD)
5 # Set port 1 as an output
6 GPIO.setup(1, GPIO.OUT)
7 # Send outputs to port 1
8 GPIO.output(1, 1)
9 GPIO.output(1, 0)
```

### A.3.3 Electrode bundles and microdrives

We build electrode bundles with 32 nichrome-formvar microwires (0.0015 inch diameter, from [a-msystems](#)), a 200 $\mu$  fiber for optogenetics (optionally), and 3D printed microdrives. Our custom built drives cost about \$50 and their designs are freely available for use and modification at the [Katz lab website](#).

### A.3.4 Electrophysiology hardware

We use an open-source ephys recording system from [Intan Technologies](#) for neural recordings. The RHD2000 series ephys recording headstages connect to electrode bundles implanted in the animal's brain and contain 32-128 amplifiers and ADCs. The Intan data acquisition system offers an open-source C++ based graphical interface that can record up to 512 electrodes (4 headstages) simultaneously at sampling rates of up to 30kHz/channel. This recording system is relatively robust to AC noise, because the electrode signals are digitized right on the headstage itself, but we additionally encase the animal's behavior and recording chamber in a Faraday cage constructed with standard aluminum insect netting.

## A.4 Electrophysiology in systems neuroscience

In-vivo ephys is unique in systems neuroscience in the temporal and spatial view it provides into the role of the brain in generating accurate behavioral responses. Ephys typically involves the placement of a bundle ([Sadacca et al. \(2016\)](#)) or spatially structured [array](#) ([Wang et al. \(2015\)](#)) of electrodes in a brain region of interest. After the animal recovers from the surgical implantation of electrodes, its behavior in tightly controlled experimental paradigms is correlated with neural activity in the brain region being recorded from. The study of sensory systems (vision, somatosensation, olfaction, taste, etc) in the brain, for instance, involves an awake, behaving animal experiencing different sensory stimuli while ephys recordings are performed in the corresponding sensory cortex (or other involved regions). In addition, ephys electrodes are often implanted in multiple brain regions in the same animal in order to understand the role of inter-regional coordination in the animal's behavior.

In our lab, we study taste processing in adult mice and rats - [Figure A.1](#) shows a typical experimental setup. We surgically implant bundles of 64 microwire electrodes bilaterally (32

wires in each hemisphere) in the taste cortex (among many other regions). Our basic experimental paradigm involves the animal tasting solutions of different kinds (sweet - sucrose, salty - NaCl or bitter - quinine, for instance) while its behavioral responses to the tastes are being recorded (Li et al. (2016)). All this while, we record electrical activity in the taste cortex using the implanted electrodes and eventually try to understand the animals behavior in the light of the activity of the neurons being recorded from.

The essential step in the analysis of ephys data, therefore, is to isolate (and identify) the activity of single neurons from the raw voltage recordings from the implanted electrodes. As shown in Figure A.1, this involves high-pass filtering of the raw voltage signals (see next section for more details) to identify putative action potentials (or ‘*spikes*’). These spikes can originate either from a single neuron or multiple neurons. We thus need to sort them into groups, based on how they are inferred to originate (spikes inferred to be from single neurons are called ‘*single units*’ and those from multiple neurons are called ‘*multi units*’). This entire pipeline is, therefore, called ‘*spike sorting*’. Typically, we are able to isolate 10-40 neurons from our recordings with 64 electrodes - we then go on to correlate the responses of this population of recorded units with the animal’s behavior in our experimental paradigms (see Sadacca et al. (2016), Li et al. (2016) as examples, and Figure A.1).

## A.5 Scientific Python stack for data analysis – spike sorting

The recent push in ephys experiments towards increased channel counts and longer recordings poses significant data handling and analysis challenges. Each of the implanted electrodes needs to be sampled at frequencies in the range of 20-30kHz if it is to clearly render action potentials (the signature binary voltage waveforms, about 1ms in duration, that neurons produce when active – also called ‘*spikes*’, hence the name ‘*spike sorting*’). In our experiments, we sample signals coming from 64 electrodes at 30kHz for up to 2 hours, generating datasets that total 10-30GB in size. Datasets of such sizes cannot be loaded into memory and processed in serial – there is evidently a need to convert the data to a format that allows access to specific parts of the data and can support a parallel computing framework.

The Hierarchical Data Format (HDF5) is ideal for dealing with such big numerical datasets. We use the [Pytables](#) package to build, structure and modify HDF5 files at ev-

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

ery point in our spike sorting and analysis toolchain. Pytables allows data to be stored and extracted from HDF5 files in the convenient form of [numpy arrays](#). We decided to use individual electrodes as storage and computation splits, storing the voltage recording from each electrode as a separate array in the HDF5 file with its analysis assigned to a separate process.

We adopt a semi-supervised approach to spike sorting, starting with a (parallelized) set of automated filtering and clustering steps that can be fine-tuned by the experimenter (who presumably comes equipped with expert knowledge about action potential shapes actually observed in the brain). Our setup therefore involves 3 distinct steps (all the code is available on [Github](#)):

1. Pre-processing (**blech\_clust.py**) – Constructs a HDF5 file post-experiment with the raw binary data recorded by the Intan system, acquires the clustering parameters from the user and creates a shell file that runs the actual processing step in parallel.
2. Processing (**blech\_process.py**) – Runs filtering and clustering steps on the voltage data from every electrode and plots out the results.
3. Post-processing (**blech\_post\_process.py**) – Removes raw recordings from the HDF5 file and compresses it, and then allows the user to sieve out real spikes from the putative spikes plotted in step 2.

### A.5.1 Pre-processing

The pre-processing starts by building a HDF5 file for the ephys dataset with separate nodes for raw neural electrodes, digital inputs and outputs. This structuring of different aspects of the data into separate nodes is a recurrent feature of our toolchain. The Pytables library provides a convenient set of functions for this purpose:

```
1 # modified from blech_clust.py
2 import tables
3 # Create hdf5 file , and make group for raw data
4 hf5 = tables.open_file(hdf5_name[-1]+'.h5', 'w',
5                       title = hdf5_name[-1])
6 # Node for raw electrode data
7 hf5.create_group('/', 'raw')
8 # Node for digital inputs
```

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

```
9 hf5.create_group('/', 'digital_in')
10 #Node for digital outputs
11 hf5.create_group('/', 'digital_out')
12 hf5.close()
```

We have set up Pytables *extendable arrays* (EArrays) to read the electrode and digital input data saved by the Intan system. Extendable arrays are akin to standard Python lists in the sense that their size can be ‘extended’ as data is appended to them – unlike lists, however, they are a homogeneous data class and cannot store different types together. The Intan system saves all the data as integers in binary files and therefore, EArrays of type `int` (defined by `IntAtom` in Pytables) are perfect for this purpose. These EArrays can be constructed and filled as follows:

```
1 # Modified from create_hdf_arrays() in read_file.py
2 # Open HDF5 file with read and write permissions -r+
3 hf5 = tables.open_file(file_name, 'r+')
4 # 2 ports/headstages each with 32
5 # electrodes in our experiments
6 n_electrodes = len(ports)*32
7 # All the data is stored as integers
8 atom = tables.IntAtom()
9 # Create arrays for neural electrodes
10 for i in range(n_electrodes):
11     e1 = hf5.create_earray('/raw',
12                             'electrode%i' % i,
13                             atom, (0,))
14 hf5.close()
15
16 # Modified from read_files() in read_file.py
17 # Open HDF5 file with read and write permissions - r+
18 hf5 = tables.open_file(file_name, 'r+')
19 # Fill data from electrode 1 on port A
20 # Electrode data are stored in binary files
21 # as 16 bit signed integers
22 # Filenames of binary files as defined
23 # by the Intan system
24 data = np.fromfile('amp-A-001.dat',
25                    dtype = np.dtype('int16'))
26 hf5.flush()
```

```
27 hf5.close()
```

To facilitate the spike sorting process, we use the [easygui](#) package to integrate user inputs through a simple graphical interface. Finally, we use GNU [Parallel](#) ([Tange \(2011\)](#)) to run filtering and clustering on every electrode in the dataset in a separate process. GNU Parallel is a great parallelization tool on .nix systems, and allows us to: 1) assign a minimum amount of RAM to every process and 2) resume failed processes by reading from a log file.

## A.5.2 Processing

The voltage data from the electrodes are stored as signed integers in the HDF5 file in the pre-processing step – they need to be converted into actual voltage values (in microvolts) as floats. The datasheet of the [Intan RHD2000](#) system gives the transformation as:

$$voltage(\mu V) = 0.195 * voltage(int) \quad (\text{A.1})$$

Spikes are high frequency events that typically last for 1-1.5 ms – we therefore remove low frequency transients by bandpass filtering the data in 300-3000 Hz using a 2-pole Butterworth filter as follows:

```
1 # Modified from get_filtered_electrode()
2 # in clustering.py
3 from scipy.signal import butter
4 from scipy.signal import filtfilt
5 m, n = butter(2, [300.0/(sampling_rate/2.0),
6                 3000.0/(sampling_rate/2.0)],
7              btype = 'bandpass')
8 filt_el = filtfilt(m, n, el)
```

Depending on the position of the electrode in relation to neurons in the brain, action potentials appear as transiently large positive or negative deflections from the mean voltage detected on the electrode. Spike sorting toolchains thus typically impose an amplitude threshold on the voltage data to detect spikes. In our case (i.e., cortical neurons recorded extracellularly with microwire electrodes), the wide swath of action potentials appear as negative voltage deflections from the average – we therefore need to choose segments of the recording that go *below* a predefined threshold. The threshold we define is based on the median of the electrode’s absolute voltage (for details, see [Quiroga et al. \(2004\)](#)):

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

```
1 # Modified from extract_waveforms() in clustering.py
2 m = np.mean(filt_e1)
3 th = 5.0*np.median(np.abs(filt_e1)/0.6745)
4 pos = np.where(filt_e1 <= m - th)[0]
```

We treat each of these segments as a ‘*putative spike*’. We locate the minimum of each segment and slice out 1.5ms (0.5ms before the minimum, 1ms after = 45 samples at 30kHz) of data around it. These segments, having been recorded digitally, are eventually approximations of the actual analog signal with repeated samples. Even at the relatively high sampling rates that we use in our experiments, it is possible that these segments are significantly ‘jittered’ in time and their shapes do not line up exactly at their minima due to sampling approximation. In addition, due to a variety of electrical noise that seeps into such a recording, we pick up a large number of segments that have multiple troughs (or minima) and are unlikely to be action potentials. To deal with these issues, we ‘dejitter’ the set of potential spikes by interpolating their shapes (using `scipy.interpolate.interp1d`), up-sampling them 10-fold using the interpolation, and finally picking just the segments that can be lined up by their unique minimum.

This set of 450-dimensional putative spikes now needs to be sorted into two main groups: one that consists of actual action potentials recorded extracellularly and the other that consists of noise (this is high-frequency noise that slips in despite the filtering and amplitude thresholding steps). In addition, an electrode can record action potentials from multiple neurons - the group consisting of real spikes, therefore, needs to be further sorted into one or more groups depending upon the number of neurons that were recorded on the electrode. We start this process by first splitting up the set of putative spikes into several *clusters* by fitting a Gaussian Mixture Model (GMM) (Lewicki (1998)). GMM is an unsupervised clustering technique that assumes that the data originate from several different groups, each defined by a Gaussian distribution (in our case over the 450 dimensions of the putative spikes). Classifying the clusters that the GMM picks as noise or real spikes is eventually a subjective decision (explained in the post-processing section). The user picks the best solution with their expert knowledge in the manual part of our semi-automated spike sorting toolchain (which is potentially time consuming for recordings with large numbers of electrodes, see *Discussion* for more details).

Each putative spike waveform picked by the procedure above consists of 450 samples after interpolation – there can be more than a million such waveforms in a 2 hour recording

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

from each electrode. Fitting a GMM in such a high dimensional space is both processor time and memory consuming (and can potentially run into the [curse-of-dimensionality](#)). We therefore reduce the dimensionality of the dataset by picking the first 3 components produced through principal component analysis (PCA) ([Bro and Smilde \(2014\)](#)) using the scikit-learn package ([Pedregosa et al. \(2011\)](#)). These principal components, however, are known to depend mostly on the amplitude-induced variance in shapes of recorded action potential waveforms – to address this possibility, we scale each waveform by its energy (modified from [Fee et al. \(1996\)](#)), defined as follows, before performing the PCA:

$$Energy = \frac{1}{n} \sqrt{\sum_{i=1}^{450} X_i^2} \quad (\text{A.2})$$

where  $X_i = i^{th}$  component of the waveform

Finally, we feed in the energy and maximal amplitude of each waveform as features into the GMM in addition to the first 3 principal components. Using scikit-learn’s GMM API, we fit GMMs with cluster numbers varying from 2 to a user-specified maximum number (usually 7 or 8). Each of these models is fit to the data several times (usually 10) and the best fit is chosen according to the Bayesian Information Criterion (BIC) ([Bhat and Kumar \(2010\)](#)).

The clustering results need to be plotted for the user to be able to pick action potentials from the noise in the post-processing step. The most important in these sets of plots are the actual waveforms of the spikes clustered together by the GMM and the distribution of their inter-spike-intervals (ISIs) (more details in the post-processing step). Plotting the waveforms of the putative spikes in every cluster produced by the GMM together, however, is the most memory-expensive step of our toolchain. Each putative spike is 1.5ms (or 45 samples) long, and there can be tens of thousands of spikes in every cluster (see [Figure A.2](#), [Figure A.3](#)). For a 2 hour recording with 64 electrodes, the plotting step with matplotlib ([Hunter \(2007\)](#)) can consume up to 6GB of memory although the PNG files that are saved to disk are only of the order of 100KB. High memory consumption during plotting also limits the possibility of applying this spike sorting framework to recordings that are several hours long – as a potential substitute, we have preliminarily set up a live plotting toolchain using [Bokeh](#) that can be used during the post-processing step. We are currently trying to work out a more memory-efficient plotting framework, and any suggestions to that end are welcome.

### A.5.3 Post-processing

Once the parallelized processing step outlined above is over, we start the post-processing step by first deleting the raw electrode recordings (under the ‘raw’ node) and compressing the HDF5 file using `ptrepack` as follows:

```

1 # Modified from blech_post_process.py
2 hf5.remove_node('/raw', recursive = True)
3 # Use ptrepack with compression level = 9 and
4 # compression library = blosc
5 os.system("ptrepack --chunkshape=auto --propindexes
6           --complevel=9 --complib=blosc " + hdf5_name
7           + " " + hdf5_name[: -3] + "_repacked.h5")

```

The logic of the post-processing step revolves around allowing the user to look at the GMM solutions for the putative spikes from every electrode, pick the solution that best splits the noise and spike clusters, and choose the cluster numbers that corresponds to spikes. The GMM clustering step, being unsupervised in nature, can sometimes put spikes from two (or more) separate neurons (with very similar energy-scaled shapes, but different amplitudes) in the same cluster or split the spikes from a single neuron across several clusters. In addition, the actual action potential waveform observed on an electrode depends on the timing of the activity of the neurons in its vicinity – co-active neurons near an electrode can additively produce spike waveforms that have smaller amplitude and are noisier (called ‘multi’ units) (Figure A.3) than single, isolated neurons (called ‘single’ units, Figure A.2). Therefore, we set up utilities to merge and split clusters in the post-processing step – users can choose to merge clusters when the spikes from a single neuron have been distributed across clusters or split (with a GMM clustering using the same features as in the processing step) a single cluster if it contains spikes from separate neurons.

HDF5, once again, provides a convenient format to store the single and multi units that the user picks from the GMM results. We make a ‘sorted\_units’ node in the file to which units are added in the order that they are picked by the user. In addition, we make a ‘unit\_descriptor’ table that contains metadata about the units that are picked – these metadata are essential in all downstream analyses of the activity of the neurons in the dataset. To set up such a table through Pytables, we first need to create a class describing the datatypes that the columns of the table will hold and then use this class as the description while creating the table.

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

```
1 # Modified from blech_post_process.py
2 # Define a unit_descriptor class to be used
3 # to add things (anything!) about the sorted
4 # units to a pytables table
5 class UnitDescriptor(tables.IsDescription):
6     electrode_number = tables.Int32Col()
7     single_unit = tables.Int32Col()
8     regular_spiking = tables.Int32Col()
9     fast_spiking = tables.Int32Col()
10
11 # Make a table describing the sorted units.
12 # If unit_descriptor already exists, just open it up
13 try:
14     table = hf5.create_table('/', 'unit_descriptor',
15                             description = UnitDescriptor)
16 except Exception:
17     table = hf5.root.unit_descriptor
```

Cortical neurons (including gustatory cortical neurons that we record from in our experiments) fall into two major categories – 1) excitatory pyramidal cells that define cortical layers and have long range connections across brain regions, and 2) inhibitory interneurons that have short range connections. In ephys records, pyramidal cells produce relatively large and slow action potentials at rates ranging from 5-20 Hz (spikes/s) (Figure A.2, top). Interneurons, on the other hand, have much higher spiking rates (usually from 25-50Hz, and sometimes up to 70 Hz) and much faster (and hence, narrower) action potentials (Figure A.2, bottom). Therefore, in the `unit_descriptor` table, we save the type of cortical neuron that the unit corresponds to in addition to the electrode number it was located on and whether its a single unit. In keeping with classical ephys terminology, we refer to putative pyramidal neuron units as ‘regular spiking units (RSU)’ and interneuron units as ‘fast spiking units (FS)’ (McCormick et al. (1985), Hengen et al. (2013)). In addition, anatomically, pyramidal cells are much larger and more abundant than interneurons in cortical regions (Yokota et al. (2011), Adachi et al. (2013), Peng et al. (2017)) – expectedly, in a typical gustatory cortex recording, 60-70% of the units we isolate are RSUs. This classification of units is in no way restrictive – new descriptions can simply be added to the `UnitDescriptor` class to account for recordings in a sub-cortical region that contains a different electrophysiological unit.

Apart from the shape of the spikes (look at Figure A.2, Figure A.3 and Figure A.4 to

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

compare spikes and typical noise) in a cluster, the distribution of their inter-spike-intervals (ISIs) (plotted in the processing step) is another important factor in differentiating single units from multi units or noise. Due to electrochemical constraints, after every action potential, neurons enter a ‘*refractory period*’ - most neurons cannot produce another spike for about 2ms. We, therefore, advise a relatively conservative ISI threshold while classifying single units – in our recordings, we designate a cluster as a single unit only if  $<0.01\%$  ( $<1$  in 10000) spikes fall within 2ms of another spike.

Finally, we consider the possibility that since the processing of the voltage data from each electrode happens independently in a parallelized manner, we might pick up action potentials from the same neuron on different electrodes (if they are positioned close to each other). We, therefore, calculate ‘*similarity*’ between every pair of units in the dataset – this is the percentage of spikes in a unit that are within 1ms of spikes in a different unit. This metric should ideally be very close to 0 for two distinct neurons that are spiking independently – in our datasets, we consider units that have similarity greater than 20% as the same neuron and discard one of them from our downstream analysis. To speed up this analysis, especially for datasets that have 20-40 neurons each with  $>10000$  spikes, we use [Numba](#)’s just-in-time compilation (JIT) feature:

```
1 # Modified from blech_units_distance.py
2 from numba import jit
3 @jit(nogil = True)
4 def unit_distance(this_unit_times, other_unit_times):
5     this_unit_counter = 0
6     other_unit_counter = 0
7     for i in range(len(this_unit_times)):
8         for j in range(len(other_unit_times)):
9             if np.abs(this_unit_times[i]
10                      - other_unit_times[j])
11                <= 1.0:
12                 this_unit_counter += 1
13                 other_unit_counter += 1
14     return this_unit_counter, other_unit_counter
```

## A.6 Discussion

In-vivo extracellular electrophysiology in awake, behaving animals provides a unique spatiotemporal glimpse into the activity of populations of neurons in the brain that underlie the animals' behavioral responses to complex stimuli. Recording, detecting, analyzing and isolating action potentials of single neurons in a brain region in an awake animal poses a variety of technical challenges, both at the hardware and software levels. Rodent and primate electrophysiologists have classically used proprietary hardware and software solutions in their experiments – these closed-source technologies are expensive, not suited to specific experimental contexts and hard to adapt to sharing and collaboration. The push towards open, collaborative and reproducible science has spurred calls for affordable, scalable open-source experimental setups. In this paper, we have outlined a Raspberry Pi and scientific Python-based solution to these technical challenges and described its successful use in electrophysiological and optogenetic experiments in the taste cortex of awake mice and rats. Our setup can scale as data sizes grow with increasingly longer recordings and larger number of electrodes, and costs \$5000 (compared to up to \$100k for a comparable proprietary setup).

Our approach uses the HDF5 data format, which allows us to organize all of the data (and their associated metadata) under specific nodes in the same file. This approach has several advantages over traditional practices of organizing ephys data. Firstly, HDF5 is a widely used cross-platform data format that has convenient APIs in all major programming languages. Secondly, having all the data from an experimental session in the same file (that can be easily compressed – we use `ptrepack` in the post-processing step) makes data sharing and collaboration easier. Thirdly, HDF5 files allow quick access to desired parts of the data during analysis – as a consequence, larger than memory workflows can easily be supported without worrying about the I/O overhead involved. Lastly, in our setup, we splice the storage and processing of the data by individual electrodes – this allows us to run the processing step in parallel on several electrodes together bringing down processing time significantly.

The standard approach of picking units in ephys studies involves arbitrary, user-defined amplitude threshold on spike waveforms during ephys recordings and manually drawing polygons around spikes from a putative unit in principal component (PC) space. This process is very time consuming for the experimenter and is prone to human errors. Our semi-automated approach to spike sorting is faster and more principled than the standard approach - we automate both these steps of the traditional spike sorting toolchain by using

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

an amplitude threshold that depends on the median voltage recorded on an electrode and clustering putative spikes with a Gaussian Mixture Model (GMM). The user’s expertise only enters the process in the last step of our workflow — they label the clusters picked out by the GMM as noise, single unit or multi unit based on the shapes of the spike waveforms and their ISI distributions. As the number of electrodes in an electrophysiological recording is already starting to run into the hundreds and thousands, there is a need to automate this last manual step as well – this can be achieved by fitting supervised classifiers to the units (and their types) picked out manually in a few training datasets. As the waveforms of spikes can depend upon the brain region being recorded from, such an approach would likely have to be applied to every brain region separately.

During the pre-processing step, we restrict our setup to pick only *negative* spikes – those in which the voltage deflection goes *below* a certain threshold. While most extracellular spikes will appear as negative voltage deflections (due to the fact that they are being mostly recorded from outside the axons of neurons), sometimes an electrode, depending on the brain region, ends up being close enough to the cell body of a neuron to record positive spikes. Our pre-processing step requires only trivial modifications to include positive deflections ‘*above*’ a threshold as spikes as well.

The use of the HDF5 format and the ease of supporting larger-than-memory workflows allows our toolchain to scale to longer recordings and increased electrode counts. However, as explained previously, plotting all the spike waveforms in a cluster together during the processing step using matplotlib is a major memory bottleneck in our workflow. We are working on still more efficient workarounds, and have devised a live plotting setup with Bokeh (that plots 50 waveforms at a time) that can be used during post processing instead. In addition, recordings running for several hours (or days) have to account for the change in spike waveforms induced by ‘*electrode drift*’ - the electrode moves around in the fluid medium of the brain with time. The live plotting module is potentially useful in such longer recordings as well – it can be used to look at spikes recorded in small windows of time (30 minutes say) to see if their shapes change with time.

We are currently attempting to fold our Python based ephys analysis setup into the format of a Python package that can be used by electrophysiologists (using the Intan recording system) to analyze their data with ease on a shared computing resource or on personal workstations. We think that using the scientific Python stack will make previously hidden *under-the-hood* spike sorting principles clearer to the average electrophysiologist, and will

make implementing downstream analyses on these data easier.

## A.7 Figures

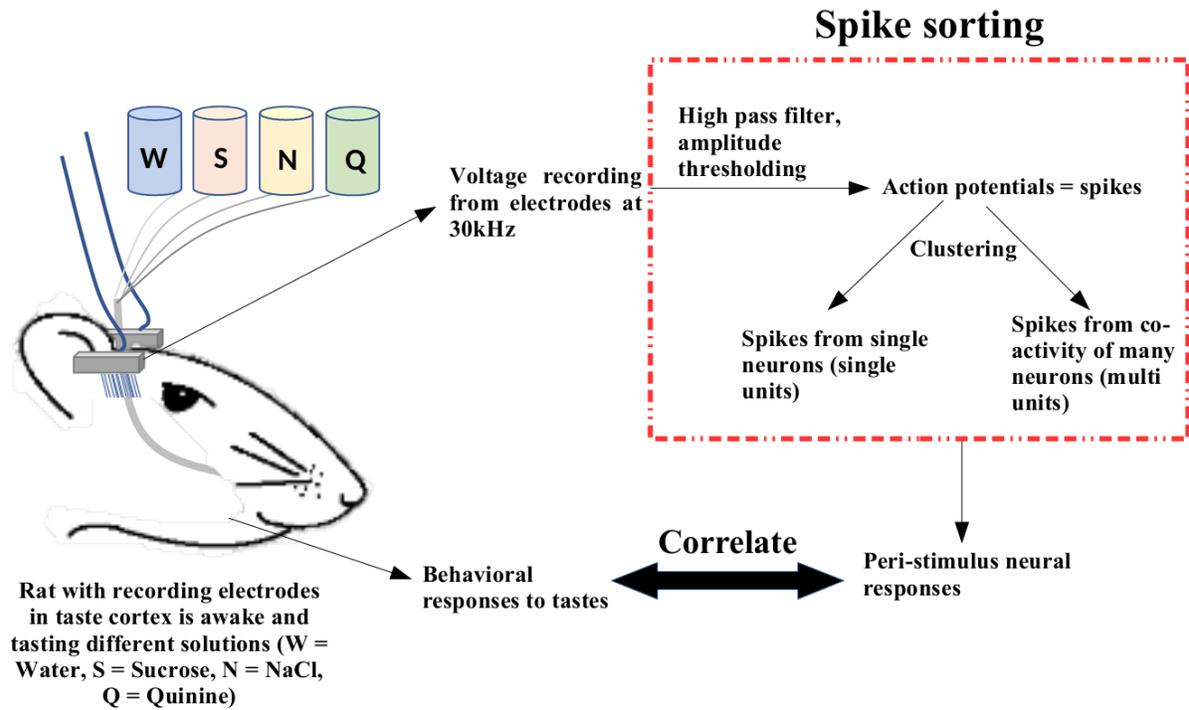


Figure A.1: **An example of a sensory systems experimental setup.** The animal (rodent, primate, etc) experiences sensory stimuli (taste, in this case) while cortical (or other) neurons are being recorded. Eventually, the activity of the recorded population of neurons (also called units) is analyzed in the context of the animal's behavioral responses.

## APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

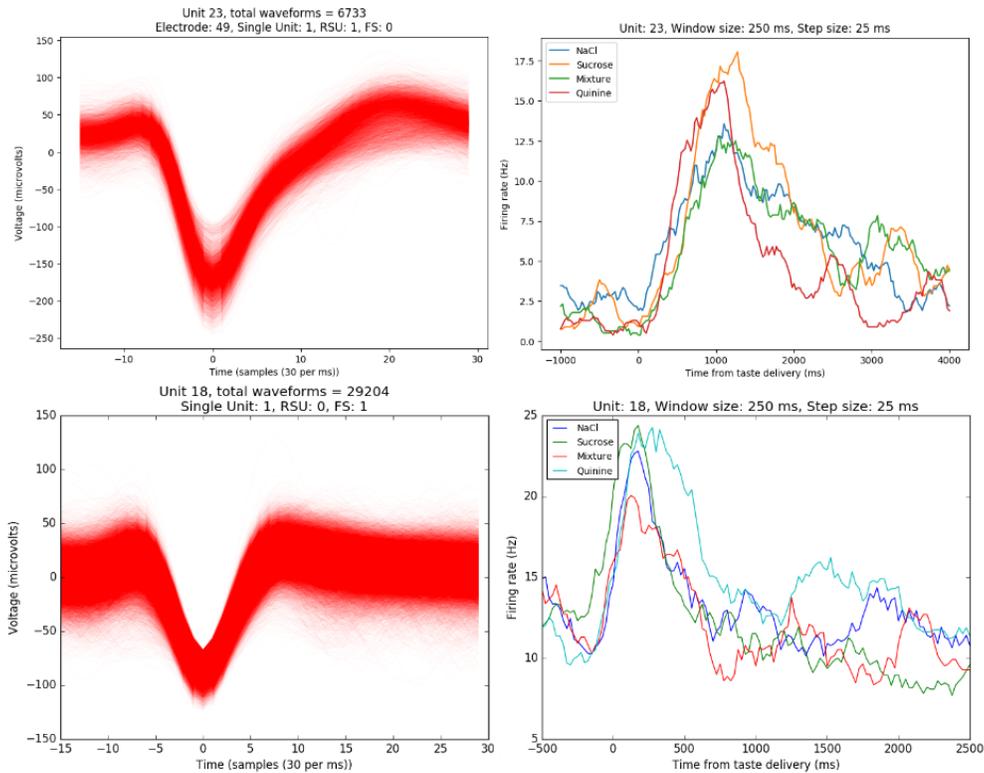


Figure A.2: **Two types of single units isolated from taste cortex recordings.** Spike waveforms on the left, and responses to the taste stimuli on the right. **Top-left:** Spikes waveforms of a regular spiking unit (RSU) - 45 samples (1.5ms) on the time/x axis. Note the 2 inflection points as the spikes go back to baseline from their minimum - this is characteristic of the shape of RSUs. RSUs represent the activity of excitatory cortical pyramidal neurons on ephys records - these spikes are slow and take about 1ms (20-30 samples) to go back up to baseline from their minimum (with 2 inflection points). **Bottom-left:** Spike waveforms of a fast spiking unit (FS) - 45 samples (1.5ms) on the time/x axis. Compare to the spike waveforms of the RSU in the top-left figure and note that this unit has narrower/faster spikes that take only 5-10 samples (1/3 ms) to go back up to baseline from their minimum. FSs represent the activity of (usually inhibitory) cortical interneurons on ephys records. **Top-Right:** Peri-stimulus time histogram (PSTH) - Plot of the activity of the RSU around the time of stimulus (taste) delivery (0 on the time/x axis). Note the dramatic increase in firing rate (spikes/second) that follows taste delivery. **Bottom-Right:** Peri-stimulus time histogram (PSTH) - Plot of the activity of the FS around the time of stimulus (taste) delivery (0 on the time/x axis). Note the dramatic increase in firing rate (spikes/second) that follows taste delivery. Also compare to the PSTH of the RSU in the figure above and note that the FS has a higher firing rate (more spikes) than the RSU. 0.1M Sodium Chloride (NaCl), 0.15M Sucrose, 1mM Quinine-HCl and a 50:50 mixture of 0.1M NaCl and 0.15M Sucrose were used as the taste stimuli.

APPENDIX A. PYTHON MEETS SYSTEMS NEUROSCIENCE

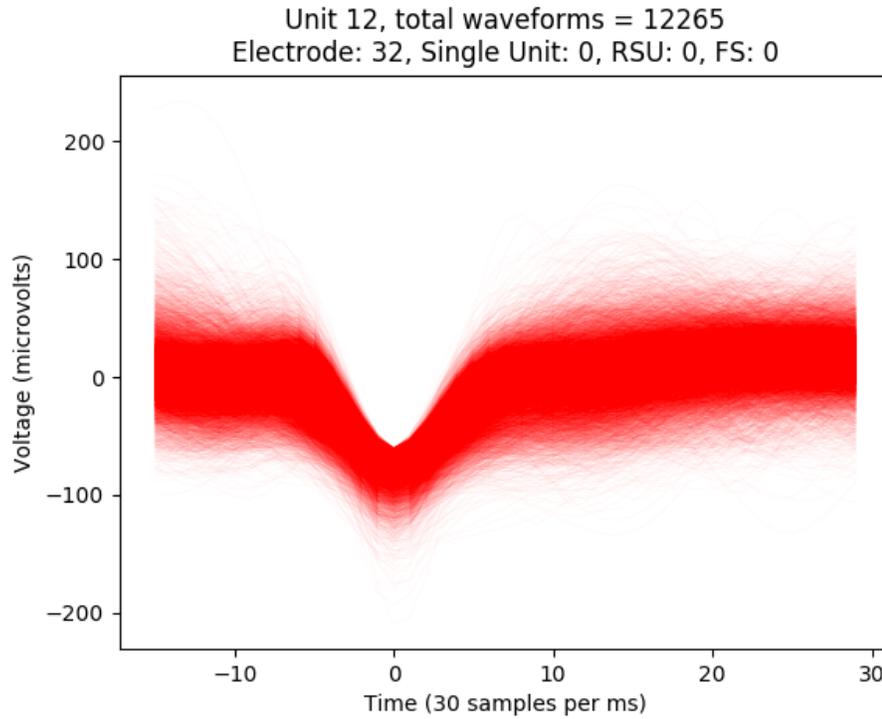


Figure A.3: **A multi unit - 45 samples (1.5ms) on the time/x axis.** Compare to the single units in [Figure A.2](#) and note that these spikes have smaller amplitudes and are noisier. Multi units are produced by the co-activity of multiple neurons near the electrode.

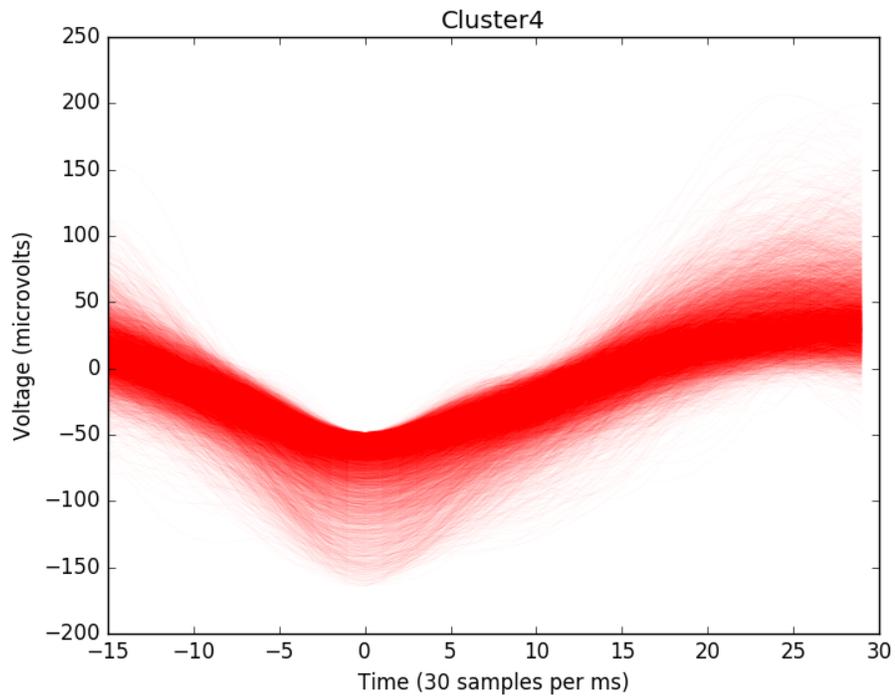


Figure A.4: **A noise cluster - 45 samples (1.5ms) on the time/x axis.** This is high frequency noise that seeps in despite the filtering and thresholding steps used in the processing step. Compare to the single units in [Figure A.2](#) and multi unit in [Figure A.3](#) and note that these waveforms are much smoother and do not have the characteristics of a unit.

# Bibliography

- Moshe Abeles, Hagai Bergman, Itay Gat, Isaac Meilijson, Eyal Seidemann, Naftali Tishby, and Eilon Vaadia. Cortical activity flips among quasi-stationary states. *Proceedings of the National Academy of Sciences*, 92(19):8616–8620, 1995.
- Kazunori Adachi, Satoshi Fujita, Atsushi Yoshida, Hiroshi Sakagami, Noriaki Koshikawa, and Masayuki Kobayashi. Anatomical and electrophysiological mechanisms for asymmetrical excitatory propagation in the rat insular cortex: In vivo optical imaging and whole-cell patch-clamp studies. *Journal of Comparative Neurology*, 521(7):1598–1613, 2013.
- Chinnakkaruppan Adaikkan and Kobi Rosenblum. A molecular mechanism underlying gustatory memory trace for an association in the insular cortex. *Elife*, 4:e07582, 2015.
- Robert H Ahlers and Phillip J Best. Novelty vs temporal contiguity in learned taste aversions. *Psychonomic Science*, 25(1):34–36, 1971.
- Brian D Allen, Annabelle C Singer, and Edward S Boyden. Principles of designing interpretable optogenetic behavior experiments. *Learning & Memory*, 22(4):232–238, 2015.
- Gary V. Allen, Clifford B. Saper, Karen M. Hurley, and David F. Cechetto. Organization of visceral and limbic connections in the insular cortex of the rat. *The Journal of Comparative Neurology*, 311(1):1–16, sep 1991. doi: 10.1002/cne.903110102. URL <https://doi.org/10.1002%2Fcne.903110102>.
- Daniel J Amit. *Modeling brain function: The world of attractor neural networks*. Cambridge university press, 1992.
- Justin St. Andre, Katie Albanos, and Steve Reilly. C-fos expression in the rat brain following

## BIBLIOGRAPHY

- lithium chloride-induced illness. *Brain Research*, 1135:122–128, mar 2007. doi: 10.1016/j.brainres.2006.12.010. URL <https://doi.org/10.1016%2Fj.brainres.2006.12.010>.
- Dominik F Aschauer, Sebastian Kreuz, and Simon Rumpel. Analysis of transduction efficiency, tropism and axonal transport of aav serotypes 1, 2, 5, 6, 8 and 9 in the mouse brain. *PloS one*, 8(9):e76310, 2013.
- Madelyn A Baez-Santiago, Emily E Reid, Anan Moran, Joost X Maier, Yasmin Marrero-Garcia, and Donald B Katz. Dynamic taste responses of parabrachial pontine neurons in awake rats. *Journal of neurophysiology*, 115(3):1314–1323, 2016.
- Amir Bahar, Nimrod Dorfman, and Yadin Dudai. Amygdalar circuits required for either consolidation or extinction of taste aversion memory are not required for reconsolidation. *European Journal of Neuroscience*, 19(4):1115–1118, feb 2004. doi: 10.1111/j.0953-816x.2004.03215.x. URL <https://doi.org/10.1111%2Fj.0953-816x.2004.03215.x>.
- F. Ballarini, D. Moncada, M.C. Martinez, N. Alen, and H. Viola. Behavioral tagging is a general mechanism of long-term memory formation. *Proceedings of the National Academy of Sciences*, 106(34):14599–14604, 2009.
- Peter D Balsam, Michael R Drew, and Cynthia Yang. Timing at the start of associative learning. *Learning and Motivation*, 33(1):141–155, 2002.
- Leonard E Baum, Ted Petrie, George Soules, and Norman Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *The annals of mathematical statistics*, 41(1):164–171, 1970.
- Kasia M Berlau and Norman M Weinberger. Learning strategy determines auditory cortical plasticity. *Neurobiology of learning and memory*, 89(2):153–166, 2008.
- D. E. Berman. Memory extinction, learning anew, and learning the new: Dissociations in the molecular machinery of learning in cortex. *Science*, 291(5512):2417–2419, mar 2001. doi: 10.1126/science.1058165. URL <https://doi.org/10.1126%2Fscience.1058165>.
- Kent C Berridge and Harvey J Grill. Isohedonic tastes support a two-dimensional hypothesis of palatability. *Appetite*, 5(3):221–231, 1984.

## BIBLIOGRAPHY

- Kent C Berridge and Elliot S Valenstein. What psychological process mediates feeding evoked by electrical stimulation of the lateral hypothalamus? *Behavioral neuroscience*, 105(1):3, 1991.
- Harish S Bhat and Nitesh Kumar. On the derivation of the bayesian information criterion. *School of Natural Sciences, University of California*, 2010.
- Mit Balvantray Bhavsar, Ralf Heinrich, and Andreas Stumpner. Multielectrode recordings from auditory neurons in the brain of a small grasshopper. *Journal of Neuroscience Methods*, 256:63 – 73, 2015. ISSN 0165-0270. doi: <https://doi.org/10.1016/j.jneumeth.2015.08.024>. URL <http://www.sciencedirect.com/science/article/pii/S0165027015003155>.
- Christopher M Bishop. *Pattern recognition and machine learning*. Springer-Verlag New York, 2016.
- Rafal Bogacz, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D Cohen. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, 113(4):700, 2006.
- Anil Bollimunta, Douglas Totten, and Jochen Ditterich. Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex. *Journal of Neuroscience*, 32(37):12684–12701, 2012.
- Mark E. Bouton. Context, time, and memory retrieval in the interference paradigms of pavlovian learning. *Psychological Bulletin*, 114(1):80–99, 1993. doi: 10.1037/0033-2909.114.1.80. URL <https://doi.org/10.1037%2F0033-2909.114.1.80>.
- Nicholas E Bowman, Konrad P Kording, and Jay A Gottfried. Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. *Neuron*, 75(5):916–927, 2012.
- Jan Brascamp, Randolph Blake, and Tomas Knäuper. Negligible fronto-parietal bold activity accompanying unreportable switches in bistable perception. *Nature neuroscience*, 18(11):1672, 2015.
- Norman S. Braveman. Preexposure to feeding-related stimuli reduces neophobia. *Animal Learning & Behavior*, 6(4):417–422, dec 1978. doi: 10.3758/bf03209638. URL <https://doi.org/10.3758%2Fbf03209638>.

## BIBLIOGRAPHY

- PA Breslin, JM Kaplan, AC Spector, CM Zambito, and HJ Grill. Lick rate analysis of sodium taste-state combinations. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 264(2):R312–R318, 1993.
- Paul A Breslin, Alan C Spector, and Harvey J Grill. A quantitative comparison of taste reactivity behaviors to sucrose before and after lithium chloride pairings: a unidimensional account of palatability. *Behavioral neuroscience*, 106(5):820, 1992.
- G Larry Bretthorst. *Bayesian spectrum analysis and parameter estimation*, volume 48. Springer Science & Business Media, 2013.
- Rasmus Bro and Age K Smilde. Principal component analysis. *Analytical Methods*, 6(9): 2812–2831, 2014.
- Jan Bures, Federico Bermúdez-Rattoni, and Takashi Yamamoto. *Conditioned taste aversion: Memory of a special kind*. Oxford University Press, 1998.
- PATRICK J. CAPRETTA, J. TIMOTHY PETERSIK, and DANIEL J. STEWART. Acceptance of novel flavours is increased after early experience of diverse tastes. *Nature*, 254(5502):689–691, apr 1975. doi: 10.1038/254689a0. URL <https://doi.org/10.1038/2F254689a0>.
- Kathleen C Chambers. Conditioned taste aversions. *World journal of otorhinolaryngology-head and neck surgery*, 2018.
- Xiaoke Chen, Mariano Gabitto, Yueqing Peng, Nicholas JP Ryba, and Charles S Zuker. A gustotopic map of taste qualities in the mammalian brain. *Science*, 333(6047):1262–1266, 2011.
- Zhixiong Chen and Joseph B Travers. Inactivation of amino acid receptors in medullary reticular formation modulates and suppresses ingestion and rejection responses in the awake rat. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 285(1):R68–R83, 2003.
- Young K Cho, Cheng-Shu Li, and David V Smith. Descending influences from the lateral hypothalamus and amygdala converge onto medullary taste neurons. *Chemical Senses*, 28(2):155–171, 2003.

## BIBLIOGRAPHY

- Ming-Yi Chou, Ryunosuke Amo, Masae Kinoshita, Bor-Wei Cherng, Hideaki Shimazaki, Masakazu Agetsuma, Toshiyuki Shiraki, Tazu Aoki, Mikako Takahoko, Masako Yamazaki, Shin-ichi Higashijima, and Hitoshi Okamoto. Social conflict resolution regulated by two dorsal habenular subregions in zebrafish. *Science*, 352(6281):87–90, 2016. ISSN 0036-8075. doi: 10.1126/science.aac9508. URL <http://science.sciencemag.org/content/352/6281/87>.
- Emily Wilkins Clark and Ilene L. Bernstein. Boosting cholinergic activity in gustatory cortex enhances the salience of a familiar conditioned stimulus in taste aversion learning. *Behavioral Neuroscience*, 123(4):764–771, 2009. doi: 10.1037/a0016398. URL <https://doi.org/10.1037/a0016398>.
- M. Contreras, F. Ceric, and F. Torrealba. Inactivation of the interoceptive insula disrupts drug craving and malaise induced by lithium. *Science*, 318(5850):655–658, oct 2007. doi: 10.1126/science.1145590. URL <https://doi.org/10.1126/science.1145590>.
- Kevin M Crisp and Karen A Mesce. A cephalic projection neuron involved in locomotion is dye coupled to the dopaminergic neural network in the medicinal leech. *Journal of Experimental Biology*, 207(26):4535–4542, 2004.
- Sébastien M Crouzet, Niko A Busch, and Kathrin Ohla. Taste quality decoding parallels taste sensations. *Current Biology*, 25(7):890–896, 2015.
- Andrew M Dacks and Klaudiusz R Weiss. Latent modulation: a basis for non-disruptive promotion of two incompatible behaviors by a single network state. *Journal of Neuroscience*, 33(9):3786–3798, 2013.
- Christina Dalla and Tracey J. Shors. Sex differences in learning processes of classical and operant conditioning. *Physiology & Behavior*, 97(2):229–238, may 2009. doi: 10.1016/j.physbeh.2009.02.035. URL <https://doi.org/10.1016/j.physbeh.2009.02.035>.
- JOHN D. DAVIS. The microstructure of ingestive behavior. *Annals of the New York Academy of Sciences*, 575(1 The Psychobio):106–121, dec 1989. doi: 10.1111/j.1749-6632.1989.tb53236.x. URL <https://doi.org/10.1111/j.1749-6632.1989.tb53236.x>.
- Ivan E de Araujo, Paul Geha, and Dana M Small. Orosensory and homeostatic functions of the insular taste cortex. *Chemosensory perception*, 5(1):64–79, 2012.

## BIBLIOGRAPHY

- Victor de Lafuente and Ranulfo Romo. Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proceedings of the National Academy of Sciences of the United States of America*, 103(39):14266–71, Sep 2006.
- Gustavo Deco, Mar Pérez-Sanagustín, Victor de Lafuente, and Ranulfo Romo. Perceptual detection as a dynamical bistability phenomenon: a neurocomputational correlate of sensation. *Proceedings of the National Academy of Sciences of the United States of America*, 104(50):20073–7, Dec 2007.
- G. Delacasa and R.E. Lubow. Latent inhibition in conditioned taste aversion: The roles of stimulus frequency and duration and the amount of fluid ingested during preexposure. *Neurobiology of Learning and Memory*, 64(2):125–132, sep 1995. doi: 10.1006/nlme.1995.1051. URL <https://doi.org/10.1006/nlme.1995.1051>.
- Alexander J Denman, Joshua D Sammons, Jonathan D Victor, and Patricia M Di Lorenzo. Heterogeneity of neuronal responses in the nucleus of the solitary tract suggests sensorimotor integration in the neural code for taste. *Journal of Neurophysiology*, 2018.
- Aline Desmedt, Shoshi Hazvi, and Yadin Dudai. Differential pattern of cAMP response element-binding protein activation in the rat brain after conditioned aversion as a function of the associative process engaged: Taste versus context association. *The Journal of Neuroscience*, 23(14):6102–6110, jul 2003. doi: 10.1523/jneurosci.23-14-06102.2003. URL <https://doi.org/10.1523/jneurosci.23-14-06102.2003>.
- Patricia M Di Lorenzo. The neural code for taste in the brain stem: response profiles. *Physiology & behavior*, 69(1-2):87–96, 2000.
- Patricia M Di Lorenzo and Christian H Lemon. The neural code for taste in the nucleus of the solitary tract of the rat: effects of adaptation. *Brain research*, 852(2):383–397, 2000.
- Lisa A Dinardo and Joseph B Travers. Hypoglossal neural activity during ingestion and rejection in the awake rat. *Journal of neurophysiology*, 72(3):1181–1191, 1994.
- Flavio Donato, Santiago Belluco Rompani, and Pico Caroni. Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning. *Nature*, 504(7479):272–276, dec 2013. doi: 10.1038/nature12866. URL <https://doi.org/10.1038/nature12866>.

## BIBLIOGRAPHY

- Guy Doron and Kobi Rosenblum. c-fos expression is elevated in GABAergic interneurons of the gustatory cortex following novel taste learning. *Neurobiology of Learning and Memory*, 94(1):21–29, jul 2010. doi: 10.1016/j.nlm.2010.03.003. URL <https://doi.org/10.1016%2Fj.nlm.2010.03.003>.
- Martha L. Escobar and Federico Bermúdez-Rattoni. Long-term potentiation in the insular cortex enhances conditioned taste aversion retention. *Brain Research*, 852(1):208–212, jan 2000. doi: 10.1016/s0006-8993(99)02134-4. URL <https://doi.org/10.1016%2Fs0006-8993%2899%2902134-4>.
- Sean Escola, Alfredo Fontanini, Don Katz, and Liam Paninski. Hidden markov models for the stimulus-response relationships of multistate neural systems. *Neural computation*, 23(5):1071–1132, 2011.
- Michael S Fanselow and Andrew M Poulos. The neuroscience of mammalian associative learning. *Annu. Rev. Psychol.*, 56:207–234, 2005.
- Michale S. Fee, Partha P. Mitra, and David Kleinfeld. Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-gaussian variability. *Journal of Neuroscience Methods*, 69(2):175 – 188, 1996. ISSN 0165-0270. doi: [https://doi.org/10.1016/S0165-0270\(96\)00050-7](https://doi.org/10.1016/S0165-0270(96)00050-7). URL <http://www.sciencedirect.com/science/article/pii/S0165027096000507>.
- G. Ferreira, M.I. Miranda, V. Cruz, C.J. Rodríguez-Ortiz, and F. Bermúdez-Rattoni. Basolateral amygdala glutamatergic activation enhances taste aversion through NMDA receptor activation in the insular cortex. *European Journal of Neuroscience*, 22(10):2596–2604, nov 2005. doi: 10.1111/j.1460-9568.2005.04440.x. URL <https://doi.org/10.1111%2Fj.1460-9568.2005.04440.x>.
- Yazmín Figueroa-Guzmán and Steve Reilly. NMDA receptors in the basolateral amygdala and gustatory neophobia. *Brain Research*, 1210:200–203, may 2008. doi: 10.1016/j.brainres.2008.02.074. URL <https://doi.org/10.1016%2Fj.brainres.2008.02.074>.
- Veronica L. Flores, Anan Moran, Max Bernstein, and Donald B. Katz. Preexposure to salty and sour taste enhances conditioned taste aversion to novel sucrose. *Learning & Memory*,

## BIBLIOGRAPHY

- 23(5):221–228, apr 2016. doi: 10.1101/lm.040360.115. URL <https://doi.org/10.1101/2Flm.040360.115>.
- Veronica L Flores, Tamar Parmet, Narendra Mukherjee, Sacha Nelson, Donald B Katz, and David Levitan. The role of the gustatory cortex in incidental experience-evoked enhancement of later taste learning. *Learning & Memory*, 25(11):587–600, 2018.
- Esmeralda Fonseca, Victor de Lafuente, Sidney A Simon, and Ranier Gutierrez. Sucrose intensity coding and decision-making in rat gustatory cortices. *eLife*, 7:e41152, 2018.
- Alfredo Fontanini and Donald B Katz. State-dependent modulation of time-varying gustatory responses. *Journal of neurophysiology*, 96(6):3183–3193, 2006.
- Alfredo Fontanini, Stephen E Grossman, Joshua A Figueroa, and Donald B Katz. Distinct subtypes of basolateral amygdala taste neurons reflect palatability and reward. *Journal of Neuroscience*, 29(8):2486–2495, 2009.
- Yaihara Fortis-Santiago, Benjamin A Rodwin, Selin Neseliler, Caitlin E Piette, and Donald B Katz. State dependence of olfactory perception as a function of taste cortical inactivation. *Nature neuroscience*, 13(2):158, 2010.
- Joseph J. Franchina and Antoinette B. Dyer. Aversion conditioning and enhanced neophobia: role of test stimuli. *Behavioral and Neural Biology*, 44(1):122–131, jul 1985. doi: 10.1016/s0163-1047(85)91271-3. URL <https://doi.org/10.1016%2Fs0163-1047%2885%2991271-3>.
- Joseph J. Franchina and David W. Gilley. Effects of pretraining on conditioning-enhanced neophobia: Evidence for separable mechanisms of neophobia and aversion conditioning. *Animal Learning & Behavior*, 14(2):155–162, jun 1986. doi: 10.3758/bf03200050. URL <https://doi.org/10.3758%2Fbf03200050>.
- Joseph J. Franchina and Kristine L. Slank. Saliency and the effects of CS preexposure on aversion conditioning. *Behavioral and Neural Biology*, 50(3):367–373, nov 1988. doi: 10.1016/s0163-1047(88)91114-4. URL <https://doi.org/10.1016%2Fs0163-1047%2888%2991114-4>.

## BIBLIOGRAPHY

- Paul W Frankland, Sheena A Josselyn, Stephan G Anagnostaras, Jeffrey H Kogan, Eiki Takahashi, and Alcino J Silva. Consolidation of cs and us representations in associative fear conditioning. *Hippocampus*, 14(5):557–569, 2004.
- Bennett G Galef. Social interaction modifies learned aversions, sodium appetite, and both palatability and handling-time induced dietary preference in rats (*rattus norvegicus*). *Journal of Comparative Psychology*, 100(4):432, 1986.
- Milagros Gallo, Gabriel Roldan, and Jan Bureš. Differential involvement of gustatory insular cortex and amygdala in the acquisition and retrieval of conditioned taste aversion in rats. *Behavioural Brain Research*, 52(1):91–97, nov 1992. doi: 10.1016/s0166-4328(05)80328-6. URL <https://doi.org/10.1016%2Fs0166-4328%2805%2980328-6>.
- John Garcia, Frank R. Ervin, and Robert A. Koelling. Learning with prolonged delay of reinforcement. *Psychonomic Science*, 5(3):121–122, mar 1966a. doi: 10.3758/bf03328311. URL <https://doi.org/10.3758%2Fbf03328311>.
- John Garcia, Frank R Ervin, and Robert A Koelling. Learning with prolonged delay of reinforcement. *Psychonomic Science*, 5(3):121–122, 1966b.
- P. Garcia-DeLaTorre, C. J. Rodriguez-Ortiz, J. L. Arreguin-Martinez, P. Cruz-Castaneda, and F. Bermudez-Rattoni. Simultaneous but not independent anisomycin infusions in insular cortex and amygdala hinder stabilization of taste memory when updated. *Learning & Memory*, 16(9):514–519, aug 2009. doi: 10.1101/lm.1356509. URL <https://doi.org/10.1101%2F1m.1356509>.
- Itay Gat, Naftali Tishby, and Moshe Abeles. Hidden markov modelling of simultaneously recorded cells in the associative cortex of behaving monkeys. *Network: Computation in neural systems*, 8(3):297–322, 1997.
- Andrew Gelman, Kenneth Shirley, et al. Inference from simulations and monitoring convergence. *Handbook of markov chain monte carlo*, pages 163–174, 2011.
- Andrew Gelman, Hal S Stern, John B Carlin, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 2013.

## BIBLIOGRAPHY

- Asif A. Ghazanfar, Joost X. Maier, Kari L. Hoffman, and Nikos K. Logothetis. Multisensory integration of dynamic faces and voices in rhesus monkey auditory cortex. *Journal of Neuroscience*, 25(20):5004–5012, 2005. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.0799-05.2005. URL <http://www.jneurosci.org/content/25/20/5004>.
- Emma Granqvist, Giles ED Oldroyd, and Richard J Morris. Automated bayesian model development for frequency detection in biological time series. *BMC systems biology*, 5(1):97, 2011.
- Emma Granqvist, Matthew Hartley, and Richard J Morris. Basar—a tool in r for frequency detection. *Biosystems*, 110(1):60–63, 2012.
- Kenneth F. Green and Linda A. Parker. Gustatory memory: Incubation and interference. *Behavioral Biology*, 13(3):359–367, mar 1975. doi: 10.1016/s0091-6773(75)91416-9. URL <https://doi.org/10.1016%2Fs0091-6773%2875%2991416-9>.
- Harvey J Grill and Ralph Norgren. The taste reactivity test. i. mimetic responses to gustatory stimuli in neurologically normal rats. *Brain research*, 143(2):263–279, 1978a.
- Harvey J Grill and Ralph Norgren. The taste reactivity test. ii. mimetic responses to gustatory stimuli in chronic thalamic and chronic decerebrate rats. *Brain research*, 143(2):281–297, 1978b.
- Stephen E Grossman, Alfredo Fontanini, Jeffrey S Wieskopf, and Donald B Katz. Learning-related plasticity of temporal coding in simultaneously recorded amygdala–cortical ensembles. *Journal of Neuroscience*, 28(11):2864–2873, 2008.
- H Gurden, J-P Tassin, and TM Jay. Integrity of the mesocortical dopaminergic system is necessary for complete expression of in vivo hippocampal–prefrontal cortex long-term potentiation. *Neuroscience*, 94(4):1019–1027, 1999.
- Ranier Gutiérrez, Carlos J Rodriguez-Ortiz, Vanesa De La Cruz, Luis Núñez-Jaramillo, and Federico Bermudez-Rattoni. Cholinergic dependence of taste memory formation: Evidence of two distinct processes. *Neurobiology of Learning and Memory*, 80(3):323–331, nov 2003. doi: 10.1016/s1074-7427(03)00066-2. URL <https://doi.org/10.1016%2Fs1074-7427%2803%2900066-2>.

## BIBLIOGRAPHY

- M. Hadamitzky, K. Bösche, A. Engler, M. Schedlowski, and H. Engler. Extinction of conditioned taste aversion is related to the aversion strength and associated with c-fos expression in the insular cortex. *Neuroscience*, 303:34–41, sep 2015. doi: 10.1016/j.neuroscience.2015.06.040. URL <https://doi.org/10.1016%2Fj.neuroscience.2015.06.040>.
- B P Halpern and D N Tapper. Taste stimuli: quality coding time. *Science (New York, N.Y.)*, 171(3977):1256–8, Mar 1971.
- Xue Han, Brian Y Chow, Huihui Zhou, Nathan C Klapoetke, Amy Chuong, Reza Rajimehr, Aimei Yang, Michael V Baratta, Jonathan Winkle, Robert Desimone, et al. A high-light sensitivity optical neural silencer: development and application to optogenetic control of non-human primate cortex. *Frontiers in systems neuroscience*, 5:18, 2011.
- Doug P Hanes and Jeffrey D Schall. Neural control of voluntary movement initiation. *Science*, 274(5286):427–430, 1996.
- Alexis Hattox, Ying Li, and Asaf Keller. Serotonin regulates rhythmic whisking. *Neuron*, 39(2):343–352, 2003.
- HR Heekeren, S. Marrett, PA Bandettini, and LG Ungerleider. A general mechanism for perceptual decision-making in the human brain. *Nature*, 431(7010):859–862, 2004.
- Keith B Hengen, Mary E Lambo, Stephen D Van Hooser, Donald B Katz, and Gina G Turrigiano. Firing rate homeostasis in visual cortex of freely behaving rodents. *Neuron*, 80(2):335–342, 2013.
- Arnold J. Heynen and Mark F. Bear. Long-term potentiation of thalamocortical transmission in the adult visual cortex in vivo. *Journal of Neuroscience*, 21(24):9801–9813, 2001. ISSN 0270-6474. URL <http://www.jneurosci.org/content/21/24/9801>.
- Anh Son Ho, Etsuro Hori, Phuong Hong Thi Nguyen, Susumu Urakawa, Takashi Kondoh, Kunio Torii, Taketoshi Ono, and Hisao Nishijo. Hippocampal neuronal responses during signaled licking of gustatory stimuli in different contexts. *Hippocampus*, 21(5):502–519, 2011.
- Matthew D Hoffman and Andrew Gelman. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.

## BIBLIOGRAPHY

- John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- MacKenzie Allen Howard, John LR Rubenstein, and Scott C Baraban. Bidirectional homeostatic plasticity induced by interneuron cell death and transplantation in vivo. *Proceedings of the National Academy of Sciences*, 111(1):492–497, 2014.
- David H Hubel and Torsten N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology*, 160(1):106–154, 1962.
- Alexander C Huk and Michael N Shadlen. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience*, 25(45):10420–10436, 2005.
- J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing In Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.
- Stephen J Huston and Vivek Jayaraman. Studying sensorimotor integration in insects. *Current opinion in neurobiology*, 21(4):527–534, 2011.
- S. Inberg, A. Elkobi, E. Edri, and K. Rosenblum. Taste familiarity is inversely correlated with arc/arg3.1 hemispheric lateralization. *Journal of Neuroscience*, 33(28):11734–11743, jul 2013. doi: 10.1523/jneurosci.0801-13.2013. URL <https://doi.org/10.1523%2Fjneurosci.0801-13.2013>.
- Sharon Inberg, Eyal Jacob, Alina Elkobi, Efrat Edry, Akiva Rappaport, T. Ian Simpson, J. Douglas Armstrong, Noam Shomron, Metsada Pasmanik-Chor, and Kobi Rosenblum. Fluid consumption and taste novelty determines transcription temporal dynamics in the gustatory cortex. *Molecular Brain*, 9(1), feb 2016. doi: 10.1186/s13041-016-0188-4. URL <https://doi.org/10.1186%2Fs13041-016-0188-4>.
- Darrel C Ince, Leslie Hatton, and John Graham-Cumming. The case for open computer programs. *Nature*, 482(7386):485–488, 2012.
- Andrea Insabato, Mario Pannunzi, Edmund T Rolls, and Gustavo Deco. Confidence-related decision making. *Journal of neurophysiology*, 104(1):539–47, Jul 2010. doi: 10.1152/jn.01068.2009.

## BIBLIOGRAPHY

- Ahmad Jezzini, Luca Mazzucato, Giancarlo La Camera, and Alfredo Fontanini. Processing of hedonic and chemosensory features of taste in medial prefrontal and insular networks. *Journal of Neuroscience*, 33(48):18966–18978, 2013.
- Joshua P Johansen, Christopher K Cain, Linnaea E Ostroff, and Joseph E LeDoux. Molecular mechanisms of fear learning and memory. *Cell*, 147(3):509–524, 2011.
- Lauren M Jones, Alfredo Fontanini, Brian F Sadacca, Paul Miller, and Donald B Katz. Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proceedings of the National Academy of Sciences*, 104(47):18772–18777, 2007.
- James W. Kalat and Paul Rozin. Role of interference in taste-aversion learning. *Journal of Comparative and Physiological Psychology*, 77(1):53–58, 1971. doi: 10.1037/h0031585. URL <https://doi.org/10.1037/h0031585>.
- Eric R Kandel, Yadin Dudai, and Mark R Mayford. The molecular and systems biology of memory. *Cell*, 157(1):163–186, 2014.
- R.E. Kass, V. Ventura, and E.N. Brown. Statistical issues in the analysis of neuronal data. *Journal of Neurophysiology*, 94(1):8–25, 2005.
- Donald B Katz and Brian F Sadacca. 6 taste. *Neurobiology of Sensation and Reward*, page 127, 2011.
- Donald B Katz, SA Simon, and Miguel AL Nicolelis. Dynamic and multimodal responses of gustatory cortical neurons in awake rats. *Journal of Neuroscience*, 21(12):4478–4489, 2001.
- Donald B Katz, SA Simon, and Miguel AL Nicolelis. Taste-specific neuronal ensembles in the gustatory cortex of awake rats. *Journal of Neuroscience*, 22(5):1850–1857, 2002.
- Caleb Kemere, Gopal Santhanam, M Yu Byron, Afsheen Afshar, Stephen I Ryu, Teresa H Meng, and Krishna V Shenoy. Detecting neural state transitions using hidden markov models for motor cortical prostheses. *Journal of neurophysiology*, 2008.
- Adam Kepecs, Naoshige Uchida, and Zachary F Mainen. The sniff as a unit of olfactory processing. *Chemical senses*, 31(2):167–179, 2006.

## BIBLIOGRAPHY

- Roозbeh Kiani, Timothy D Hanks, and Michael N Shadlen. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *Journal of Neuroscience*, 28(12):3017–3029, 2008.
- Stephen W Kiefer and Missy R Orr. Taste avoidance, but not aversion, learning in rats lacking gustatory cortex. *Behavioral neuroscience*, 106(1):140, 1992.
- Camille Tessitore King, Koji Hashimoto, Ginger D Blonde, and Alan C Spector. Unconditioned oromotor taste reactivity elicited by sucrose and quinine is unaffected by extensive bilateral damage to the gustatory zone of the insular cortex in rats. *Brain research*, 1599: 9–19, 2015.
- Ming Teng Koh and Ilene L. Bernstein. Mapping conditioned taste aversion associations using c-fos reveals a dynamic role for insular cortex. *Behavioral Neuroscience*, 119(2):388–398, 2005. doi: 10.1037/0735-7044.119.2.388. URL <https://doi.org/10.1037%2F0735-7044.119.2.388>.
- Ming Teng Koh, Emily E. Wilkins, and Ilene L. Bernstein. Novel tastes elevate c-fos expression in the central amygdala and insular cortex: Implication for taste aversion learning. *Behavioral Neuroscience*, 117(6):1416–1422, 2003. doi: 10.1037/0735-7044.117.6.1416. URL <https://doi.org/10.1037%2F0735-7044.117.6.1416>.
- Nicholas Kolodiy, Gary M. Brosvic, David Pak, and Sheryl Loeffler. Taste preference behavior in long-evans rats and egyptian spiny mice. *Bulletin of the Psychonomic Society*, 31(4):307–310, apr 1993. doi: 10.3758/bf03334937. URL <https://doi.org/10.3758%2Fbf03334937>.
- Eva Kosar, Harvey J Grill, and Ralph Norgren. Gustatory cortex in the rat. i. physiological properties and cytoarchitecture. *Brain research*, 379(2):329–341, 1986.
- Alexei A Koulakov, Sridhar Raghavachari, Adam Kepecs, and John E Lisman. Model for a robust neural integrator. *Nature neuroscience*, 5(8):775, 2002.
- John Kounios, Jennifer L Frymiare, Edward M Bowden, Jessica I Fleck, Karuna Subramaniam, Todd B Parrish, and Mark Jung-Beeman. The prepared mind neural activity prior to problem presentation predicts subsequent solution by sudden insight. *Psychological Science*, 17(10):882–890, 2006.

## BIBLIOGRAPHY

- John Kounios, Jessica I Fleck, Deborah L Green, Lisa Payne, Jennifer L Stevenson, Edward M Bowden, and Mark Jung-Beeman. The origins of insight in resting-state brain activity. *Neuropsychologia*, 46(1):281–291, 2008.
- Ikue Kusumoto-Yoshida, Haixin Liu, Billy T Chen, Alfredo Fontanini, and Antonello Bonci. Central role for the insular cortex in mediating conditioned responses to anticipatory cues. *Proceedings of the National Academy of Sciences*, 112(4):1190–1195, 2015.
- Dorothy W. S. Kwok, Evan J. Livesey, and Robert A. Boakes. Serial overshadowing of taste aversion learning by stimuli preceding the target taste. *Learning & Behavior*, 40(4):427–438, jan 2012. doi: 10.3758/s13420-011-0064-0. URL <https://doi.org/10.3758/2Fs13420-011-0064-0>.
- Clay O Lacefield, Eftychios A Pnevmatikakis, Liam Paninski, and Randy M Bruno. Reinforcement learning recruits somata and apical dendrites across layers of primary sensory cortex. *Cell reports*, 26(8):2000–2008, 2019.
- Kenneth W Latimer, Jacob L Yates, Miriam LR Meister, Alexander C Huk, and Jonathan W Pillow. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244):184–187, 2015.
- Vernon Lawhern, Wei Wu, Nicholas Hatsopoulos, and Liam Paninski. Population decoding of motor cortical activity using a generalized linear model with hidden states. *Journal of neuroscience methods*, 189(2):267–80, Jun 2010. doi: 10.1016/j.jneumeth.2010.03.024.
- Marianne Leger, Eleni Paizanis, Kwamivi Dzahini, Anne Quiedeville, Valentine Bouet, Jean-Christophe Cassel, Thomas Freret, Pascale Schumann-Bard, and Michel Boulouard. Environmental enrichment duration differentially affects behavior and neuroplasticity in adult mice. *Cerebral Cortex*, 25(11):4048–4061, jun 2014. doi: 10.1093/cercor/bhu119. URL <https://doi.org/10.1093/cercor/bhu119>.
- D. Levitan, Y. Fortis-Santiago, J. A. Figueroa, E. E. Reid, T. Yoshida, N. C. Barry, A. Russo, and D. B. Katz. Memory retrieval has a dynamic influence on the maintenance mechanisms that are sensitive to -inhibitory peptide (ZIP). *Journal of Neuroscience*, 36(41):10654–10662, oct 2016. doi: 10.1523/jneurosci.1568-16.2016. URL <https://doi.org/10.1523/jneurosci.1568-16.2016>.

## BIBLIOGRAPHY

- Michael S Lewicki. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network: Computation in Neural Systems*, 9(4):R53–R78, 1998.
- Cheng-Shu Li and Young K Cho. Efferent projection from the bed nucleus of the stria terminalis suppresses activity of taste-responsive neurons in the hamster parabrachial nuclei. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 291(4):R914–R926, 2006.
- Cheng-Shu Li, Young K Cho, and David V Smith. Modulation of parabrachial taste neurons by electrical and chemical stimulation of the lateral hypothalamus and amygdala. *Journal of neurophysiology*, 93(3):1183–1196, 2005.
- Jennifer X Li, Takashi Yoshida, Kevin J Monk, and Donald B Katz. Lateral hypothalamus contains two types of palatability-related taste responses with distinct dynamics. *Journal of Neuroscience*, 33(22):9462–9473, 2013.
- Jennifer X Li, Joost X Maier, Emily E Reid, and Donald B Katz. Sensory cortical activity is related to the selection of a rhythmic motor action pattern. *Journal of Neuroscience*, 36(20):5596–5607, 2016.
- Jinrong Li, Jianqun Yan, Ke Chen, Bo Lu, Qian Wang, Wei Yan, and Xiaolin Zhao. Lesions of the central nucleus of the amygdala decrease taste threshold for sodium chloride in rats. *Brain Research Bulletin*, 89(1-2):8–15, oct 2012. doi: 10.1016/j.brainresbull.2012.06.013. URL <https://doi.org/10.1016%2Fj.brainresbull.2012.06.013>.
- William A Liberti III, Jeffrey E Markowitz, L Nathan Perkins, Derek C Liberti, Daniel P Leman, Grigori Guitchounts, Tarciso Velho, Darrell N Kotton, Carlos Lois, and Timothy J Gardner. Unstable neurons underlie a stable learned behavior. *Nature Neuroscience*, 19(12):1665–1671, 2016.
- Jian-You Lin, Chris Roman, Joe Arthurs, and Steve Reilly. Taste neophobia and c-fos expression in the rat brain. *Brain Research*, 1448:82–88, apr 2012. doi: 10.1016/j.brainres.2012.02.013. URL <https://doi.org/10.1016%2Fj.brainres.2012.02.013>.
- Annie Liu and Nathaniel N Urban. Prenatal and early postnatal odorant exposure heightens odor-evoked mitral cell responses in the mouse olfactory bulb. *Eneuro*, 4(5), 2017.

## BIBLIOGRAPHY

- Gerald E Loeb and Carl Gans. *Electromyography for experimentalists*. University of Chicago Press, 1986.
- A. W. Logue. Taste aversion and the generality of the laws of learning. *Psychological Bulletin*, 86(2):276–296, 1979. doi: 10.1037/0033-2909.86.2.276. URL <https://doi.org/10.1037/2F0033-2909.86.2.276>.
- Peter F Lovibond, GC Preston, and NJ Mackintosh. Context specificity of conditioning, extinction, and latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*, 10(3):360, 1984.
- RE Lubow and AU Moore. Latent inhibition: the effect of nonreinforced pre-exposure to the conditional stimulus. *Journal of comparative and physiological psychology*, 52(4):415, 1959.
- Robert E. Lubow. Latent inhibition. *Psychological Bulletin*, 79(6):398–407, 1973. doi: 10.1037/h0034425. URL <https://doi.org/10.1037%2Fh0034425>.
- Kourosh Maboudi, Etienne Ackermann, Laurel Watkins de Jong, Brad E Pfeiffer, David Foster, Kamran Diba, and Caleb Kemere. Uncovering temporal structure in hippocampal output patterns. *eLife*, 7:e34467, 2018.
- Arianna Maffei, Melissa Haley, and Alfredo Fontanini. Neural processing of gustatory information in insular circuits. *Current opinion in neurobiology*, 22(4):709–716, 2012.
- Joost X. Maier and Donald B. Katz. Neural dynamics in response to binary taste mixtures. *Journal of Neurophysiology*, 109(8):2108–2117, apr 2013a. doi: 10.1152/jn.00917.2012. URL <https://doi.org/10.1152%2Fjn.00917.2012>.
- Joost X Maier and Donald B Katz. Neural dynamics in response to binary taste mixtures. *Journal of neurophysiology*, 109(8):2108–2117, 2013b.
- Joost X Maier, Meredith L Blankenship, Jennifer X Li, and Donald B Katz. A multisensory network for olfactory processing. *Current Biology*, 25(20):2642–2650, 2015.
- Eve Marder and Dirk Bucher. Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs. *Annu. Rev. Physiol.*, 69:291–316, 2007.

## BIBLIOGRAPHY

- Mark Mayford and Leon Reijmers. Exploring memory representations with activity-based genetics. *Cold Spring Harbor Perspectives in Biology*, 8(3):a021832, dec 2015. doi: 10.1101/cshperspect.a021832. URL <https://doi.org/10.1101%2Fcshperspect.a021832>.
- Luca Mazzucato, Alfredo Fontanini, and Giancarlo La Camera. Dynamics of multistable states during ongoing and evoked cortical activity. *Journal of Neuroscience*, 35(21):8214–8231, 2015.
- David A McCormick, Barry W Connors, James W Lighthall, and David A Prince. Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of neurophysiology*, 54(4):782–806, 1985.
- Richard McElreath. Statistical rethinking. texts in statistical science, 2015.
- M. Merhav and K. Rosenblum. Facilitation of taste memory acquisition by experiencing previous novel taste is protein-synthesis dependent. *Learning & Memory*, 15(7):501–507, jul 2008. doi: 10.1101/lm.986008. URL <https://doi.org/10.1101%2F1m.986008>.
- Paul Miller. Dynamical systems, attractors, and neural circuits. *F1000Research*, 5, 2016.
- Paul Miller and Donald B Katz. Stochastic transitions between neural states in taste processing and decision-making. *Journal of Neuroscience*, 30(7):2559–2570, 2010.
- Paul Miller and Donald B Katz. Accuracy and response-time distributions for decision-making: linear perfect integrators versus nonlinear attractor-based neural circuits. *Journal of computational neuroscience*, 35(3):261–294, 2013.
- Ralph R. Miller and Arnold D. Holzman. Neophobias and conditioned taste aversions in rats following exposure to novel flavors. *Animal Learning & Behavior*, 9(1):89–100, mar 1981. doi: 10.3758/bf03212030. URL <https://doi.org/10.3758%2Fbf03212030>.
- M Miranda. Role of cholinergic system on the construction of memories: Taste memory encoding. *Neurobiology of Learning and Memory*, 80(3):211–222, nov 2003. doi: 10.1016/s1074-7427(03)00061-3. URL <https://doi.org/10.1016%2Fs1074-7427%2803%2900061-3>.

Mari

## BIBLIOGRAPHY

- a Isabel Miranda, Leticia Ramırez-Lugo, and Federico Bermúdez-Rattoni. Cortical cholinergic activity is related to the novelty of the stimulus. *Brain Research*, 882(1-2):230–235, nov 2000. doi: 10.1016/s0926-6410(00)00050-1. URL <https://doi.org/10.1016%2Fs0926-6410%2800%2900050-1>.
- Mikael Molet and Ralph R Miller. Timing: an attribute of associative learning. *Behavioural processes*, 101:4–14, 2014.
- D. Moncada, F. Ballarini, M. C. Martinez, J. U. Frey, and H. Viola. Identification of transmitter systems and learning tag molecules involved in behavioral tagging during memory formation. *Proceedings of the National Academy of Sciences*, 108(31):12931–12936, jul 2011. doi: 10.1073/pnas.1104495108. URL <https://doi.org/10.1073%2Fpnas.1104495108>.
- K. J. Monk, B. D. Rubin, J. C. Keene, and D. B. Katz. Licking microstructure reveals rapid attenuation of neophobia. *Chemical Senses*, 39(3):203–213, dec 2013. doi: 10.1093/chemse/bjt069. URL <https://doi.org/10.1093%2Fchemse%2Fbjt069>.
- Jeffrey D Moore, David Kleinfeld, and Fan Wang. How the brainstem controls orofacial behaviors comprised of rhythmic actions. *Trends in neurosciences*, 37(7):370–380, 2014.
- Rodrigo Moraga-Amaro, Andrés Cortés-Rojas, Felipe Simon, and Jimmy Stehberg. Role of the insular cortex in taste familiarity. *Neurobiology of learning and memory*, 109:37–45, Mar 2014. doi: 10.1016/j.nlm.2013.11.012.
- A. Moran and D. B. Katz. Sensory cortical population dynamics uniquely track behavior across learning and extinction. *Journal of Neuroscience*, 34(4):1248–1257, jan 2014a. doi: 10.1523/jneurosci.3331-13.2014. URL <https://doi.org/10.1523%2Fjneurosci.3331-13.2014>.
- Anan Moran and Donald B Katz. Sensory cortical population dynamics uniquely track behavior across learning and extinction. *Journal of Neuroscience*, 34(4):1248–1257, 2014b.
- Ruben Moreno-Bote, John Rinzel, and Nava Rubin. Noise-induced alternations in an attractor network model of perceptual bistability. *Journal of neurophysiology*, 98(3):1125–39, Sep 2007.

## BIBLIOGRAPHY

- Ruben Moreno-Bote, David C Knill, and Alexandre Pouget. Bayesian sampling in visual perception. *Proceedings of the National Academy of Sciences*, 108(30):12491–12496, 2011.
- I. Morón and M. Gallo. Effect of previous taste experiences on taste neophobia in young-adult and aged rats. *Physiology & Behavior*, 90(2-3):308–317, feb 2007. doi: 10.1016/j.physbeh.2006.09.036. URL <https://doi.org/10.1016%2Fj.physbeh.2006.09.036>.
- Narendra Mukherjee, Joseph Wachutka, and Donald B Katz. Python meets systems neuroscience: affordable, scalable and open-source electrophysiology in awake, behaving rodents. 2017.
- Narendra Mukherjee, Joseph Wachutka, and Donald B Katz. Dynamical structure of cortical taste responses revealed by precisely-timed optogenetic perturbation. *bioRxiv*, 2019. doi: 10.1101/486043. URL <https://www.biorxiv.org/content/early/2019/02/09/486043>.
- Marvin Nachman. Learned taste and temperature aversions due to lithium chloride sickness after temporal delays. *Journal of Comparative and Physiological Psychology*, 73(1):22, 1970.
- Marvin Nachman and John H. Ashe. Learned taste aversions in rats as a function of dosage, concentration, and route of administration of LiCl. *Physiology & Behavior*, 10(1):73–78, jan 1973. doi: 10.1016/0031-9384(73)90089-9. URL <https://doi.org/10.1016%2F0031-9384%2873%2990089-9>.
- Montserrat Navarro, Kristina J. Spray, Inmaculada Cubero, Todd E. Thiele, and Ilene L. Bernstein. cFos induction during conditioned taste aversion expression varies with aversion strength. *Brain Research*, 887(2):450–453, dec 2000. doi: 10.1016/s0006-8993(00)03032-8. URL <https://doi.org/10.1016%2Fs0006-8993%2800%2903032-8>.
- Karly N. Neath, Cheryl L. Limebeer, Steve Reilly, and Linda A. Parker. Increased liking for a solution is not necessary for the attenuation of neophobia in rats. *Behavioral Neuroscience*, 124(3):398–404, 2010. doi: 10.1037/a0019505. URL <https://doi.org/10.1037%2Fa0019505>.
- Selin Neseliler, Darshana Narayanan, Yaihara Fortis-Santiago, Donald B. Katz, and Susan J. Birren. Genetically induced cholinergic hyper-innervation enhances taste learn-

## BIBLIOGRAPHY

- ing. *Frontiers in Systems Neuroscience*, 5, 2011. doi: 10.3389/fnsys.2011.00097. URL <https://doi.org/10.3389%2Ffnsys.2011.00097>.
- Hisao Nishijo, Teruko Uwano, Ryoji Tamura, and Taketoshi Ono. Gustatory and multimodal neuronal responses in the amygdala during licking and discrimination of sensory stimuli in awake rats. *Journal of neurophysiology*, 79(1):21–36, 1998.
- Ralph Norgren. Taste pathways to hypothalamus and amygdala. *Journal of Comparative Neurology*, 166(1):17–30, 1976.
- Hiroshi Okamoto, Yoshikazu Isomura, Masahiko Takada, and Tomoki Fukai. Temporal integration by stochastic recurrent network dynamics with bimodal neurons. *Journal of neurophysiology*, 2007.
- Chethan Pandarinath, Daniel J O’Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D Stavisky, Jonathan C Kao, Eric M Trautmann, Matthew T Kaufman, Stephen I Ryu, Leigh R Hochberg, et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature methods*, page 1, 2018.
- Ryan G. Parsons. Behavioral and neural mechanisms by which prior experience impacts subsequent learning. *Neurobiology of Learning and Memory*, 154:22–29, oct 2018. doi: 10.1016/j.nlm.2017.11.008. URL <https://doi.org/10.1016%2Fj.nlm.2017.11.008>.
- Erika Pastrana. Optogenetics: controlling cell function with light. *Nature Methods*, 8(1):24, 2011.
- George Paxinos and Charles Watson. *The Rat Brain in Stereotaxic Coordinates in Stereotaxic Coordinates*. Elsevier, 2007.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- R. Pedroza-Llinás, L. Ramírez-Lugo, K. Guzmán-Ramos, S. Zavala-Vega, and F. Bermúdez-Rattoni. Safe taste memory consolidation is disrupted by a protein synthesis inhibitor in the nucleus accumbens shell. *Neurobiology of Learning and Memory*, 92(1):45–52, jul

## BIBLIOGRAPHY

2009. doi: 10.1016/j.nlm.2009.02.011. URL <https://doi.org/10.1016%2Fj.nlm.2009.02.011>.
- Yangfan Peng, Barreda Tomás, J Federico, Constantin Klisch, Imre Vida, and Jörg RP Geiger. Layer-specific organization of local excitatory and inhibitory synaptic connectivity in the rat presubiculum. *Cerebral Cortex*, 27(4):2435–2452, 2017.
- Yueqing Peng, Sarah Gillis-Smith, Hao Jin, Dimitri Tränkner, Nicholas JP Ryba, and Charles S Zuker. Sweet and bitter taste in the brain of awake behaving animals. *Nature*, 527(7579):512, 2015.
- Isaac O Perez, Miguel Villavicencio, Sidney A Simon, and Ranier Gutierrez. Speed and accuracy of taste identification and palatability: impact of learning, reward expectancy, and consummatory licking. *American journal of physiology. Regulatory, integrative and comparative physiology*, 305(3):R252–70, Aug 2013. doi: 10.1152/ajpregu.00492.2012.
- Marios G Philiastides and Paul Sajda. Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral cortex (New York, N. Y. : 1991)*, 16(4):509–18, apr 2006. ISSN 1047-3211. doi: 10.1093/cercor/bhi130.
- MI Phillips and RE Norgren. A rapid method for permanent implantation of an intraoral fistula in rats. *Behavior Research Methods & Instrumentation*, 2(3):124–124, 1970.
- C. E. Piette, M. A. Baez-Santiago, E. E. Reid, D. B. Katz, and A. Moran. Inactivation of basolateral amygdala specifically eliminates palatability-related information in cortical sensory responses. *Journal of Neuroscience*, 32(29):9981–9991, jul 2012a. doi: 10.1523/jneurosci.0669-12.2012. URL <https://doi.org/10.1523%2Fjneurosci.0669-12.2012>.
- Caitlin E Piette, Madelyn A Baez-Santiago, Emily E Reid, Donald B Katz, and Anan Moran. Inactivation of basolateral amygdala specifically eliminates palatability-related information in cortical sensory responses. *Journal of Neuroscience*, 32(29):9981–9991, 2012b.
- Patricia Pliner, Marcia Pelchat, and Marius Grabski. Reduction of neophobia in humans by exposure to novel foods. *Appetite*, 20(2):111–123, apr 1993. doi: 10.1006/appe.1993.1013. URL <https://doi.org/10.1006%2Fappe.1993.1013>.

## BIBLIOGRAPHY

- Adrián Ponce-Alvarez, Verónica Nácher, Rogelio Luna, Alexa Riehle, and Ranulfo Romo. Dynamics of cortical neuronal ensembles transit from decision making to storage for later report. *Journal of Neuroscience*, 32(35):11956–11969, 2012.
- Wolfgang Prinz. Perception and action planning. *European journal of cognitive psychology*, 9(2):129–154, 1997.
- R Quian Quiroga, Zoltan Nadasdy, and Yoram Ben-Shaul. Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural computation*, 16(8):1661–1687, 2004.
- Jovita F Randall-Thompson and Anthony L Riley. Morphine-induced conditioned taste aversions: assessment of sexual dimorphism. *Pharmacology Biochemistry and Behavior*, 76(2):373–381, sep 2003. doi: 10.1016/j.pbb.2003.08.010. URL <https://doi.org/10.1016%2Fj.pbb.2003.08.010>.
- Roger Ratcliff, Anil Cherian, and Mark Segraves. A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of simple two-choice decisions. *Journal of neurophysiology*, 2003.
- Samuel H Revusky. Aversion to sucrose produced by contingent x-irradiation: Temporal and dosage parameters. *Journal of Comparative and Physiological Psychology*, 65(1):17, 1968.
- Walter H. Riege. Environmental influences on brain and behavior of year-old rats. *Developmental Psychobiology*, 4(2):157–167, 1971. doi: 10.1002/dev.420040207. URL <https://doi.org/10.1002%2Fdev.420040207>.
- Christopher A Riley and Michael S King. Differential effects of electrical stimulation of the central amygdala and lateral hypothalamus on fos-immunoreactive neurons in the gustatory brainstem and taste reactivity behaviors in conscious rats. *Chemical senses*, 38(8):705–717, 2013.
- J RINKER, G BUSSE, and A RILEY. An assessment of sex differences in nicotine-induced conditioned taste aversions. *Pharmacology Biochemistry and Behavior*, 88(4):427–431, feb 2008. doi: 10.1016/j.pbb.2007.09.016. URL <https://doi.org/10.1016%2Fj.pbb.2007.09.016>.

## BIBLIOGRAPHY

- Luis F. Rodríguez-Durán, Diana V. Castillo, Minerva Moguel-González, and Martha L. Escobar. Conditioned taste aversion modifies persistently the subsequent induction of neocortical long-term potentiation in vivo. *Neurobiology of Learning and Memory*, 95(4): 519–526, may 2011. doi: 10.1016/j.nlm.2011.03.003. URL <https://doi.org/10.1016%2Fj.nlm.2011.03.003>.
- E.T. Rolls, F. Grabenhorst, and G. Deco. Decision-making, errors, and confidence in the brain. *Journal of neurophysiology*, 104(5):2359, 2010.
- Christopher Roman and Steve Reilly. Effects of insular cortex lesions on conditioned taste aversion and latent inhibition in the rat. *European Journal of Neuroscience*, 26(9):2627–2632, oct 2007. doi: 10.1111/j.1460-9568.2007.05872.x. URL <https://doi.org/10.1111%2Fj.1460-9568.2007.05872.x>.
- Ranulfo Romo, Adrian Hernandez, Antonio Zainos, Carlos Brody, and Emilio Salinas. Exploring the cortical evidence of a sensory-discrimination process. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 357(1424):1039–51, Aug 2002.
- Kobi Rosenblum, Noam Meiri, and Yadin Dudai. Taste memory: The role of protein synthesis in gustatory cortex. *Behavioral and Neural Biology*, 59(1):49–56, jan 1993. doi: 10.1016/0163-1047(93)91145-d. URL <https://doi.org/10.1016%2F0163-1047%2893%2991145-d>.
- Brian F Sadacca, Jason T Rothwax, and Donald B Katz. Sodium concentration coding gives way to evaluative coding in cortex and amygdala. *Journal of Neuroscience*, 32(29): 9999–10011, 2012.
- Brian F Sadacca, Narendra Mukherjee, Tony Vladusich, Jennifer X Li, Donald B Katz, and Paul Miller. The behavioral relevance of cortical neural ensemble responses emerges suddenly. *Journal of Neuroscience*, 36(3):655–669, 2016.
- John Salvatier, Thomas V Wiecki, and Christopher Fonnesebeck. Probabilistic programming in python using pymc3. *PeerJ Computer Science*, 2:e55, 2016.

## BIBLIOGRAPHY

- Chad L Samuelsen and Alfredo Fontanini. Processing of intraoral olfactory and gustatory signals in the gustatory cortex of awake rats. *Journal of Neuroscience*, pages 1926–16, 2016.
- Chad L Samuelsen, Matthew PH Gardner, and Alfredo Fontanini. Effects of cue-triggered expectation on cortical processing of taste. *Neuron*, 74(2):410–422, 2012.
- Kazushige Sasamoto, Guixin Zhang, and Makoto Iwasaki. Two types of rhythmical jaw movements evoked by stimulation of the rat cortex. *Japanese Journal of Oral Biology*, 32(1):57–68, 1990.
- G SCALERA. Acid taste thresholds assessed by conditioned taste aversion and two-bottle preference in rats. *Physiology & Behavior*, 82(2-3):411–423, sep 2004. doi: 10.1016/j.physbeh.2004.04.052. URL <https://doi.org/10.1016%2Fj.physbeh.2004.04.052>.
- G E Schafe and I L Bernstein. Forebrain contribution to the induction of a brainstem correlate of conditioned taste aversion. ii. insular (gustatory) cortex. *Brain research*, 800(1):40–7, Jul 1998.
- L. A. Schier, K. Hashimoto, M. B. Bales, G. D. Blonde, and A. C. Spector. High-resolution lesion-mapping strategy links a hot spot in rat insular cortex with impaired expression of taste aversion learning. *Proceedings of the National Academy of Sciences*, 111(3):1162–1167, jan 2014. doi: 10.1073/pnas.1315624111. URL <https://doi.org/10.1073%2Fpnas.1315624111>.
- Lindsey A. Schier, Ginger D. Blonde, and Alan C. Spector. Bilateral lesions in a specific subregion of posterior insular cortex impair conditioned taste aversion expression in rats. *Journal of Comparative Neurology*, 524(1):54–73, jul 2015. doi: 10.1002/cne.23822. URL <https://doi.org/10.1002%2Fcne.23822>.
- Steven J Schiff. Kalman meets neuron: the emerging intersection of control theory with neuroscience. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 3318–3321. IEEE, 2009.
- Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin

## BIBLIOGRAPHY

- Schmid, Jean-Yves Tinevez, Daniel James White, Volker Hartenstein, Kevin Eliceiri, Pavel Tomancak, and Albert Cardona. Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9(7):676–682, jul 2012. doi: 10.1038/nmeth.2019. URL <https://doi.org/10.1038/nmeth.2019>.
- Michael N Shadlen and William T Newsome. Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of neurophysiology*, 86(4):1916–1936, 2001.
- SJ Shammah-Lagnado, MSMO Costa, and JA Ricardo. Afferent connections of the parvocellular reticular formation: a horseradish peroxidase study in the rat. *Neuroscience*, 50(2):403–425, 1992.
- David A.T. Siddle and Nigel W. Bond. Avoidance learning, pavlovian conditioning, and the development of phobias. *Biological Psychology*, 27(2):167–183, oct 1988. doi: 10.1016/0301-0511(88)90048-8. URL [https://doi.org/10.1016/0301-0511\(88\)90048-8](https://doi.org/10.1016/0301-0511(88)90048-8).
- Joshua H Siegle, Gregory J Hale, Jonathan P Newman, and Jakob Voigts. Neural ensemble communities: open-source approaches to hardware for large-scale electrophysiology. *Current opinion in neurobiology*, 32:53–59, 2015.
- David V Smith and Steven J St John. Neural coding of gustatory information. *Current opinion in neurobiology*, 9(4):427–435, 1999.
- James C Smith, John D Davis, and Gwendolen B O’Keefe. Lack of an order effect in brief contact taste tests with closely spaced test trials. *Physiology & Behavior*, 52(6):1107–1111, 1992.
- Virginia L Smith-Swintosky, Carlos R Plata-Salaman, and THOMAS R Scott. Gustatory neural coding in the monkey cortex: stimulus quality. *Journal of Neurophysiology*, 66(4):1156–1165, 1991.
- Alberto Soto, Patricia Gasalla, Azucena Begega, and Matías López. c-fos activity in the insular cortex, nucleus accumbens and basolateral amygdala following the intraperitoneal injection of saccharin and lithium chloride. *Neuroscience Letters*, 647:32–37, apr 2017.

## BIBLIOGRAPHY

- doi: 10.1016/j.neulet.2017.03.025. URL <https://doi.org/10.1016%2Fj.neulet.2017.03.025>.
- Alan C Spector, Paul Breslin, and Harvey J Grill. Taste reactivity as a dependent measure of the rapid formation of conditioned taste aversion: a tool for the neural analysis of taste-visceral associations. *Behavioral neuroscience*, 102(6):942, 1988.
- Kristina J. Spray, Christopher B. Halsell, and Ilene L. Bernstein. c-fos induction in response to saccharin after taste aversion learning depends on conditioning method. *Brain Research*, 852(1):225–227, jan 2000. doi: 10.1016/s0006-8993(99)02203-9. URL <https://doi.org/10.1016%2Fs0006-8993%2899%2902203-9>.
- Jennifer R Stapleton, Michael L Lavine, Robert L Wolpert, Miguel AL Nicolelis, and Sidney A Simon. Rapid taste responses in the gustatory cortex during licking. *Journal of Neuroscience*, 26(15):4126–4138, 2006.
- Jimmy Stehberg, Rodrigo Moraga-Amaro, and Felipe Simon. The role of the insular cortex in taste function. *Neurobiology of learning and memory*, 96(2):130–5, Sep 2011. doi: 10.1016/j.nlm.2011.03.005.
- Robert J Sternberg and Janet E Davidson. *The Nature of Insight*. MIT Press, Cambridge, Mass, 1995. ISBN 0262193450.
- M. E. Stone. Hippocampal inactivation enhances taste learning. *Learning & Memory*, 12(6):579–586, nov 2005. doi: 10.1101/lm.32305. URL <https://doi.org/10.1101%2F1m.32305>.
- Mark Stopfer, Vivek Jayaraman, and Gilles Laurent. Intensity versus identity coding in an olfactory system. *Neuron*, 39(6):991 – 1004, 2003. ISSN 0896-6273. doi: <https://doi.org/10.1016/j.neuron.2003.08.011>. URL <http://www.sciencedirect.com/science/article/pii/S089662730300535X>.
- Michael W. Swank and Ilene L. Bernstein. c-fos induction in response to a conditioned stimulus after single trial taste aversion learning. *Brain Research*, 636(2):202–208, feb 1994. doi: 10.1016/0006-8993(94)91018-9. URL <https://doi.org/10.1016%2F0006-8993%2894%2991018-9>.

## BIBLIOGRAPHY

- O. Tange. Gnu parallel - the command-line power tool. *login: The USENIX Magazine*, 36(1):42–47, Feb 2011. doi: <http://dx.doi.org/10.5281/zenodo.16303>. URL <http://www.gnu.org/s/parallel>.
- Kirk G Thompson, Doug P Hanes, Narcisse P Bichot, and Jeffrey D Schall. Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *Journal of Neurophysiology*, 76(6):4040–4055, 1996.
- Joseph Travers, Lisa DiNardo, and Hamid Karimnamazi. Medullary reticular formation activity during ingestion and rejection in the awake rat. *Experimental brain research*, 130(1):78–92, 2000.
- Joseph B Travers and Ralph Norgren. Electromyographic analysis of the ingestion and rejection of sapid stimuli in the rat. *Behavioral neuroscience*, 100(4):544, 1986.
- Joseph B Travers, Lisa A Dinardo, and Hamid Karimnamazi. Motor and premotor mechanisms of licking. *Neuroscience & Biobehavioral Reviews*, 21(5):631–647, 1997.
- Naoshige Uchida and Zachary F Mainen. Speed and accuracy of olfactory discrimination in the rat. *Nature neuroscience*, 6(11):1224, 2003.
- Akira Uematsu, Akihiko Kitamura, Ken Iwatsuki, Hisayuki Uneyama, and Tomokazu Tsugizawa. Correlation between activation of the prelimbic cortex, basolateral amygdala, and agranular insular cortex during taste memory formation. *Cerebral Cortex*, 25(9):2719–2728, apr 2014. doi: [10.1093/cercor/bhu069](https://doi.org/10.1093/cercor/bhu069). URL <https://doi.org/10.1093/cercor/bhu069>.
- Sharmila Venugopal, Joseph B Travers, and David H Terman. A computational model for motor pattern switching between taste-induced ingestion and rejection oromotor behaviors. *Journal of computational neuroscience*, 22(2):223–238, 2007.
- Roberto Vincis and Alfredo Fontanini. Associative learning changes cross-modal representations in the gustatory cortex. *Elife*, 5:e16420, 2016a.
- Roberto Vincis and Alfredo Fontanini. A gustocentric perspective to understanding primary sensory cortices. *Current opinion in neurobiology*, 40:118–124, 2016b.

## BIBLIOGRAPHY

- Edgar T Walters and John H Byrne. Associative conditioning of single sensory neurons suggests a cellular mechanism for learning. *Science*, 219(4583):405–408, 1983.
- Xiao-Jing Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–968, 2002.
- Yingxue Wang, Sandro Romani, Brian Lustig, Anthony Leonardo, and Eva Pastalkova. Theta sequences are essential for internally generated hippocampal firing fields. *Nature neuroscience*, 18(2):282–288, 2015.
- Ina Weiner. What the brain teaches us about latent inhibition (li): the neural substrates of the expression and prevention of li. *Latent Inhibition: Cognition, Neuroscience and Applications to Schizophrenia*. Cambridge University Press. New York, pages 372–415, 2010.
- Michael S Weiss and Patricia M Di Lorenzo. Not so fast: taste stimulus coding time in the rat revisited. *Frontiers in integrative neuroscience*, 6:27, 2012. doi: 10.3389/fnint.2012.00027.
- Emily E. Wilkins and Ilene L. Bernstein. Conditioning method determines patterns of c-fos expression following novel taste-illness pairing. *Behavioural Brain Research*, 169(1):93–97, apr 2006. doi: 10.1016/j.bbr.2005.12.006. URL <https://doi.org/10.1016%2Fj.bbr.2005.12.006>.
- Daniel M Wolpert and Zoubin Ghahramani. Computational principles of movement neuroscience. *Nature neuroscience*, 3(11s):1212, 2000.
- Daniel M Wolpert and Mitsuo Kawato. Multiple paired forward and inverse models for motor control. *Neural networks*, 11(7-8):1317–1329, 1998.
- Daniel M Wolpert, Zoubin Ghahramani, and Michael I Jordan. An internal model for sensorimotor integration. *Science*, 269(5232):1880–1882, 1995.
- Kong-Fatt Wong and Xiao-Jing Wang. A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience*, 26(4):1314–1328, 2006.
- Frank Wood, Michael J Black, Carlos Vargas-Irwin, Matthew Fellows, and John P Donoghue. On the variability of manual spike sorting. *IEEE Transactions on Biomedical Engineering*, 51(6):912–918, 2004.

## BIBLIOGRAPHY

- T. Yamamoto, N. Yuyama, T. Kato, and Y. Kawamura. Gustatory responses of cortical neurons in rats. III. neural and behavioral measures compared. *Journal of Neurophysiology*, 53(6):1370–1386, jun 1985. doi: 10.1152/jn.1985.53.6.1370. URL <https://doi.org/10.1152%2Fjn.1985.53.6.1370>.
- Takashi Yamamoto, Noriyuki Yuyama, Takuya Kato, and Yojiro Kawamura. Gustatory responses of cortical neurons in rats. i. response characteristics. *Journal of neurophysiology*, 51(4):616–635, 1984.
- Takashi Yamamoto, Ryuji Matsuo, Yoshitaka Kiyomitsu, and Ryuji Kitamura. Response properties of lateral hypothalamic neurons during ingestive behavior with special reference to licking of various taste solutions. *Brain research*, 481(2):286–297, 1989.
- Y. Yasoshima, N. Sako, E. Senba, and T. Yamamoto. Acute suppression, but not chronic genetic deficiency, of c-fos gene expression impairs long-term memory in aversive taste learning. *Proceedings of the National Academy of Sciences*, 103(18):7106–7111, apr 2006. doi: 10.1073/pnas.0600869103. URL <https://doi.org/10.1073%2Fpnas.0600869103>.
- Ofer Yizhar, Lief E. Fenno, Thomas J. Davidson, Murtaza Mogri, and Karl Deisseroth. Optogenetics in neural systems. *Neuron*, 71(1):9–34, jul 2011a. doi: 10.1016/j.neuron.2011.06.004. URL <https://doi.org/10.1016%2Fj.neuron.2011.06.004>.
- Ofer Yizhar, Lief E Fenno, Thomas J Davidson, Murtaza Mogri, and Karl Deisseroth. Optogenetics in neural systems. *Neuron*, 71(1):9–34, 2011b.
- T Yokota, K Eguchi, and K Hiraba. Functional properties of putative pyramidal neurons and inhibitory interneurons in the rat gustatory cortex. *Cerebral Cortex*, 21(3):597–606, 2011.
- Takashi Yoshida and Donald B Katz. Control of prestimulus activity related to improved sensory coding within a discrimination task. *The Journal of Neuroscience*, 31(11):4101–4112, 2011.
- Itay Zelcer, Hagit Cohen, Gal Richter-Levin, Tom Lebiosn, Tomer Grossberger, and Edi Barkai. A cellular correlate of learning-induced metaplasticity in the hippocampus. *Cerebral Cortex*, 16(4):460–468, jun 2005. doi: 10.1093/cercor/bhi125. URL <https://doi.org/10.1093%2Fcercor%2Fbhi125>.

## BIBLIOGRAPHY

- Feng Zhang, Viviana Gradinaru, Antoine R Adamantidis, Remy Durand, Raag D Airan, Luis de Lecea, and Karl Deisseroth. Optogenetic interrogation of neural circuits: technology for probing mammalian brain structures. *Nature Protocols*, 5(3):439–456, feb 2010. doi: 10.1038/nprot.2009.226. URL <https://doi.org/10.1038/nprot.2009.226>.
- Guixin Zhang and Kazushige Sasamoto. Projections of two separate cortical areas for rhythmic jaw movements in the rat. *Brain research bulletin*, 24(2):221–230, 1990.