UNIVERSITY OF CALIFORNIA SAN DIEGO

Volitional control of neuromodulation

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Physics

by

Conrad Foo

Committee in charge:

Professor David Kleinfeld, Chair
Professor Brenda Bloodgood
Professor Alexander Groisman
Professor Douglas Smith
Professor Jing Wang

2021

The dissertation of Conrad Foo is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2021

DEDICATION

I dedicate this thesis to my mother and father, who have provided the resources necessary for me to pursue this degree.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF SUPPLEMENTAL FILES

Foo_Code.zip

ACKNOWLEDGEMENTS

I would like to thank my committee, Professors Brenda Bloodgood, Alexander Groisman, Doug Smith and Jing Wang. In particular, I would like to especially thank Professor Jing Wang for his deep insight into the project and suggestions on how to interpret my data. I would also like to thank Professor Alexander Groisman for his encouraging words and help with my writing.

I would like to specifically acknowledge Professor David Kleinfeld for his support as my advisor. David was involved in the project since its inception, and was supportive of the direction I took with the project. He was available if I had any issues and helped guide the direction of the project if I ever hit a dead end.

Chapters 1-4 contain parts of a manuscript in preparation: Foo C, Lozada A, Aljadeff J, Wang J, Li Y, Slesinger P, Kleinfeld D, “Volitional control links spontaneous dopamine transients to reward” 2021. The dissertation author is the primary author of this work.

Appendix 1 contains unpublished work done in collaboration with Adrian Lozada, Paul Slesinger, and David Kleinfeld. The dissertation author is the primary author of this work.

Appendix 2 contains unpublished work done in collaboration with Paul Slesinger and David Kleinfeld. The dissertation author is the primary author of this work.

VITA

2021 Ph.D., Physics
University of California San Diego

2012 Bachelor of Science in Engineering Physics
University of Illinois, Urbana-Champaign

ABSTRACT OF THE DISSERTATION

Volitional control of neuromodulation

by

Conrad Foo

Doctor of Philosophy in Physics

University of California San Diego, 2021

Professor David Kleinfeld, Chair

Neuromodulatory neurons release neurotransmitters extrasynaptically to modulate large groups of neurons. Although much work has been done to understand neuromodulation, this has been primarily measured indirectly via

action potentials recorded electrically from neuromodulatory somata. In this work, we use genetically modified CNiFER cells, a technique developed in the Kleinfeld lab, to directly measure changes in neuromodulator concentration in real-time. We focus our investigation primarily on cortical noradrenaline, dopamine and acetylcholine, for which CNiFER cells have already been developed. We ask whether spontaneous neuromodulator release occurs in cortex of mice. We find that, indeed, spontaneous dopamine, noradrenaline, and acetylcholine transients occur in cortex of mice. Furthermore, we use real-time feedback of cortical neuromodulation and reinforcement to show that mice can volitionally increase cortical dopamine and noradrenaline levels. We show that mice are able to volitionally link spontaneous dopamine transients to future reward.

Chapter 1. Introduction

Neuromodulation is the process by which a small subset of neurons exerts influence over large brain regions by extrasynaptic release of neurotransmitters. Extrasynaptic neurotransmitters are known as neuromodulators in this context. Previous work on neuromodulation has focused primarily on electrical recordings of subcortical neuromodulatory nuclei, rather than direct measurements of neuromodulator concentration in target areas.

However, more recent studies that measure neuromodulator concentration directly have found that the electrical activity of these nuclei might not necessarily be causally linked to increases in neuromodulator concentration. In this work, we focus primarily on directly measuring cortical neuromodulation using genetically modified CNiFER cells.

We focus primarily on dopamine, which has been observed to be involved in learning to associate sensory cues with reward. We ask whether cortical dopamine release occurs spontaneously without reward, and if so, whether such spontaneous cortical dopamine release can be volitionally controlled.

1.1 Role of dopamine in the brain

Dopamine is one of the first neurotransmitters discovered in the mammalian brain. It has been implicated in a diverse set of functional roles, from reward related signaling to motor control. Dopaminergic neurons are localized to two major brain areas, the substantia nigra pars compacta (SNc), and the ventral tegmental area (VTA) [40].

Dopamine was first identified as a neurotransmitter by Carlsson in the late 1950s [69]. While dopamine was known as a precursor for norepinephrine synthesis, Carlsson discovered that the distribution of dopamine is significantly different from that of norepinephrine, suggesting that dopamine could work as a neurotransmitter rather than merely a precursor for norepinephrine. Furthermore, he discovered that injection of a dopamine precursor, DOPA, restored normal behavior in rabbits injected earlier with reserpine. Reserpine injection is known to deplete dopamine and noradrenaline, and leads to sedation. Carlsson later observed that the largest concentration of dopamine in the brain occurred in the striatum, suggesting that dopamine played an important role in motor control.

Later studies [68] found that lesions of SNc led to symptoms associated with Parkinson's disease: bradykinesia and unstable posture. These defects were ameliorated when L-DOPA, a dopamine precursor, is given orally. However, these lesions did not lead to the tremors associated with the disease. A new model, the MPTP monkey model of Parkinson's disease, was able to reproduce the tremors that other models did not [67]. MPTP given to monkeys is oxidized into MPP+ in the brain. MPP+ is selectively taken up by dopaminergic neurons, where it interferes with mitochondrial function and causes cell death. Examining the brains of monkeys given MPTP, Langston found that nearly complete degeneration of dopaminergic neurons in SNc and VTA, as well as variable amounts of damage to noradrenergic neurons in locus coeruleus (LC). Monkeys showed improvement with L-DOPA treatment, consistent with the effect in Parkinson's disease patients. These results suggest that dopaminergic activity in

SNc and VTA is vital to fine motor control and defects in dopamine release lead to observable behavioral defects.

Seminal work by Wolfram Schultz has shown that dopamine also signals reward prediction error [66]. He found that in monkeys, dopaminergic neurons in the VTA were mostly silent, but receipt of a juice reward elicited a short burst of activity in these cells. Interestingly, if the juice reward was preceded reliably by a cue (i.e. a tone), the neuronal activity shifted from the reward to the cue. If the reward typically associated with a cue fails to appear, the firing rate of dopaminergic neurons decreased. This bidirectional association has led to the idea that phasic bursts of dopamine encode reward prediction errors.

The temporal shift of dopaminergic neuronal activity from the reward to the cue has been the subject of much theoretical interest. The most prominent theoretical explanation of this shift is the "eligibility trace" [13]. The eligibility trace hypothesis posits that slowly decaying neuronal plasticity mechanisms are modulated by dopamine. Synapses active before the receipt of reward, including those that lead to release of dopamine, activate slowly decaying spike-timing dependent plasticity mechanisms (the eligibility trace). These mechanisms are enhanced in the presence of dopamine and weakened in the absence of dopamine. Thus, reward-predictive cues that activate synapses shortly before reward delivery are enhanced by reward delivery that occurs later via the eligibility trace mechanism.

On a network scale, temporal difference models of reinforcement learning have also proved capable of shifting activation from reward to reward-predicting

cues. These models maintain a record of expected future rewards given the current state (the "value" function), and update their predictions as time passes. The propagation of value information from the reward to the reward-predicting cue is explicitly included in the temporal difference update equation. However, these models predict that this information will propagate smoothly backwards in time, which may not be consistent with firing rates recorded from neurons.

Schultz also found that bursts of dopaminergic cell activity were associated with novelty or impact of sensory stimuli [66], independent of any reward value learned. This response preceded the reward prediction error signal, which occurred approximately 50ms later.

Dopamine also seems to motivate the initiation of behaviors that lead to rewards, rather than code for the value of the reward. Koob [65] found that dopamine depletion in nucleus accumbens decreased locomotor activity associated with obtaining food in hungry rats, but did not impair the consumption of food. Female rats engaged in less proceptive behavior when dopamine is depleted.

Although dopamine has been shown to play a critical role in a diverse set of behaviors, dopaminergic release is confined to a small set of neurons located in the basal ganglia. How does this small set of neurons manage to perform such a diverse set of roles? Although the somata of dopaminergic neurons are localized to a small region, dopaminergic axons are long and pass through large regions of the brain. Recent studies have found that the firing rates of dopaminergic neurons, as measured from the somata, are not necessarily

correlated with changes in dopamine concentrations measured in their projection targets [8,9]. Threlfell et al found that cholinergic inputs to striatum regulate dopamine release [64]. Optogenetic stimulation of striatal cholinergic interneurons evoked dopamine release, measured with cyclic voltammetry. Notably, these experiments were done in coronal slices that did not contain dopaminergic somata, suggesting that dopamine release does not necessarily arise from action potentials originating in the soma.

Another study by Mohebi et al found that, while dopamine concentrations in the nucleus accumbens (NAc) core is correlated with reward rate in an operant conditioning task, the firing rate of midbrain dopamine neurons in VTA is not [63]. Mohebi trained rats to perform an operant conditioning bandit task. Rats were asked to choose between two ports. Each port had a certain probability of reward when chosen, delivered to a separate food port behind the apparatus. Each trial had several discrete steps. First, a light cue signaled that the apparatus was ready to deliver reward. After some delay, rats would poke their nose into the center port. After a variable hold period, a white noise burst signaled the rats to leave the center port and poke their nose into either the right or left port. If the rat's choice led to reward, a food-hopper click signaled the dispensation of reward.

Mohebi et al monitored dopamine release in multiple regions of the rat's brain during this task using microdialysis. They found that reward rate was correlated with dopamine in nucleus accumbens core, but was not correlated in NAc shell or dorsal-medial striatum. Mohebi repeated these experiments using

the genetically encoded optical dopamine sensor, dLight [70], to measure dopamine concentrations. While VTA cell firing rate correlated well with NAc dopamine release associated with sensory cues (the light signaling the readiness of the apparatus, and the white noise burst signaling the go-cue), it did not correlate with NAc dopamine release associated with reward expectation during task performance.

Wolfram Schultz's seminal work that showed phasic dopamine signaling corresponds to both a reward prediction error in Pavlovian conditioning tasks and sensory impact suggests that dopamine plays a critical role in learning. Further studies have found that dopamine signaling in different task contexts and regions of the brain may represent a variety of different properties. Notably, dopaminergic signaling may not be controlled globally from VTA dopaminergic cell bodies, but regulated locally, perhaps by acetylcholine.

1.2 CNiFER cells as a method to measure neuromodulation in vivo

In this work, we ask if animals are able to volitionally control cortical neuromodulation. To probe this, we need a way to measure neuromodulator concentrations in real-time. Although a number of methods have been developed to perform this measurement, all of them have some drawbacks. The earliest method developed to measure neuromodulator concentration is microdialysis [59]. This involves inserting a microdialysis probe with a semipermeable membrane into the brain. Solutes, such as neuromodulators, are able to diffuse across this membrane. The solution in the probe is collected after some time,

and analyzed to determine solute concentrations. A microdialysis probe is able to simultaneously measure multiple neuromodulator concentrations with high sensitivity simultaneously. However, the large size of the probe and the low temporal resolution limit the ability of this method to detect real-time changes in neuromodulation.

Recently, fast scan cyclic voltammetry (FSCV) has emerged as an effective method to measure real-time changes in neuromodulator concentration [58]. This technique involves rapidly oscillating the voltage of a carbon fiber electrode inserted in the brain. When the voltage on the electrode is close to the redox potential of the compound to be measured, the compound is repeatedly oxidized and reduced. This is detected as a small oscillatory current; changes in neuromodulator concentration are detected as changes in amplitudes of this current. While FSCV is able to detect subsecond changes in neuromodulator concentration (each scan is typically done at >10Hz), it is unable to distinguish between compounds with similar redox potential. An important example of this in the brain is noradrenaline and dopamine, which have nearly identical redox potentials, but exert different effects in the brain. Thus, FSCV experiments to detect either of these compounds are typically done in regions that receive input from only one of these systems (i.e. dopamine measurements in striatum).

To address the issues with current techniques, our lab has developed CNiFER cells [57]. CNiFER cells are HEK-293 cells transfected with 2 genes; one that encodes the G-protein coupled receptor (GPCR) for the neuromodulator of interest, and another that encodes a fluorescent calcium probe (TN-XXL for

the CNiFER cells). Activation of the GPCR by the presence of the neuromodulator activates the intracellular IP3 pathway that releases calcium. However, this is only true for excitatory neuromodulators coupled to the Gq-protein. For other neuromodulators, where the receptor is bound to the Gi-protein, such as dopamine, it was necessary to transfect an additional gene that encodes a chimeric GPCR that binds to the Gq-protein instead of the Gi-protein. The increase in intracellular calcium will be detected as a change in the ratio of YFP to CFP fluorescence. The presence of calcium changes the conformation of the TN-XXL protein such that YFP fluorescence is increased and CFP fluorescence is decreased by increasing the efficiency of the FRET mechanism.

CNiFER cells possess superior temporal resolution compared to microdialysis; they are able to distinguish release events that occur within 3s of each other. Due to the fact that the receptors transfected into the CNiFER cells are the same as those in the brain, the sensitivity and specificity of the response should be comparable to neurons. Notably, CNiFER cells are able to distinguish dopamine from noradrenaline, two neurotransmitters that have nearly identical redox potentials, with a 30x difference in response to these two neurotransmitters.

Recently, genetically encoded membrane-bound fluorescent sensors for dopamine [56,70] have been developed. These sensors are based on cpGFP; the D2 receptor is modified to include cpGFP. Conformational changes in the receptor caused by binding of dopamine affects the shape of cpGFP, which modulates the efficiency of fluorescence. Compared to CNiFERs, these genetically encoded sensors possess superior temporal and spatial resolution.

Temporally, these sensors do not have to activate second messenger systems to release calcium, unlike CNIFERs. Thus, they are able to resolve changes in dopamine at subsecond resolution, unlike the CNiFERs, which have a lag time of approximately 3s. In addition, as the sensors are expressed on the membrane of neurons, and not a separate cell injected in the extracellular space, these sensors are capable of distinguishing changes in dopamine concentration from sites that are within 10um apart. CNiFER cells must be injected using a quartz pipette and it is thus impossible to separate injection sites by less than ~60um.

As these sensors are genetically encoded, transgenic animals can be bred such that minimally invasive monitoring of dopamine concentrations is possible. However, the effects of dLight or GRABDA expression are not known at this time. Pavlovian conditioning experiments in mice suggest that reward and cue related dopamine release in striatum remains functionally similar to that found in intact animals.

1.3 The brain machine interface as a method to probe volitional control

Understanding which portions of the brain are responsible for volitional control of behavior is best probed with operant conditioning [11]. The operant conditioning paradigm involves associating a reward or punishment with specific behaviors. If the subject is able to volitionally control its behavior, it should avoid punished behaviors and seek out rewarded behaviors.

In mammals, behavioral patterns are generated in the brain. In a study by Carmena et al [62], monkeys were trained to control a robotic arm. Electrodes

implanted into the motor cortex of the monkeys recorded neuronal activity while they used their arm to control a joystick. The joystick controlled a cursor on a computer screen, and the monkeys would receive a juice reward if they managed to move the cursor to an indicated point. During the initial training sessions, Carmena et al used the recorded electrical activity in motor cortex to train a linear model that predicted the cursor velocity from measured brain activity.

After training the monkeys to control the joystick physically to receive juice, the joystick was uncoupled from the cursor. The cursor position was calculated by applying the previously trained linear model to the measured brain activity. The monkeys were able to successfully perform this task as well, suggesting that brain activity in motor cortex was sufficient to drive control of the joystick towards the desired position.

While this is not entirely surprising, since the monkeys were still using their limbs to control the joystick, Carmena found that even when the joystick was removed, the monkeys were still able to perform the task without moving their limbs. This suggests that either motor cortex activity does not necessarily generate arm motion, or the monkeys were able to modify motor cortex activity so that it does not evoke arm motion.

Carmena's experiment used the monkeys' innate brain activity to control the cursor; the response of the cursor to the measured brain activity after removing the joystick did not differ from that before training. Thus it is unclear whether or not the monkeys' modified their own brain activity to match the task, or simply reinforced an existing pattern. We can extend this approach to probe

whether subjects can volitionally control brain activity directly, without relying on innate behavioral correlates to train them. We modify the operant conditioning paradigm such that instead of rewarding and punishing behaviors, we reward and punish measured patterns of brain activity. The interpretation of the results is exactly the same as that of an operant conditioning experiment. If the subject avoids punished brain activity patterns and seeks out rewarded brain activity patterns, it demonstrates the ability to volitionally control its brain activity.

Fetz tested this type of brain machine interface in monkeys [61]. Fetz recorded the activity of a single neuron in motor cortex with an electrode. He used the activity of this cell to condition the animal to increase the firing rate of this neuron. The firing rate of the neuron in question was conveyed to the monkey by either an audible click for each action potential or a meter where the reading varied with the firing rate. Notably, Fetz recorded from different neurons everyday, but after sufficient training, the monkeys were able to rapidly increase the activity of newly isolated cells. This suggests that the monkeys had learned a method to increase the firing rate of neurons. Fetz himself ruled out several systematic causes of the increase in firing rate, including injury-induced response, direct response to sensory feedback, and reward presentation itself. However it is not clear whether the monkeys learned to increase the firing rate of a single neuron or a large ensemble of neurons such that any tested neuron would show an increase in firing rate.

To distinguish between an increase in firing rate of multiple cells or just a single cell, a different experiment is necessary. Neely et al. demonstrated that

mice were able to volitionally control the firing rate in distinct neuronal ensembles in primary visual cortex (V1) using a brain machine interface paradigm [60]. Unlike Fetz, Neely et al rewarded animals for increasing the difference in firing rate between 2 different neuronal ensembles, rather than the firing rate of a single neuron. Therefore, an increase in the firing rate of a large set of cells would not constitute success in this task. Neely et al found that the animals were able to perform this task successfully. Unlike the neurons measured by Fetz and Carmena, the neurons in V1 are not innately responsible for behavior; instead they process visual information from the environment. In fact, Neely et al found that animals trained in lit conditions had to be retrained when put in a dark environment and vice versa.

Neely et al. demonstrated that animals are able to volitionally modify innate neuronal activity patterns to a novel activity pattern to receive a reward. Interestingly, the animals were unable to perform the task without an external cue; Neely et al used an auditory cursor, where the frequency of the tone was proportional to the firing rate. This suggests that the animals had no internal sense of the neuronal activity, but were able to modify it once they were aware of it.

Many experiments, in both humans and animals, have shown that brain activity can be modified and reinforced volitionally [60,61,62]. This is a remarkable example of neuronal plasticity and serves to illustrate the complexity of the brain's response to reinforcement. Interestingly, it seems animals are not able to perform most neurofeedback tasks without sensory cues related to the

task. The ability to volitionally control certain behaviors without reinforcing others may lead to some understanding of the underlying circuitry. To quote Fetz: “The degree to which one neural or muscular activity can be behaviorally dissociated from another could be a relevant test for causal connections between the underlying structures.”

Acknowledgements

Chapter 1 contain parts of a manuscript in preparation: Foo C, Lozada A, Aljadeff J, Wang J, Li Y, Slesinger P, Kleinfeld D, “Volitional control links spontaneous dopamine transients to reward” 2021. The dissertation author is the primary author of this work.

Chapter 2. Methods and Materials

2.1 Animal preparation

Adult, male C57BL/6 mice, age P30 to P45, were maintained in standard cages on a natural light-dark cycle. The Institutional Animal Care and Use Committee at the University of California San Diego approved all protocols. For surgery, mice were anesthetized with isoflurane (Butler Schein). Body temperature was monitored and maintained at 37°C. Subcutaneous injections of 5 % (w/v) glucose in saline were given every 2 h for rehydration. Buprenorphine (0.02 mg/kg, Butler Schein) was administered i.p. for post-operative analgesia.

2.2 Injection of mice for CNiFER measurements and/or $GRAB_{DA}$ expression

CNiFER cells were implanted for imaging as described in previous work [57]. $GRAB_{DA}$ expression was induced by injection of 0.5 - 1.0 µL of 10^{13}µg/ml of an AAV2 virus containing the $GRAB_{DA}$ sequence under the synapsin promoter [56]. The injections were made in frontal somatosensory cortex, 1.5 M/L, 1.5 R/C using a quartz glass pipette. At least three weeks were allowed for expression of the virus. A thin skull craniotomy [55] was made for animals that did not express $GRAB_{DA}$, whereas an open craniotomy was made for animals that expressed $GRAB_{DA}$. The craniotomies were 3 - 4 mm in diameter. In $GRAB_{DA}$ expressing animals, CNiFER cells were injected as close as possible to the $GRAB_{DA}$ expression area without overlapping or rupturing pial vessels, typically 60 - 100 µm away.

2.3 TPLSM imaging

In vivo imaging [54] was performed with a custom-built two-photon laser-scanning microscope [53]. Two-photon imaging was need to image through the thin skull to minimize the potential for inflammation consistent with high signal to noise ratio [54]. Control of scanning and data acquisition was achieved through the ScanImage software platform (Vidrio Inc.). Excitation light at 840 nm was used to excite the CFP portion of TN-XXL. Fluorescence was collected by a 25X dipping objective (HCX-IRAPO, Leica). The fluorescent signal was split into two channels using a 506 nm long-pass dichroic mirror. Each channel was further bandpass filtered: 465 ± 20 nm for measurement of emission by CFP($F_{465\ nm}$) and 520 ± 20 nm for emission by YFP ($F_{520\ nm}$) and the change in FRET signal, denoted ΔR(t)/R, was calculated as

$$\frac{\Delta R(t)}{R} = \frac{F_{520\ nm}(t) \,/\, \langle F_{520\ nm} \rangle}{F_{465\ nm}(t) \,/\, \langle F_{465\ nm} \rangle} - 1$$

where $\langle F_{520\ nm} \rangle$ and $\langle F_{465\ nm} \rangle$ refer to the baseline values determined from a LOWESS fit (see Data Analysis). Running activity was recorded using a rotary encoder (TRD-S360BD, Koyo Electronics) attached to the underside of a rotation platform (Figs. 2A and 3A). Licking activity was recorded using an optical lickometer (Sanworks LLC). The analog signals were recorded using a PowerLab 8.35 device (ADInstruments Inc.) acquiring at 1 kHz and synchronized with the two-photon imaging data using the frame trigger output of ScanImage.

2.4 Imaging of naïve animals

After one day of recovery from surgery, mice were water deprived (24h/day). Imaging began the following day. The animals were placed in a stationary head-frame fixed on top of a wheel; no lick port pas present nor was a reward or cues were given. The animals were allowed to locomote freely. Each imaging session lasted for one hour, after which the animals were returned to their home cages. Animals were given unrestricted access to water for one hour after imaging. Animals were imaged under identical conditions for four successive days.

2.5 Adaptive threshold feedback training

After one day of recovery from surgery, mice were water deprived (24h/day). Imaging began the following day. During the first day of imaging, the animals were placed in a stationary head-frame fixed on top of a wheel with a lick port. No reward or cues were given, and the animals were allowed to lick and run freely. For imaging trials with feedback reinforcement, animals were placed in the same head-frame, but given a reward, 0.1 mL of 10 % (v/v) sucrose water, for increasing their measured neuromodulator concentration. Real-time readout of neuromodulator concentration and reward administration were done using custom written ScanImage user functions. Each imaging session lasted for one hour, after which the animals were returned to their home cages. Animals were

given unrestricted access to water for one hour after imaging. Animals were imaged once a day for 4 - 6 consecutive days, depending on the optical quality of the CNiFER implants.

Rewards were dispensed using a gravity fed solenoid (VAC-100 PSIG, Parker Instrumentation) system. A 50 mL reservoir of 10 % (v/v) sucrose water was hung from a stand, and a solenoid was used to control the flow of liquid coming out of the reservoir. The output of this solenoid was routed to the input of the lick port. A second solenoid controlled the vacuum which was connected to the output of the lick port. The timing and control of both solenoids used a National Instruments board (USB-6211, National Instruments) interfaced with ScanImage. The initial threshold for reward was determined by taking 80 % of the average amplitude of spontaneous neuromodulator transients from the first day of imaging data, and the threshold increment was 50 % of this initial reward threshold. Reward administration occurred 0.25 s after the detection of the neuromodulator concentration rising above the threshold level, with a minimum delay of 3.5 s between rewards. At a time of 0.5 s after reward administration, a vacuum was activated for 1 s to remove any leftover sucrose water. The reward threshold was increased by the threshold increment whenever a reward was administered. If animals did not receive a reward within 225 s of the previous reward, the reward threshold was decreased by the threshold increment, down to a minimum of the initial threshold.

2.6 Data Analysis

All data analysis was done using MATLAB. TN-XXL fluorescence traces were normalized to baseline intensities measured from an initial period of 600s at the beginning of each imaging trial during which no reward was given. The neuromodulator concentration was calculated as the fractional change in the FRET ratio[3]. The FRET response was separated into a baseline and phasic component by performing a LOWESS fit on the data using a window size of 470 s and a step size of 4.7 s. The fitted curve was defined as the baseline (low-frequency) component, and the residual (high-frequency) was defined as the transient component. $GRAB_{DA}$ fluorescence traces were also normalized to baseline intensities measured from the initial 600 s period without reward. Baseline trends in the $GRAB_{DA}$ fluorescence traces were corrected using the same method as the FRET response; a LOWESS fit using the same parameters was subtracted from the trace and the residual was used.

Licks were detected from the lickometer data by low pass filtering the lickometer signal, subtracting a 3-sample median filter of the signal and detecting the subsequent rising edges. The licking frequency was calculated by taking a 1 s window around the sample point and counting the number of licks in the window. Running speed was similarly calculated from the encoder signal by detecting both rising and falling edges and counting the number of edges in a 0.235 s window around the sample point.

Significant transient were detected from the transient component of the FRET response as those epochs whose amplitude exceeded 2-times the root-mean-square level of the noise. The amplitude was calculated as the maximum

value in each detected transient epoch, and the width was the full width at half maximum amplitude (FWHM) of the transient.

The onset of neuromodulator transients was defined as the point at which the ratio ΔR/R increased by 2-times the root-mean-square level of the baseline noise. The onsets of licking and running bouts were calculated in a similar manner. Onset times were calculated separately for each imaging trial from the average transient-triggered response. The lags of the FRET signal behind the licking and running signals were defined as the difference in the onset time.

A linear model was used to fit the FRET response to the licking and running traces [52]. The transfer function of the model was calculated directly from the power spectra of the respective traces as the cross power divided by the input power. New models were fit for each imaging trial, and the variance explained, i.e., R^2, was calculated directly from the predicted FRET. Spectral power density and spectral coherence were calculated using multi-taper methods [51] with a time-bandwidth product of 10 and the Chronux package [50].

Acknowledgements

Chapter 2 contain parts of a manuscript in preparation: Foo C, Lozada A, Aljadeff J, Wang J, Li Y, Slesinger P, Kleinfeld D, “Volitional control links spontaneous dopamine transients to reward” 2021. The dissertation author is the primary author of this work.

Chapter 3. Probing volitional control of neuromodulation with a brain machine interface

In this work, we ask whether there are spontaneous dopamine transients in somatosensory cortex of mice in the absence of reward or environmental cues. If there are transients, we ask if mice can volitionally increase cortical dopamine levels. We also examine the relationship between cortical dopamine levels and locomotion, and ask whether training mice to increase their cortical dopamine levels affects this relationship. While we focus primarily on dopamine, further experiments were done on volitional control of noradrenaline and acetylcholine.

Rodent somatosensory cortex is innervated sparsely by dopaminergic dendrites from the midbrain. Most of these dendrites are located in the deeper layers of cortex (layer 4/5). Unlike striatum, somatosensory cortex also contains dense noradrenergic innervation [10]. As FSCV cannot distinguish between noradrenaline and dopamine, we decided to use CNiFERs to measure cortical dopamine and noradrenaline.

As previous studies have shown, midbrain dopaminergic signaling appears to code for either locomotion-related or reward-related signals, but not both [25]. However, as most dopaminergic terminals are not synaptic in nature and exert their influence via volume transmission [44], it is not clear how the cortex interprets changes in dopamine concentration that are a superposition of both locomotion and reward information. Further complicating the matter, dopaminergic release itself might be regulated locally [63,64] rather than

synchronized with action potentials from neuronal somata. Thus it is not clear what changes in cortical extracellular dopamine signals code for.

3.1 Spontaneous dopamine transients

We imaged surgically implanted D2-CNiFERs in frontal cortex of mice with two-photon scanning microscopy. The mice were headfixed on a wheel with a lickport placed in front of their mouths. The mice were free to run and lick. No reward or cues were given to the animal other than the sound of the microscope scan mirrors. We measure changes in cortical dopamine by detecting changes in the fluorescence of the D2-CNiFERs (Figure 1A).

We observe that even in the absence of cues or reward, there are spontaneous dopamine transients, occurring at a rate of 0.008 +/- 0.001 Hz (Figure 1D). Spontaneous dopamine transients varied greatly in width and amplitude, spanning 2 orders of magnitude. However, due to the nonlinearity of the in vitro D2-CNiFER sensitivity curve, dopamine transients reflect an increase in dopamine of only 1-10 nM despite spanning 2 orders of magnitude in FRET space (Figure 1E). Spontaneous dopamine transients were often preceded by anticipatory licking, but not locomotion.

Other studies have shown spontaneous burst firing of dopaminergic neurons in calm, awake rats [49]. The rate of large bursts found by Hyland et al is similar to our measured transient rate, suggesting that the dopamine transients we measure are derived from neuronal somata action potentials. However,

further work must be done (e.g. simultaneous measurement of VTA dopaminergic cells and cortical DA) to determine whether this is the case or not.

3.2 Feedback training of mice

After 1 day of imaging to observe spontaneous dopamine transients, animals were given reward for increasing their cortical dopamine level. The measured dopamine level was compared, in real time, to an adaptive threshold. Dopamine levels in excess of the threshold were rewarded, and the threshold was increased. The threshold was decreased if the animal did not receive a reward for a set period of time (Figure 2A, see Methods for details). No auditory or visual cues were given to the animal to indicate the real-time dopamine level.

After 3 days of feedback training, animals demonstrated the ability to volitionally increase cortical dopamine (Figure 2C). Notably, both the basal dopamine level and dopamine transient frequency increased as the threshold was increased (Figure 3). These changes were abolished with the cessation of reward and recovered when reward was restored. Dopamine transients when the animal was given reward lasted significantly longer than transients when the animal did not receive reward (Figure 3C). Interestingly, while the amplitude of rewarded transients was significantly greater than the amplitude of spontaneous transients, the amplitude of unrewarded transients under feedback training was not significantly different than the amplitude of spontaneous transients.

3.3 Timing of motor response relative to onset of dopamine transients

The licking activity and running speed of animals were monitored concomitantly with the cortical dopamine level. The onset of licking relative to the onset of dopamine transients shifted over the 4 days of feedback training. In untrained animals, anticipatory licking preceded dopamine transients by 11.0+/-3.5s. When the animal was rewarded for increasing cortical dopamine, the onset of licking shifted to lag behind the onset of dopamine transients by 1.4+/-3.9s. Licking began to precede the onset of dopamine transients once again when feedback was removed (Figure 3D). The variance explained by a linear model fitted to predict licking behavior from dopamine release increased over training, from 0.13+/-0.06 in naïve animals to 0.19+/-0.07 on Day 3 of training, indicating that dopamine release better predicted licking activity later in training.

In contrast, running bouts did not consistently evoke, nor were they evoked by, dopamine transients (Figure 5). For trials where running was correlated with dopamine transients, the onset of running relative to the onset of dopamine transients did not shift significantly with training. Fitting a linear model to predict cortical dopamine from running speed, dopamine did not correlate strongly with running speed and the correlation did not change significantly over training. The variance explained by this model remained 0.10+/-0.05 throughout training.

3.4 Randomly distributed reward

Dopamine transients are inherently tied to reward and reward-predicting cues. It is possible that the presentation of sugar water on its own would increase cortical dopamine concentrations, thereby creating a positive feedback loop. Presentation of sugar water, the reward, would increase dopamine concentrations, triggering the presentation of more reward (as the presentation of reward is contingent on increases in dopamine concentration). To ensure that this is not the case, and the animals were not learning some systematic cue from our experimental setup (i.e. learning the sound of the solenoid opening for water reward), we rewarded animals randomly over 4 days rather than for increased dopamine levels. The random rewards were Poisson distributed with a mean inter-reward interval equal to the mean dopamine transient frequency for animals trained on feedback.

Animals given reward randomly did not show an increase in basal dopamine levels nor in the dopamine transient frequency (Figure 3). While random reward did reliably elicit dopamine transients, these transients always lagged behind the onset of licking by 14.3+/-5.8s. Dopamine release in randomly rewarded animals did not predict licking behavior well; the variance explained by our linear model was only 0.03+/-0.01.

3.5 Effect of lick port presence on dopamine release

It is possible that the presence of the lick port itself served as a reward cue and thus drove spontaneous dopamine release. To ensure that this was not

the case, we measured cortical dopamine levels in animals that were head-fixed but did not have a lick port placed in front of their mouths. Like the other conditions, animals were imaged for 4 consecutive days while water deprived.

We found that spontaneous dopamine transients still persisted even in the absence of the lick port (Figure 5). Furthermore, the rate of transients was not significantly different from that in the presence of the lick port. This suggests that the lick port itself did not serve as either a motivator or cue for reward.

Interestingly, we found that, in the absence of the lick port, animals ran more than in the presence of the lick port. This is likely because the animal did not want to run into the lick port; to ensure the animal could receive reward from the lick port, the lick port was positioned close to the mouth and could easily be sensed by the animals' whiskers. Further, while running activity was coherent with dopamine release in the presence of the lick port at low frequencies (< 0.05 Hz), in the absence of the lick port, running was not coherent with dopamine release (Figure 1C). In the absence of the lick port, animals engaged in long running bouts, and dopamine transients occurred whether or not the animal was running. The rate of dopamine transients during running did not significantly differ from the rate when the animals were still.

3.6 Genetically encoded dopamine sensors

Recently, a genetically encoded fluorescent probe for dopamine, known as GRABDA, has been developed [56]. We tested these probes by first expressing them in frontal cortex, then injecting D2-CNiFERs 60um away from

the injection site. The fluorescent signal from the GRABDA and the D2-CNiFERs show good agreement (Figure 6). GRABDA shows superior temporal resolution and sensitivity compared to the D2-CNiFERs, but was not stable over the training period. We were unable to complete the full 4 day training session using GRABDA; the GRABDA signal did not show significant fluctuations by day 3 of training. Whether this is due to surgical technique or optical damage to the neurons is unclear.

Spontaneous dopamine transients were observed simultaneously in both the D2-CNiFER implant and the GRABDA implant. The durations of these transients were shorter as measured using GRABDA (Figure 6). The average FWHM of GRABDA transients was 18.8s compared to 33.2s as measured with the D2-CNiFERs. However, the peak response of the D2-CNiFER lagged behind the GRABDA response by only 1.2s; the primary contribution to the width of the D2-CNiFER response is the slow decay after the peak. Thus, timing of DA release as measured by the D2-CNiFERs should not differ significantly from that measured with GRABDA. Additionally, the average spontaneous dopamine transient rate was 0.008 Hz, corresponding to an inter-transient interval of 125s. As this interval is significantly longer than the decay time of transients measured using D2-CNiFER, we are unlikely to have failed to detect transients.

3.7 Other neuromodulators - noradrenaline

We trained animals to increase their cortical noradrenaline in the same way as we did for dopamine. Animals were given a reward if their cortical

noradrenaline reached a target level, which would increase as the animal performed the task. Animals were able to significantly increase their basal noradrenaline levels (Figure 8). Noradrenaline transient frequency was significantly increased with feedback training ($p = 1e-11$). Interestingly, this enhancement persisted even when reward was withheld.

Unlike dopamine transients, noradrenaline transients were not modified by feedback training, except after reward was withheld. Neither transient duration ($p = 0.64$, 0.15, and 0.0003) nor transient amplitude ($p = 0.08$, 0.49, and 0.005 for Days 2, 3, and 4 respectively) changed significantly during feedback training, with the exception of Day 4. On Day 4, after feedback was withheld, transient amplitude was significantly larger and transient width significantly shorter than in naïve animals.

Noradrenaline transients were time-locked to increases in running velocity. The onset of running bouts did not significantly shift relative to noradrenaline release during feedback training. The onset of licking shifted from preceding noradrenaline release by 14.0+/-7.0s in naïve animals to lagging noradrenaline release by 14.4+/-14.1s on the second day of feedback training ($p = 0.05$). However, by Day 3 of feedback training, this effect was abolished, and licking shifted to preceding noradrenaline release once again.

Linear models fit to predict running and licking from noradrenaline transients did not explain a significant portion of the variance in running and licking. However, the variance explained by models for licking significantly increased on days with reward, and not on Day 4 when reward was withheld. In

comparison, the variance explained by models for running did not change significantly with training.

3.8 Other neuromodulators – acetylcholine

We applied the same training paradigm from sections 2.2 and 2.6 to acetylcholine. Animals were injected with M1-CNiFERs in somatosensory cortex and trained to increase cortical acetylcholine levels. Surprisingly, we found that animals were unable to significantly increase basal acetylcholine levels (Figure 9). Neither basal acetylcholine levels nor acetylcholine transient frequency increased significantly with training.

Acetylcholine transients were significantly shorter after Day 3 of feedback training and remained shorter on Day 4 after feedback was removed. However, the amplitude of acetylcholine transients did not change significantly except on Day 4 after feedback was removed. On Day 4, the amplitude of transients was significantly less than in other conditions.

Acetylcholine transients were time-locked to the onset of running bouts in naïve animals. This relationship remained consistent throughout feedback training. The difference in onset time did not significantly change over feedback training; running lagged acetylcholine by 13.4+/-3.6s. An increase in licking frequency was observed concurrent with the onset of running associated with acetylcholine transients. As with running, the lag of licking frequency increases behind acetylcholine transients did not change significantly with feedback training, remaining around 12.4+/-6.5s.

Linear models fit to predict acetylcholine from running and licking were able to explain a larger fraction of the variance compared to models for noradrenaline and dopamine. However, unlike noradrenaline and dopamine, the variance explained by these models did not change with feedback training.

Unlike dopamine, acetylcholine transients did not change significantly with feedback training. Transient width did not significantly differ across feedback training. Transient amplitude also did not increase significantly with feedback training. These results suggest that acetylcholine transients are not under volitional control. As with transients, basal acetylcholine levels did not increase significantly over feedback training. These results combined suggest that animals were unable to increase acetylcholine levels to receive reward.

Acknowledgements

Chapter 3 contain parts of a manuscript in preparation: Foo C, Lozada A, Aljadeff J, Wang J, Li Y, Slesinger P, Kleinfeld D, “Volitional control links spontaneous dopamine transients to reward” 2021. The dissertation author is the primary author of this work.

Chapter 4. Discussion

In this work, we have shown that rodents are capable of volitionally increasing both cortical dopamine and noradrenaline levels to receive a reward. In particular, animals are able to both ramp up basal cortical dopamine concentrations and increase transient frequency to match an increasing threshold value for reward. Strikingly, unlike other brain-machine interface experiments, animals learned to ramp their cortical neuromodulator levels without continuous sensory feedback associated with the brain activity we trained on, using only discrete liquid rewards.

Our feedback scheme for reinforcement learning is modeled after schemes used in brain-machine interface (BMI) experiments [60,61,62]. Yet we did not supply cues related to the feedback signal other than the reward. In typical BMI experiments, animals are provided a real-time sensory feedback signal, e.g. an auditory tone [60], that is modulated by the neuronal activity of interest. In particular, one BMI study found that animals were unable to modulate the neuronal activity towards the reward target in the absence of such cues [60]. In another study that involved learning to move a cursor based on neuronal activity in motor cortex, monkeys first needed to learn the task using their hands to control a physical joystick that moved a cursor on the screen. They then transferred this capability to pure control by neuronal activity [62]. Unlike these BMI studies, we did not supply cues related to the feedback signal other than the reward (**Figure 2A**). Yet animals were able to perform our task and receive

reward. This suggests that animals have an internal sense of $[DA]_{ex}$, most likely derived from the normal functioning of DA-GCPRs on cortical cells.

We observe spontaneous dopamine impulses in the absence of both sensory stimuli and reward in naïve, head fixed, awake animals on a treadmill. These dopamine impulses do not appear to report a reward prediction error, as there are no cues in the apparatus that would indicate the presence of reward. We found that, in naïve animals, dopamine transients were preceded by licking in the presence of a dry lick port, perhaps suggesting a role of dopamine transients in the anticipation of reward in naïve animals. However, on Day 3 of feedback training, dopamine transients preceded licking activity. This shift of timing indicates that the animals were able to use the spontaneous dopamine transients as a cue to predict rewards and were further able to initiate more transients to receive more reward. Although the transient frequency increased over training, basal dopamine concentrations also increased over the length of the trial. Although there was a weak correlation between dopamine transients and running speed, this correlation did not change significantly over training days, suggesting that motor behavior was not responsible for driving increases in dopamine.

Moreover, we found that, in the absence of a lick port, running bouts were completely uncorrelated with observed dopamine transients. Running bouts were much shorter when the lick port was present than when the lick port was absent. This is an intuitive result; no one would continue running forwards if the presence of an object is sensed ahead. This suggests that the presence of an object near the mouth changes the relationship between dopamine transients and motor

activity, and furthermore, there is no relationship between the onset of running bouts and dopamine transients in somatosensory cortex. Although this might seem to contradict studies [23,41] that show correlation between running and dopamine, these studies focused on dopamine release in striatum, where dopaminergic terminal density is high. Dopaminergic projections to somatosensory cortex are sparse and largely confined to deeper layers (Layer 5/6). Previous studies have shown that dopamine release may be locally modulated [63,64]. Therefore, it is possible that dopamine release detected in somatosensory cortex might not be correlated to that in striatum. While we only measure licking and running behavior, it should be noted that dopamine appears to invigorate rather than initiate motor behavior [27,41], and spontaneous motor behavior is, in fact, associated with a reduction in the firing rate of midbrain substantia nigra pars compacta neurons in the absence of cues or reward [23, 25].

In contrast to dopamine impulses, tonic dopamine appears to affect how animals interact with their environment and explore. In particular, manipulations of tonic levels of $[DA]_{ex}$modified how animals examined novel stimuli in various operant conditioning tasks [31-34] and the effort that animals were willing to exert to receive reward [26,30]. It is not clear whether tonic and phasic dopamine signals are independently controlled from one another [24,28]. A previous study [29] found that explicitly modeling tonic dopamine as a separate signal in the reward prediction error signal could reproduce the results of dopamine depletion on the response rate of rats rewarded on a fixed ratio reward schedule. While the

authors used the average reward rate as an estimate of tonic dopamine[29], they note "... *we should expect the tonic average reward signal to be used predictively and not only reactively, which would require it to be somewhat decoupled from the actual obtained phasic reward signal*". Therefore, our results suggest that tonic dopamine can be volitionally modulated independently of sensory stimuli.

The increase in basal dopamine level and transient frequency are unlikely to be an inherent response of the animals to reward presentation. When the animals were given reward at random time intervals (Poisson distributed around the mean inter-reward interval in trained animals), we did not observe an increase in basal dopamine level or an increase in transient frequency. Furthermore, the increase in basal dopamine level was only significant on Day 3 of feedback training. If there was an inherent ramping of dopamine in response to reward presentation, we would expect a similar ramp to be present on Day 2 of feedback training, when the animals were also rewarded for increasing dopamine. Furthermore, animals that were trained for 3 extra days (reward was restored on Day 5 and Day 7, and withheld on Day 6) showed ramping of basal dopamine levels similar to that on Day 3 in the first session after reward was restored. The temporal delay and quick restoration of basal dopamine ramping in response to reward suggests that the animals had to learn the task before they could perform it, and were able to remember how to perform the task later.

It is interesting, however, that extinction of the learning occurred rapidly, within a single trial. We noticed no ramping of dopamine in trials with feedback off. Perhaps failure to receive the expected reward early in the trial suppresses

dopamine release, consistent with dopamine dynamics in striatum [63,66]. Ramping was restored when feedback was reintroduced the next day, suggesting that the animals retained knowledge of how to increase dopamine from trial to trial. Perhaps multiple shorter trials within a single day might shed some light on how this occurs. However, the ramping we observe in response to reward presentation occurs on a time scale of 1000s, which would be the lower bound of trial length.

We were also able to successfully train animals to increase cortical noradrenaline levels to receive reward. However, the variability in animals trained to increase noradrenaline levels was larger than that in animals trained to increase dopamine levels – while some animals were able to significantly increase cortical noradrenaline, other animals were unable to perform the task at all. Perhaps cortical noradrenaline transients, unlike dopamine, are not naturally perceived as a cue for reward in all animals, preventing animals from understanding the task. Repeating the experiment with a sensory cue, such as an auditory cursor that indicates the cortical noradrenaline level, might decrease variability in this case.

Noradrenaline transients were not consistently time-locked to the onset of licking bouts, and were only weakly correlated with licking. However, the correlation of noradrenaline transients with licking frequency increased significantly above that in naïve animals in rewarded trials (Days 2 and 3). This is to be expected; the animals will lick when presented with reward, and we are pairing noradrenaline with reward. Consistent with this result, the correlation of

noradrenaline transients with licking frequency decreased back to naïve conditions on Day 4, when no reward was given. We also see a shift of licking from preceding noradrenaline to lagging noradrenaline only on Day 2 of training. This shift might indicate that animals are able to use noradrenaline transients as cues for reward. However, this shift was not present on Days 3 and 4 of feedback training. This might be due to the increase in basal noradrenaline levels. It is known that high levels of noradrenaline inhibits task performance [16]; by Day 3 of training it is possible that elevated levels of basal noradrenaline inhibits the ability to use noradrenaline transients as a reward cue.

Noradrenaline transients tended to precede increases in running speed, but not all transients elicited motion. The lag of running behind noradrenaline transients did not shift significantly with training, suggesting that the correlation of noradrenaline release with running did not change with feedback training. In addition, the average running speed of animals did not change significantly with feedback training, despite the fact that the frequency of noradrenaline transients increased significantly with training. This suggests that animals were not simply increasing their motor activity to increase noradrenaline.

Unlike dopamine transients, noradrenaline transients did not change in width or amplitude with feedback training. However, the frequency of noradrenaline transients did increase during feedback training relative to naïve animals. This suggests that animals were unable to modulate the duration or intensity of noradrenaline transients, but were able to increase the frequency of noradrenaline transients.

If noradrenaline transients were not cues for reward, perhaps animals did not modulate noradrenaline transients, but only increased basal noradrenaline. As noradrenergic neurons in locus coeruleus also express noradrenergic receptors to modulate their own activity, increases in basal noradrenaline should modulate the frequency of transients. However, pharmacological manipulation of basal noradrenaline [46] suggests that increasing basal noradrenaline would decrease the frequency of noradrenaline transients. This suggests that the increase in the frequency of noadrenaline transients was due to animals volitionally initiating more noradrenaline transients, rather than simply a side effect of increasing basal noradrenaline levels. However, it is unclear whether the cortical noradrenaline concentration reflects the concentration in locus coeruleus. Experiments done with an auditory cursor indicating noradrenaline levels would elucidate whether animals could modify and volitionally initiate noradrenaline transients to use as a cue for reward.

Although animals were able to significantly increase cortical noradrenaline and dopamine, we were unable to train animals to increase cortical acetylcholine significantly. Acetylcholine transients were strongly time-locked to motor activity. Both running bouts and increases in licking frequency were accompanied by acetylcholine transients. During the course of feedback training, animals did not significantly increase their basal acetylcholine levels, and the amplitude of acetylcholine transients remained constant. Furthermore, the frequency of acetylcholine transients did not change with feedback training.

Acetylcholine has been implicated in control of dopamine release in striatum. It is interesting to see that animals cannot control cortical acetylcholine despite being able to control cortical dopamine. This might be due to the faster timescale of acetylcholine reuptake compared to that of dopamine or noradrenaline. Cholinergic neurons also tend to form more conventional synapses [42-45] than dopaminergic and noradrenergic neurons. In addition, as neither noradrenaline nor dopamine transient amplitudes increased significantly with feedback training, it is likely that transient amplitudes are difficult for the animal to control. Thus, not only will ramping of basal extracellular acetylcholine due to changes in tonic firing rates be more difficult to detect, but the amplitude of acetylcholine transients might not be flexible enough to follow the adaptive threshold. Similar to noradrenaline transients, it might also be the case that acetylcholine transients cannot be used as a cue for reward without external sensory feedback.

Experiments that train the animal to directly increase the tonic firing rate of cholinergic neurons in nucleus basalis would elucidate whether we are simply unable to detect ramping of extracellular acetylcholine or the animal is unable to increase the tonic firing rate. As with noradrenaline, further experiments done using an auditory cursor to indicate acetylcholine levels would show if animals could learn to associate acetylcholine transients with reward.

Conclusion

In summary, we have found that animals are able to volitionally increase extracellular dopamine and noradrenaline, but not acetylcholine in somatosensory cortex to receive a reward. All 3 neuromodulators studied displayed unique responses to our training paradigm. The neuromodulators studied also differed in how they related to licking and running behavior.

Although animals predominantly increased their basal dopamine levels to match the ever increasing threshold for reward during training, dopamine transients also changed with feedback training. Dopamine transients in naïve animals were preceded by the onset of licking bouts. After training, dopamine transients preceded licking bouts, suggesting that animals were able to use these spontaneous dopamine transients as a cue to receive a reward. Although dopamine transients were weakly correlated with running, this correlation did not change significantly over feedback training, suggesting that the animals were not simply performing some motor behavior to increase dopamine.

Dopamine transients after feedback training lasted significantly longer than those in naïve animals but were not larger in amplitude. Thus, we have shown that animals are able to use spontaneous dopamine transients as a cue to receive reward. We further show that they are able to ramp up their dopamine levels to match an adaptive reward threshold by increasing their basal dopamine levels, presumably by increasing the tonic firing rate of dopaminergic neurons.

Animals were able to increase cortical noradrenaline levels to receive a reward, but did not seem to be cognizant of noradrenaline transients in the same

way as they were of dopamine transients. Notably, the relationship between noradrenaline transients and motor behavior did not change significantly with feedback training. Although noradrenaline transient frequency did change with feedback training, it is unclear if this is simply a side effect of an increase in basal noradrenaline or a conscious effort by the animal.

We found that animals were not able to increase basal acetylcholine levels in the same way as noradrenaline and dopamine levels. Acetylcholine transients were time-locked to the onset of running bouts and increases in licking frequency. Unlike noradrenaline and dopamine, acetylcholine transients did not change significantly with training. Whether this is due to the lack of increase in basal acetylcholine levels or simply an inability to volitionally control transients is unknown.

Acknowledgements

Chapter 4 contain parts of a manuscript in preparation: Foo C, Lozada A, Aljadeff J, Wang J, Li Y, Slesinger P, Kleinfeld D, “Volitional control links spontaneous dopamine transients to reward” 2021. The dissertation author is the primary author of this work.

Appendix 1. Pupil diameter and noradrenaline during feedback training

Previous studies [22] have shown that calcium in axons of neuromodulatory neurons is positively correlated with pupil diameter. Increases in both noradrenergic and cholinergic axon calcium concentrations precede pupil dilations in awake mice. Reimer et al found that both noradrenaline and acetylcholine release precedes small rapid dilations when the animal was still. The correlation between noradrenaline and pupil dilation was higher than that between acetylcholine and pupil dilation. Initiation of locomotion was preceded by pupil dilation. This dilation was preceded by a phasic burst of noradrenaline, followed by a longer period of cholinergic release that varied linearly with pupil diameter.

Pupil diameter also seems to indicate choice in a rewarded discrimination task [21]. During a texture discrimination task, Lee et al found that large pupil dilation accompanied licking in hit and false alarm responses, while smaller dilation followed by constriction followed a correct rejection response. In particular, large pupil dilations only occurred when the animal decided to lick (i.e. the animal believes it will receive a reward), and not if the animal did not lick. Other studies have found pupil diameter changes correlate with reward uncertainty [14]. We ask whether pupil dynamics, associated with both noradrenaline release and reward expectation, change when we pair noradrenaline with reward.

We examine whether or not training animals to volitionally modulate noradrenaline affects pupil dynamics. We modified the brain machine interface

apparatus to include a fast infrared camera (Basler acA1300-200um), and a dim white LED behind a light diffuser. The LED and diffuser were placed such that the emitted light would illuminate the animal's eye. The 2 photon excitation light incident on the brain scattered through the pupil and the pupil appeared as a bright spot on the infrared camera image. The pupil diameter was estimated by thresholding the camera image, and calculating the diameter of a disk that would have the same area as the detected bright spot. The animal was trained for 7 days this time, as opposed to 4. Feedback training continued for 3 days (Days 2, 3, and 4). Feedback was removed on Day 5, restored on Day 6, and removed again on Day 7.

We find that, in the naïve animal and during early training (Days 1 to 3), tonic noradrenaline corresponded well with tonic pupil diameter – increases in tonic noradrenaline led to an increase in tonic pupil diameter and vice versa (Figure A1). However, tonic noradrenaline did not correspond with tonic pupil diameter on Days 4 and 5, when pupil diameter was at its maximum. Notably, while a slow increase in noradrenaline levels was observed on Day 4, tonic pupil diameter did not change significantly. Similarly, a slow decrease in noradrenaline level on Day 5 did not elicit any changes in tonic pupil diameter. Days 6 and 7 (after reward was withheld on Day 5), the pupil remained large at the start of the trial. However, decreases in tonic noradrenaline were associated with significant pupil constriction unlike Day 4. Increases in tonic noradrenaline on Day 6 correlated with a decrease in pupil constriction amplitude, but the tonic pupil diameter did not increase. Interestingly, the animal's tonic pupil diameter

remained large even after feedback was removed on Days 5 and 7; we did not observe tonic noradrenaline levels increase during these trials.

While the relationship between tonic noradrenaline and pupil diameter did not remain consistent through training, noradrenaline transients were reliably correlated with pupil dilation. However, the temporal relationship and magnitude of response varied over training. In the naïve animal, pupil dilations preceded noradrenaline transients and were larger than dilations in trials with feedback training. With training, pupil dilation lagged noradrenaline transients and was diminished in magnitude. The decrease in magnitude is likely due to the increase in tonic pupil diameter with training; in later training trials, the pupil was large, essentially remaining maximally dilated. It was thus physically impossible for the pupil to dilate much further.

While the tonic noradrenaline concentration did not correlate well with tonic pupil diameter in later training trials, the animal's average running speed was well correlated with tonic pupil diameter throughout feedback training. We find that a linear regression between the tonic pupil diameter for a trial and the average running speed for the same trial fit the data well, with an R^2 of 0.76.

Noradrenaline transients were reliably preceded by pupil dilations in the naïve animal. This might seem inconsistent with previous studies [1] that showed that noradrenaline axonal calcium spikes precede pupil dilation by approximately 335ms. However, the temporal lag of the peak CNiFER response from neuromodulator application is approximately 7s, and we observe a peak in the cross-correlation of pupil diameter with noradrenaline around 7s; thus we can not

conclude that pupil dilations precede noradrenaline transients as our measurement of noradrenaline is dominated by the CNiFER response time.

Interestingly, after feedback training, pupil dilation lagged noradrenaline transients. This suggests that the mechanism by which noradrenaline is coupled to the pupil diameter is plastic and can be volitionally controlled by the animal. It is unlikely that this shift in timing was caused by increases in tonic noradrenaline, as the shift was already present on Day 2 of feedback training, before the tonic pupil diameter began to increase.

While noradrenaline transients were always accompanied by pupil dilation, pupil dilations were not always associated with noradrenaline transients. This suggests that noradrenaline transients are sufficient but not necessary to drive pupil dilation. There are likely a variety of different mechanisms that control pupil diameter (i.e., the pupillary light response), some of which would not require noradrenergic input.

A key limitation of the CNiFER technology is the inability to measure absolute neuromodulator concentrations. Changes in fluorescence can only be measured relative to the fluorescence at the beginning of a trial. Thus, the variability in the tonic pupil diameter would not be reflected in our noradrenaline measurements if the noradrenaline level varied from day to day during feedback training.

The breakdown of the relationship between pupil diameter and tonic noradrenaline during feedback training suggests that noradrenaline can be driven to a level that saturates the pupil dilation response. It is likely that the tonic

noradrenaline level is proportional to both the tonic pupil diameter and average running speed. Specifically, average running speed is linearly correlated with tonic pupil diameter with an R^2 of 0.76, suggesting that both running and pupil diameter might be good measures of tonic noradrenaline in our paradigm. If we assume that the tonic pupil diameter at the beginning of each trial is proportional to the tonic noradrenaline level at the beginning of each trial, the highest tonic noradrenaline level occurred on Days 4 and 5. On these days, the pupil has been driven to saturation; thus neither increases nor decreases of tonic noradrenaline affect the pupil dynamics. On Days 6 and 7, where the tonic noradrenaline is lower than Days 4 and 5, the pupil is close to the saturation point. Thus, the pupil responds readily to decreases in noradrenaline, and only slightly to increases, as further increases beyond the saturation point would not elicit a response. Unlike tonic pupil diameter, the average running speed showed no saturation response, peaking on Day 4, and decreasing from Days 5 to 7.

Interestingly, even after feedback was removed on Day 5, the tonic pupil diameter remained elevated compared to the naïve animal, despite a lack of ramping tonic noradrenaline. Notably, pupil diameter decreased slightly on Day 5, and continued to decrease on Days 6 and 7, even though the animal was rewarded on Day 6 and successfully increased its noradrenaline level. This would suggest that the animal learned to increase tonic noradrenaline relative to the start of trial, and did not seem to control the absolute concentration of noradrenaline.

On average, animals trained on noradrenaline did not show a significant increase in running speed over feedback training. While we did not record pupil diameter from these animals, if the relationship between running speed and pupil diameter is consistent, we can conjecture that pupil diameter was not increased significantly with training in these animals. However, running speed varied significantly between animals. This discrepancy suggests that there are multiple strategies for increasing cortical noradrenaline to perform our task. Furthermore, strategies to learn the task, as well as the learning process, might be altered by variations in tonic noradrenaline concentrations at the beginning of each trial that CNiFER technology is not sensitive to.

Acknowledgements

Appendix 1 contains unpublished work done in collaboration with Adrian Lozada, Paul Slesinger, and David Kleinfeld. The dissertation author is the primary author of this work.

Appendix 2. Effect of optogenetic stimulation of locus coeruleus on vasodynamics

Noradrenaline is known to act as a vasoconstrictor in the peripheral nervous system by binding to α1-adrenergic receptors [5,6]. However, previous studies on the effect of noradrenaline on cortical vasculature show mixed results, with some studies [4] reporting an increase of cerebral blood flow, while others reporting a decrease in blood flow [1].

These studies measured changes in blood flow using intrinsic optical imaging [1] or laser Doppler flowmetry [4], methods that are unable to distinguish individual blood vessels. Furthermore, these studies used different methods to increase noradrenaline. Bekar et. al. injected an α2a-antagonist, increasing the basal firing rate of noradrenergic neurons, while Toussay et al used electrical stimulation of locus coeruleus to stimulate phasic noradrenaline release. Here, we examine how optogenetically stimulated noradrenaline release affects pial vessel diameter in somatosensory cortex by directly imaging vessels with 2-photon microscopy and simultaneously measuring noradrenaline release with α 1a-CNiFER cells.

We first injected 0.5-1 µL of an AAV2-FLEX-rev-ChR2-mCherry virus (titer of 10^{13}) in the locus coeruleus of TH-Cre mice stereotaxically. After waiting 3 weeks for the virus to express, a thin skull window was made over somatosensory cortex, and α1a-CNiFER cells are injected into the window as in the main text. An optic fiber cannula (Doric Lenses, 200 µm core, 0.48 NA) is then inserted into the brain over the locus coeruleus. After 1 day of post-surgery

recovery, we inject 0.2 µL of dextran-conjugated tetramethylrhodamine (100 kDa) in the ocular vein. We then use our TPLSM apparatus to simultaneously image blood vessel diameter and CNiFER implants during optogenetic stimulation of locus coeruleus. A fiber coupled laser (Coherent OBIS, 445 nm) was coupled to the optic fiber cannula using a press fit sleeve, and was triggered by a 5V TTL pulse. Pulse trains of 100 pulses at 100Hz, with a pulse width of 0.5ms were fed into the laser with an ADInstruments PowerLab board, and used to stimulate the locus coeruleus. The laser intensity at the fiber optic tip was measured to be 4.7mW.

To increase the time resolution of the data and reduce the size of data files stored, we used an arbitrary scan path rather than a raster scan of the field of view. A box path was traced over CNiFER implants to capture as much fluorescence over the implant as possible. Lines perpendicular to blood flow were traced over blood vessels to measure changes in vessel diameter. The fluorescence intensity along the box path was thresholded, and the average intensity of pixels above threshold was used as the intensity of the CNiFER implant. Cortical noradrenaline levels were calculated from these intensities as in the main text (see Methods). Blood vessel diameters were calculated by finding the two coordinates along the line that minimizes the variance in both the exterior and interior of the coordinates (see Appendix 3.1).

Blood vessels responded in an inconsistent manner to ChR2 stimulation of locus coeruleus. Although we measured a robust and significant increase in cortical noradrenaline with stimulation, vessel diameter did not consistently

decrease. In fact, vessels within the same field of view would respond in different ways to the stimulation (Fig A2). Notably, we observed 4 distinct vasoactive responses. 4 of the 23 vessels measured constricted in response to the stimulation. On the other hand, 4 other vessels dilated in response to stimulation, and 14 vessels showed either an inconsistent or no response to stimulation. 1 vessel showed a unique response, rapidly dilating in response to stimulation and then constricting.

The consistent increase in noradrenaline observed with stimulation in all animals and trials suggests that both the CNiFERs and the optogenetic stimulation are functioning as intended. However, the lack of a consistent vasodynamic response could be due to a myriad of factors. It is known that the craniotomy procedure will induce pathological vasomotion [38,39,40]. While we used a thin skull preparation, a small incision still had to be made to inject the α1a-CNiFER cells. Perhaps this incision caused a decoupling of noradrenaline and vasomotion in vessels near the incision.

Although noradrenaline acts as a vasoconstrictor directly on smooth muscle, it also regulates neuronal activity [4]. Activation of neurons from noradrenaline release could initiate vasodilation rather than vasoconstriction. Neuronal responses to noradrenaline might differ based on the animals' brain state and sensory environment. If the spatial pattern of neuronal activation initiated by noradrenaline release changes from one brain state to another, it stands to reason that the neurons, and by extension, the blood vessels, in a given field of view would respond differently to locus coeruleus stimulation. Thus,

changes in brain state might underlie the inconsistent response to locus coeruleus stimulation.

Other studies have implicated noradrenaline as a factor in astrocytic endfeet activation [2]. Bekar et al showed that noradrenaline release elicited a robust increase in astrocytic calcium concentrations. As astrocytes are known to extend processes onto blood vessels [3], it is possible that some vessels that respond to noradrenaline release are not responding to neuronal activity or noradrenaline, but rather astrocytic activity associated with noradrenaline release.

Further complicating issues, the absolute noradrenaline concentration at the start of a trial, a variable the CNiFERs are not sensitive to, might change from day to day, as suggested by pupil diameter measurements and changes in average running velocity (see Appendix 1). As with the pupil, it is possible that vessels are only sensitive to transient noradrenaline release within a certain range of tonic noradrenaline, or respond differently to transients at different tonic noradrenaline levels. These considerations would extend to both astrocytic and neuronal contributions to the vasodynamic response as well.

Acknowledgements

Appendix 2 contains unpublished work done in collaboration with Paul Slesinger and David Kleinfeld. The dissertation author is the primary author of this work.

Appendix 3. Computational analysis methods for general purpose use

Over the course of my studies, I have created many coding subroutines for analysis of various types of data. While not all are original, they have proven useful to others in the lab. I have added a brief description of each subroutine; the code files are uploaded as supplementary files to this thesis.

A3.1 Method for determining blood vessel diameter when signal to noise ratio is low (otsuDiameter.m)

This routine calculates the diameter of a blood vessel scanned with an arbitrary linescan or frame scan when the signal to noise ratio is low. We use the interclass variance (Otsu's criteria) to separate the interior of the vessel from the exterior. We assume that the center of mass of the vessel profile after subtracting the median is inside the vessel. We then iterate over all possible vessel boundaries and determine the boundaries that maximize the variance between the exterior points and the interior points. Unlike the full-width half maximum method, which relies on the profile's shape, this method utilizes statistics from all pixels in the profile to determine vessel diameter and is thus less susceptible to noise or vessel asymmetry, provided there are enough pixels in a scan (Figure A3).

A hyperparameter, **a**, is introduced to enhance performance when the vessel profile is highly asymmetric, dye is only localized to the vessel wall, or there are background features not relevant to the vessel. **a** ranges from 0 to 1 and controls the weighting of interior variance compared to exterior variance.

When **a** = 0, the algorithm seeks only to minimize the variance of the interior rather than exterior. When **a** = 1, the algorithm minimizes the variance of the exterior instead. In practice, **a** >= 0.5 will tend to work better, as the dark region around the vessel tends to be more homogenous than the vessel itself (this is especially true when only the vessel wall is labeled.

A3.2 Fast spectral analysis on high spatial resolution data (svdSpectralDat.m)

In widefield imaging, data can often have a large number of pixels. If there are oscillatory components to the signal of interest (e.g. calcium oscillations) that vary across the spatial dimension, it is often time-consuming to perform spectral analysis on each pixel, especially if multi-taper methods are used [CHRONUX]. While it is possible to decrease computational time by binning pixels, this results in a reduction of the spatial resolution.

Here I propose a method to dramatically speed up spectral analysis of wide-field imaging data using the SVD. The key factor is the linearity of the Fourier transform operation – the Fourier transform of a superposition of temporal modes is the superposition of the Fourier transforms of the individual modes. In our data, the number of time points is much less than the number of pixels. Thus, decomposing the data into spatial and temporal modes, then performing Fourier analysis on the temporal modes rather than on each pixel will reduce the computational time by approximately the ratio of pixels to time points, assuming the SVD can be calculated efficiently.

Spectral properties can then be recalculated for each pixel by multiplying the frequency modes (calculated from the temporal modes) by the spatial modes. For more complex computations, such as the global coherence, multi-taper spectral power, or coherence, further calculations with the modes must be done as these are nonlinear operations. However, the computational load is still significantly less as these properties can still be calculated using only the resultant frequency modes.

In practice, not all the temporal modes need to be included in the calculation (as the number of significant modes is far less than the rank of the matrix), and the computational runtime is significantly less (50-70x) than with conventional methods.

I have written several subroutines to perform multi-taper analysis of wide field imaging data, using Chronux. Subroutine 1 calculates the global coherence of the data across all pixels at a desired frequency. Subroutine 2 calculates the multi-taper estimate of the power spectra for all pixels. Subroutine 3 calculates F-ratio statistics, amplitude and phase of a deterministic sinusoid of a desired frequency embedded in the data.

A3.3 Efficient methods to access ScanImage data (scanMap.m; tiffMap.m)

Imaging data in our lab is acquired with the ScanImage software. However, the resultant data sets are often too large to fit into memory, and have custom header information encoded in the files. I have written methods to quickly access relevant subsections of data for analysis to save on memory, as well as extract

header information that describes acquisition parameters. There are 2 subroutines here. The first subroutine allows fast access to ScanImage data acquired from frame scan acquisitions, stored as tiff image stacks, while the second subroutine allows fast access to ScanImage data acquired from arbitrary scan acquisitions, stored as binary data files.

A3.4 Method to dispense reward and record video time-locked to ScanImage acquisition for real-time feedback control (SI_fbController)

For my thesis experiment, I had to write software to interface with both the lick port and ScanImage to control reward delivery in real time in response to changes in our imaging signals. In addition, functionality for simultaneous video recordings of pupil diameter was added at a later date.

A3.5 Ultrafast acquisition of video data using USB3 Basler cameras (baslerController)

Certain behaviors occur on fast timescales, such as whisking. To detect whisking behavior and compare it with neurological correlates requires the use of an ultrafast camera. This routine allows for streaming of ultrafast video from a USB3 Basler camera to disk.

Figures.

Figure 1. Characterization of spontaneous transient changes of $[DA]_{ex}$ in naive mice in the absence of reward.

Mice can run, lick, and otherwise move freely within the constraints of head-fixation.

A. Open-loop measurement of cortical $[DA]_{ex}$. Increases in $[DA]_{ex}$ are observed as an increase in YFP fluorescence with a concomitant decrease in CFP fluorescence, and vice versa. Inset: Image of CNiFER fluorescence in cortex.

B. Time series for spontaneous transients in $[DA]_{ex}$ along with the corresponding licking behavior.

C. Lick-triggered $[DA]_{ex}$, based on the peak of the licking signal, and lick-triggered licking,for the data in panel B.

D. The sample-averaged spectral coherence magnitude (top) and phase (bottom) for all animals (18 mice) between the licking rate and $[DA]_{ex}$;dashed line is the 0.95 confidence level. The roughly linear phase-shift (red line) suggests a time-lag can characterize the dynamics between the two signals.

E. Standard boxplot showing the lag of $[DA]_{ex}$ transients behind licking onset for all animals; the red line is the median.

F. Two-dimensional histogram of spontaneous transients in $[DA]_{ex}$ across all animals.Top row is the cumulative of all widths. The average width was 15.1 ± 1.3 s (blue line). Right column is the cumulative of amplitudes; 0.1 corresponds to $[DA]_{ex}$ ~ 20 nM. The average amplitude was 0.056 ± 0.002 (blue line).

G. Distribution of spontaneous DA transient rate for all animals. The average rate was 0.08 ± 0.01 Hz.

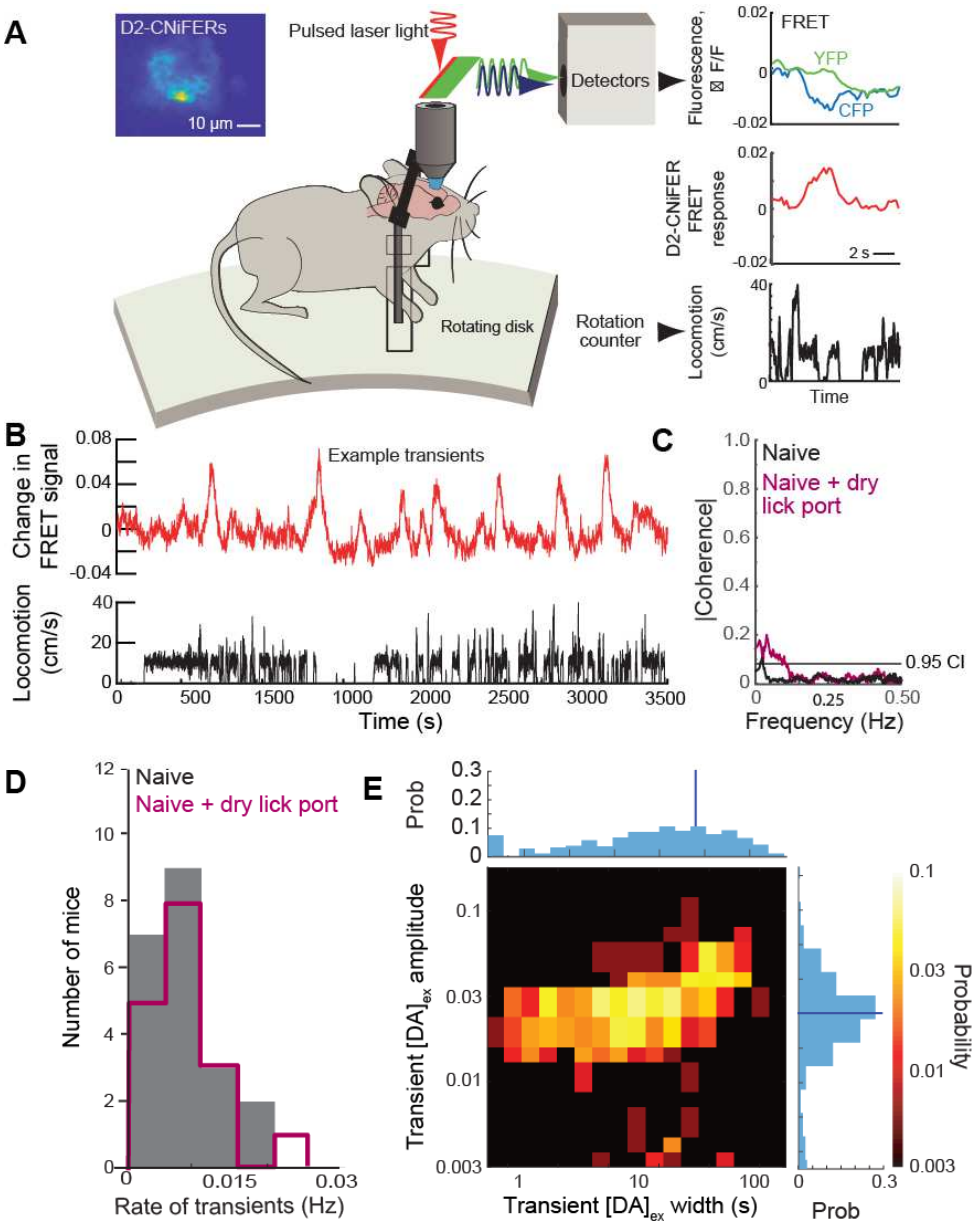

Figure 1. Foo, Lozada, Li, Wang, Slesinger* & Kleinfeld*

Figure 2. Feedback reinforcement training to volitionally link spontaneous transients in $[DA]_{ex}$ to reward.

A. The set-up is the same as in the open loop experiment (**Figure 2A**) plus a reward is supplied from the lick port. The closed-loop feedback uses the measured D2-CNiFER responses as a proxy for $[DA]_{ex}$ and drives the delivery of a liquid reward based on the $[DA]_{ex}$ signal (red) relative to an adaptive threshold (green). If $[DA]_{ex}$ exceeds the threshold, 0.1 mL of sucrose-water is released and the threshold in incremented by 0.005 signal units. The value of the threshold also exponentially decreases, in discrete steps of 0.005, with a time-constant of 225 s, since the last increment.

B. Example data shows the open-loop response on Day 1 and a ramping of the basal $[DA]_{ex}$ with closed-loop reinforcement on Day 3; the adaptive threshold follows the reward to within the decay time-constant. A 350 s snip of data highlighted by the beige band is expanded on the right.

C. Compendium rolling average of the transient $[DA]_{ex}$ for all animals (9 mice). The averaging window was 235 s. Feedback reinforcement training was activated on Day 2; $[DA]_{ex}$ ramped upwards yet did not significantly increase from that of naive mice ($p = 0.69$). By Day 3 the ramp in$[DA]_{ex}$ was significantly ($p = 0.01$). The increase extinguished when reward was withheld on Day 4 ($p = 0.73$), and rescued when reward was restored on Day 5/7 ($p = 0.02$) compared to feedback off on days 6/8.

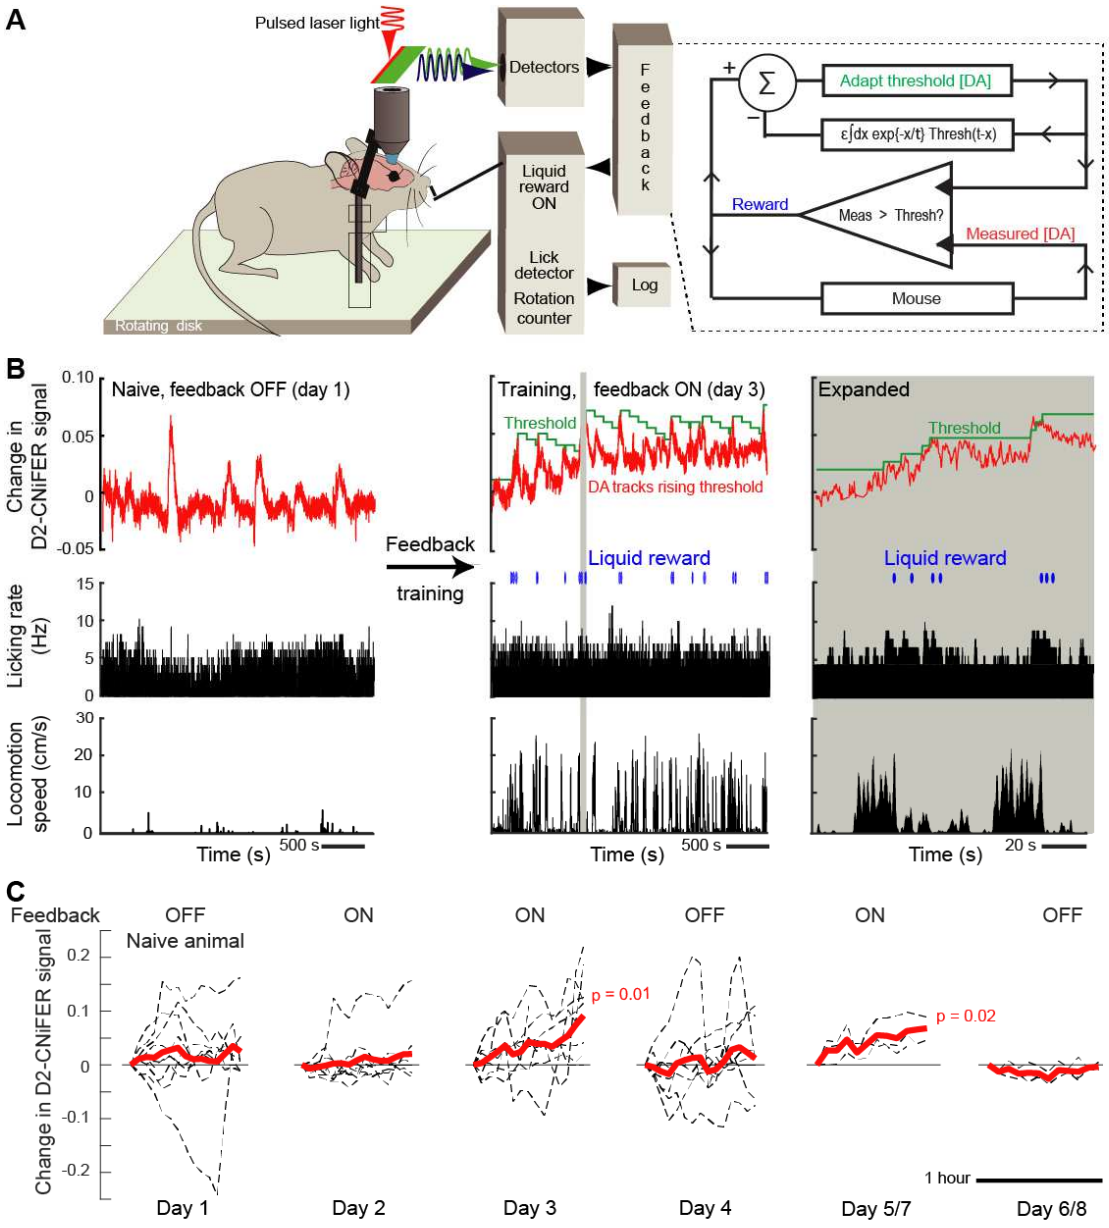

Figure 2. Foo, Lozada, Li, Wang, Slesinger* & Kleinfeld*

Figure 3. Population averaged changes in $[DA]_{ex}$ over the course of our reinforcement paradigm.

A. Tonic $[DA]_{ex}$ shows a significant increase relative to naïve, baseline animals after feedback training (Day 3; $p = 0.006$); the increase is $[DA]_{ex} = 0.088$. Day 2 and Day 4 showed no significant change in tonic $[DA]_{ex}$, $p = 0.37$ and 0.66, respectively. Randomly rewarded mice also did not show a significant change in tonic $[DA]_{ex}$ compared to naïve animals ($p = 0.61$).

B. The amplitude of $[DA]_{ex}$ impulses significantly increased over the course of the trial with feedback training (Day 3; $p = 0.01$). With feedback OFF, there was a significant decrease in amplitude over the course of the trial (Day 4; $p = 0.004$).The data for Day 1, Day 2, and the case of random reward showed no significant change in $[DA]_{ex}$ impulse amplitude over time; $p = 0.43$, $p = 0.77$, and $p = 0.67$ respectively.

C. The rate of $[DA]_{ex}$ impulses was significantly higher relative to baseline animals with feedback training (Day 3; $p = 0.001$). Day 2 and Day 4 showed no significant change in $[DA]_{ex}$ impulse frequency, $p = 0.65$ and 0.37, respectively. The rate for animals with random reward was significantly less than that for naïve animals ($p = 10^{-8}$). Alternating bars along time axis indicate intervals for binned data in panels A to C.

D. Distribution of amplitudes and widths of $[DA]_{ex}$ impulses for trained mice with feedback ON (Day 3; top), a subset of rewarded impulses when feedback is ON (Day 3; center), and when feedback reinforcement is OFF (Day 4; bottom). The amplitudes with feedback training ON were significantly larger compared to baseline animals (Day 3; $p = 10^{-7}$). The widths were significantly longer with feedback ON compared to those for baseline animals ($p = 10^{-22}$). Rewarded impulses had an amplitude of 0.078 ± 0.002 and an average width of 44.2 ± 1.6 s and were significantly longer and larger than the impulses when feedback was OFF (Day 4; $p = 10^{-8}$ and $p = 10^{-29}$, respectively).

E. Standard boxplot shows the timing of the onset of $[DA]_{ex}$ impulses relative to the onset of licking; the solid bars are the median times. The mean and standard error of the lag times of $[DA]_{ex}$ impulses relative to licking are $+11.0 \pm 3.5$ s on Day 1, $+0.6 \pm 0.9$ s on Day 2, -5.1 ± 1.4 s on Day 3, with $[DA]_{ex}$ now leading, and$+22.8 \pm 11.9$ s on Day 4. Impulses also lagged licking for randomly rewarded mice ($+10.9 \pm 7.9$ s). A two-sample t-test with respect to baseline animals yields $p = 0.04$ (Day 2), $p = 0.005$ (Day 3), $p = 0.9$ (Day 4), and $p = 0.3$ (random). The insert shows a scatter plot of the number of rewarded impulse versus lag time for individual mice on Day 2 and Day 3, together with a fit of a leaky integrator model to the data.

F. Standard boxplot shows the variance explained by an optimal linear filter that predicts licking from the measured $[DA]_{ex}$. "+" indicates an outlier excluded from the mean. Solid bars indicate median. The values of R^2 for Days 2 and 3 are significantly different than zero, with $p = 0.030$ and $p = 0.037$. The data for random reward, although very small at 0.019, is also significant ($p = 0.027$, $n = 12$) as a result of the large sample.

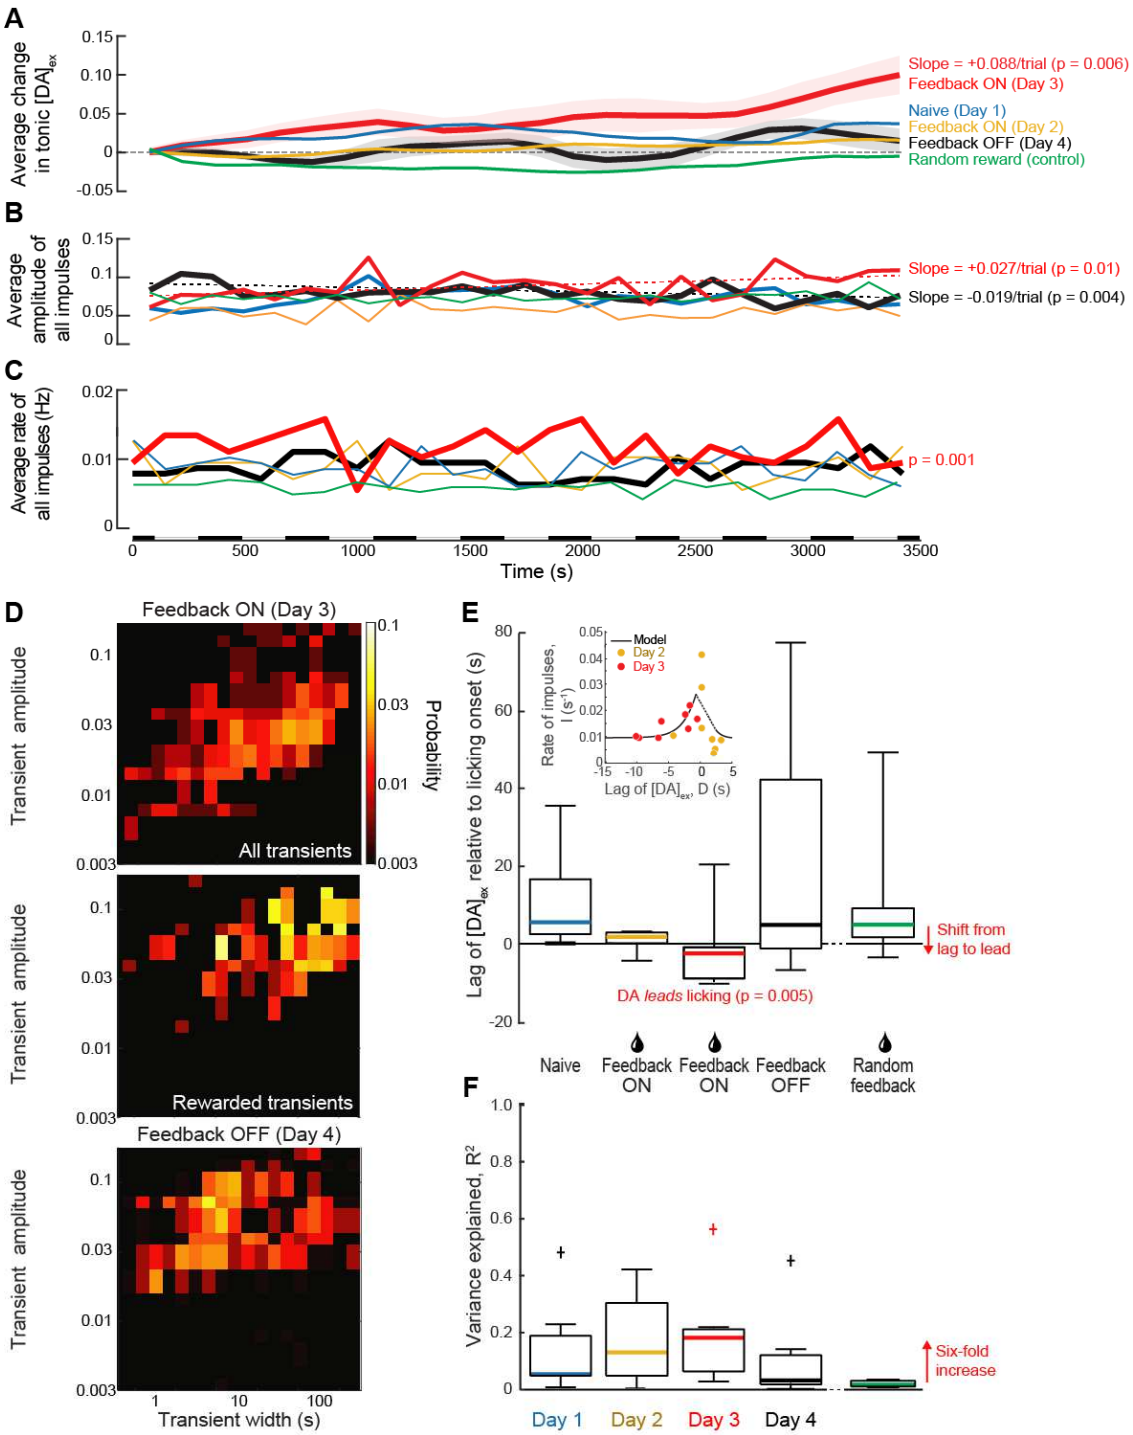

Figure 3 revised. Foo, Lozada, Aljadeff, Li, Wang, Slesinger* & Kleinfeld*

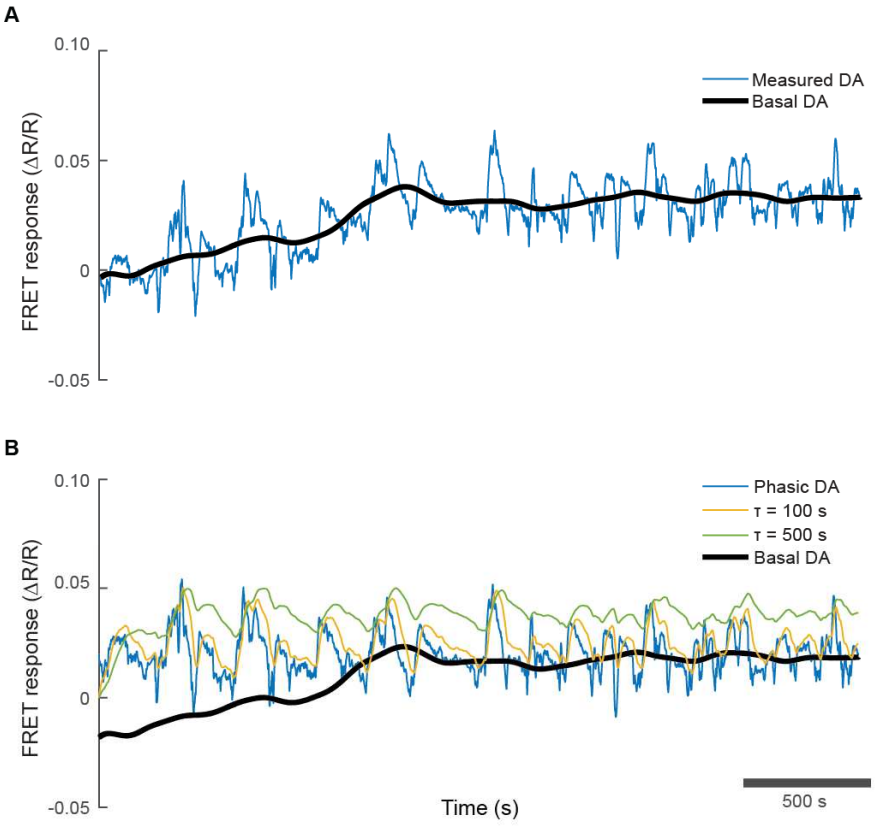

Supplemental Figure S1. Foo, Lozada, Li, Wang, Slesinger & Kleinfeld

Figure 4. Leaky integration of dopamine transients is not sufficient to explain increases in basal dopamine levels.

A. Extraction of basal DA from the measured $[DA]_{ex}$. A LOESS fit (tricubic weighting function, linear fit) was applied to the measured signal (blue) to extract DA transients. The window size was 940 s and the step size was 11 seconds. The transients were subtracted from the total DA signal to get the basal DA (black).

B. Integration of the phasic DA signal does not reproduce the ramping basal DA signal. The phasic DA signal (blue), was extracted from the measured $[DA]_{ex}$ using a LOESS fit. A leaky integrator, with exponential decay time τ, was applied to this signal (yellow, green) for values of $\tau = 100$s and 500s. The leaky integrator fails to reproduce the shift in basal DA that we observe. Although the integrator with a half-decay time of 500 s shows a similar shift upwards in the DA response, it ramps up to this level much quicker, i.e., around 100 s, than the 1000 s that we observe (black line).

Figure 5. Two-dimensional histograms of the amplitude and width of dopamine transients over feedback training.

A. Two-dimensional histograms showing the change in dopamine transient properties over training. Transient amplitudes during Day 3 of feedback training were significantly larger than those in the naïve animal. Transient widths during both Day 2 and 3 of feedback training were significantly longer than those in the naïve animal. Transient properties when feedback was turned off on Day 4 did not significantly differ from the naïve animal. The average widths of transients were 15.1 ± 1.3 s, 25.4 ± 2.1 s, 43.1 ± 2.5 s, and 18.5 ± 1.7 s; the corresponding average amplitudes were 0.056 ± 0.002, 0.045 ± 0.003, 0.081 ± 0.004, and 0.056 ± 0.002 for Days 1, 2, 3, and 4, respectively.

B. Two-dimensional histograms showing dopamine transient properties when animals were randomly rewarded. Transient width was significantly shorter on Days 3 and 4 compared to Days 1 and 2. Transient amplitudes were lower when animals rewarded compared to when they were not. The average widths of transients were 22.4 ± 2.1 s, 24.6 ± 2.3 s, 13.1 ± 1.3 s, and 14.7 ± 1.6 s; the corresponding average amplitudes were 0.052 ± 0.003, 0.029 ± 0.002, 0.027 ± 0.0005, and 0.039 ± 0.002 for Days 1, 2, 3 and 4 respectively.

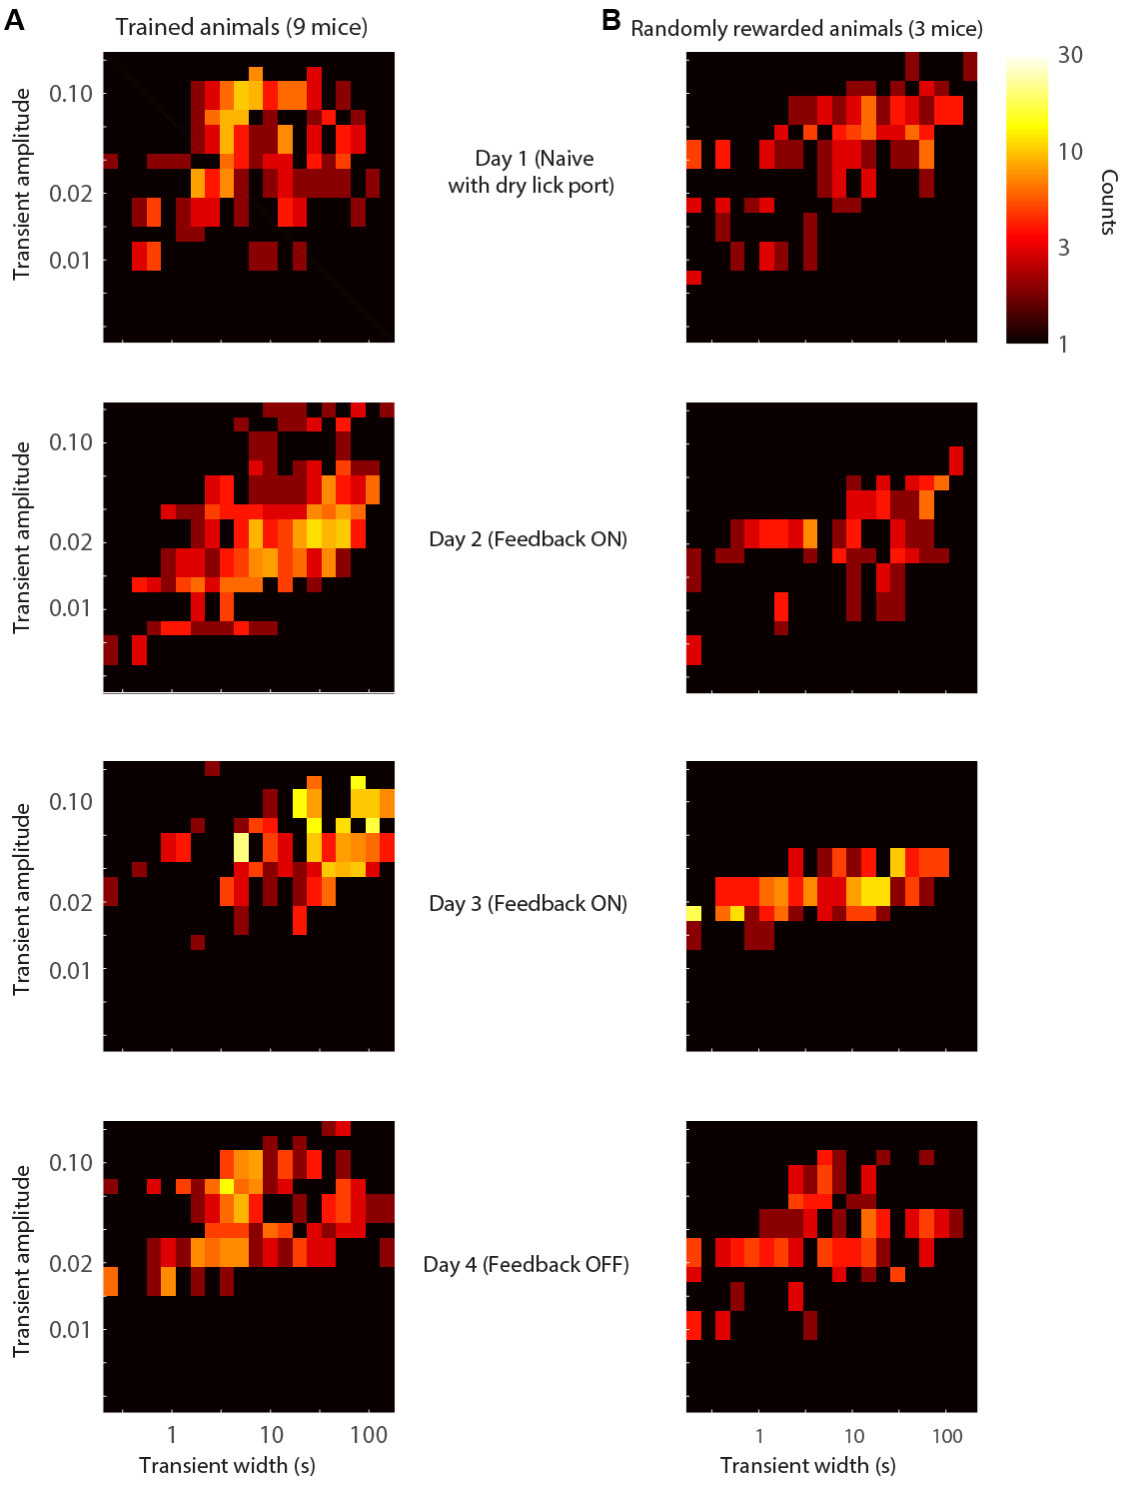

Supplemental Figure S2. Foo, Lozada, Li, Wang, Slesinger* & Kleinfeld*

Figure 6. Introduction of a dry lick port introduces a small correlation between running and dopamine release.
A. Average spectral coherence between running and phasic dopamine release across animals (9 mice with lick port, 7 mice without lick port). In the presence of a lick port, coherence was significant at frequencies below 0.2Hz. In the absence of a lick port, coherence was not significant. The coherence was calculated using the multi-taper method; the bandwidth was 0.02 Hz from averaging with 143 tapers.
B. Histograms of the predictions of a linear model of $[DA]_{ex}$ as a function of running speed versus the measured $[DA]_{ex}$ in the absence of a lick port during four consecutive days of experiments. White line shows the expected distribution of a perfectly predictive model. The model was fit to the data in the frequency domain, making use of the convolution theorem. Cross-spectral power was calculated with a multitaper estimate. A new model was fit for each trial; each histogram uses data from all trials within a given day of the experiment.
C. Same as panel B, but in the presence of a lick port. Animals were trained to increase $[DA]_{ex}$ for these data.
D. Variance explained by linear model of $[DA]_{ex}$ as a function of running speed in the absence of a lick port for different days of the experiment. This was calculated directly from the data shown in panel B. Each trial was a separate data point. R^2 was 0.011 ± 0.005, 0.009 ± 0.002, 0.010 ± 0.005, and 0.009 ± 0.003 for Days 1, 2, 3, and 4 respectively.
E. Same as panel D, but in the presence of a lick port. R^2 was 0.1 ± 0.1, 0.08 ± 0.1, 0.2 ± 0.2, and 0.04 ± 0.07 for naïve, Day 2 of training, Day 3 of training, and feedback OFF days respectively.

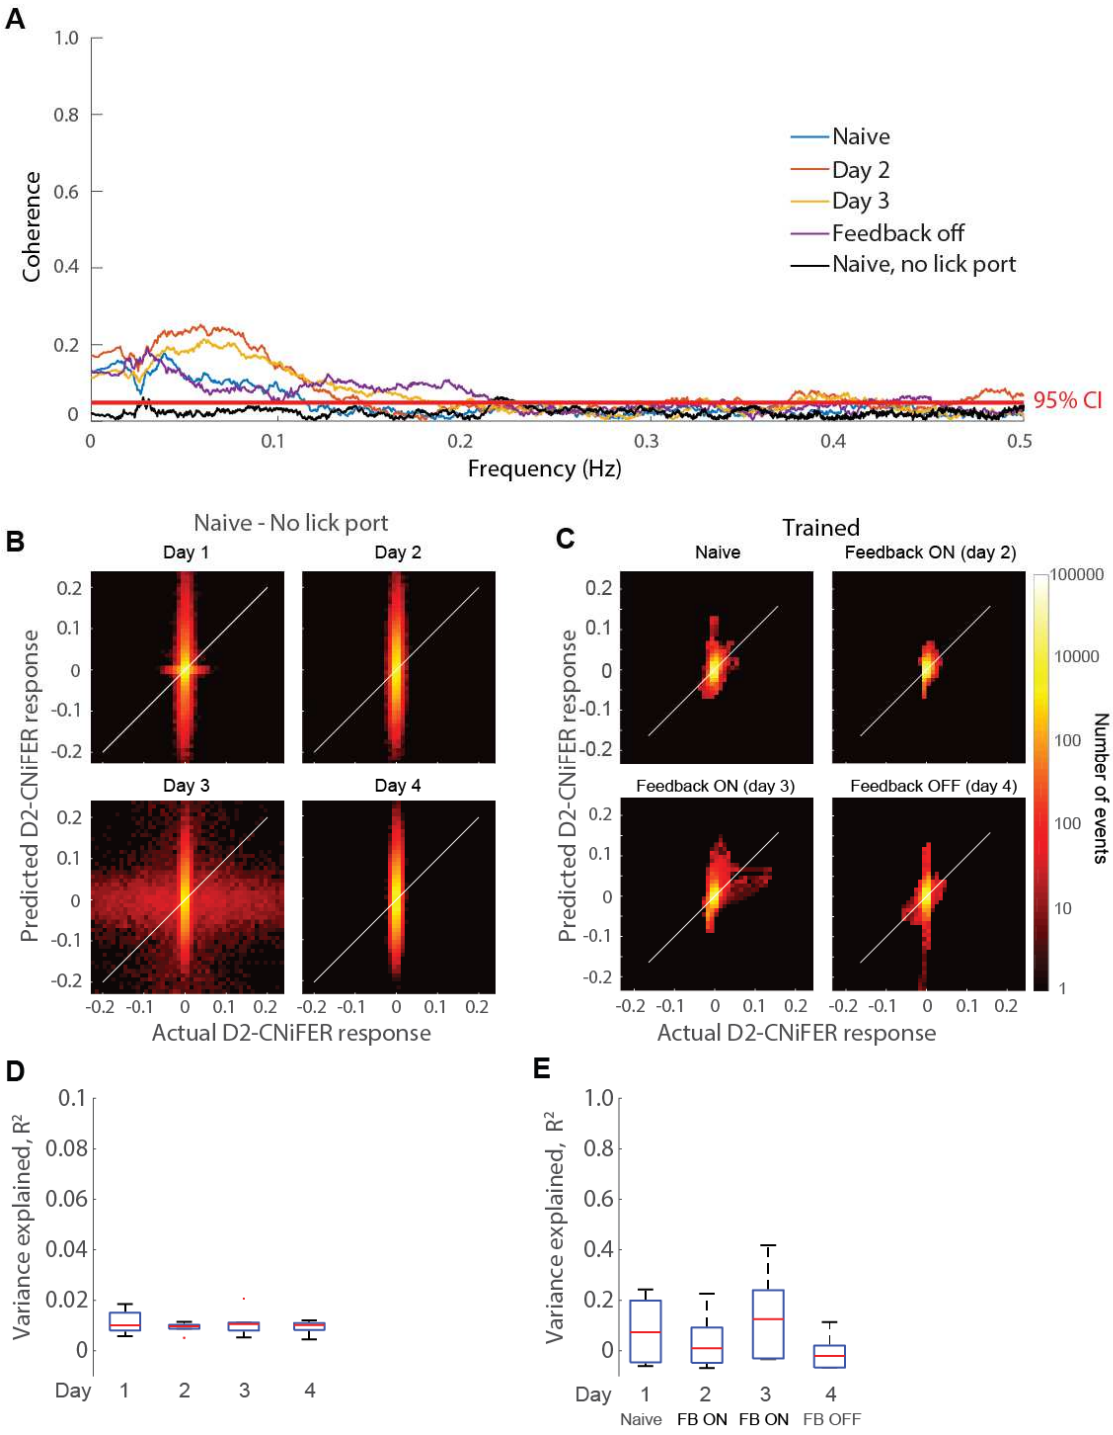

Supplemental Figure S3. Foo, Lozada, Li, Wang, Slesinger & Kleinfeld

Figure 7. Comparison of signals from genetically expressed $GRAB_{DA}$ and implanted D2-CNiFER cells.

A. Schematic of design of the genetically expressed dopamine sensor, $GRAB_{DA}$, compared to D2-CNiFERs. $GRAB_{DA}$ is constructed by inserting the dopamine binding site on the D2-GPCR into cpGFP. Binding of DA causes conformational changes in the binding site that effect the efficiency of fluorescence.

B. Averaged image showing region of $GRAB_{DA}$ expression and a D2-CNiFER implant. $GRAB_{DA}$ expression was induced using a viral vector.

C. Simultaneous measurement of genetically expressed $GRAB_{DA}$ (blue, top) and implanted D2-CNiFER cells (green, middle). A small region of interest near the center of the region of $GRAB_{DA}$ expression was averaged and a fluorescence trace was calculated. The $GRAB_{DA}$ signal had significant drift in baseline on the scale of tens of minutes, but had better signal-to-noise and temporal resolution than the D2-CNiFER signal.

D. Comparison of normalized detrended $GRAB_{DA}$ signal and normalized D2-CNiFER signal. The $GRAB_{DA}$ signal did not exhibit the decay tail that the D2-CNiFER signal had, and was about twice as bright; small transients that were detected by $GRAB_{DA}$ were not always detected by the D2-CNiFERs. Transients occurring in quick succession as observed by $GRAB_{DA}$ appeared as a single, longer transient when observed by D2-CNiFERs.

E. Normalized average transient triggered response of $GRAB_{DA}$ (blue) and D2-CNiFER (green) signals. The $GRAB_{DA}$ signal both rose and decayed more rapidly than the D2-CNiFER signal, as one would expect; the change in fluorescence in the D2-CNiFER signal requires activation of a second messenger pathway that is not necessary in $GRAB_{DA}$.

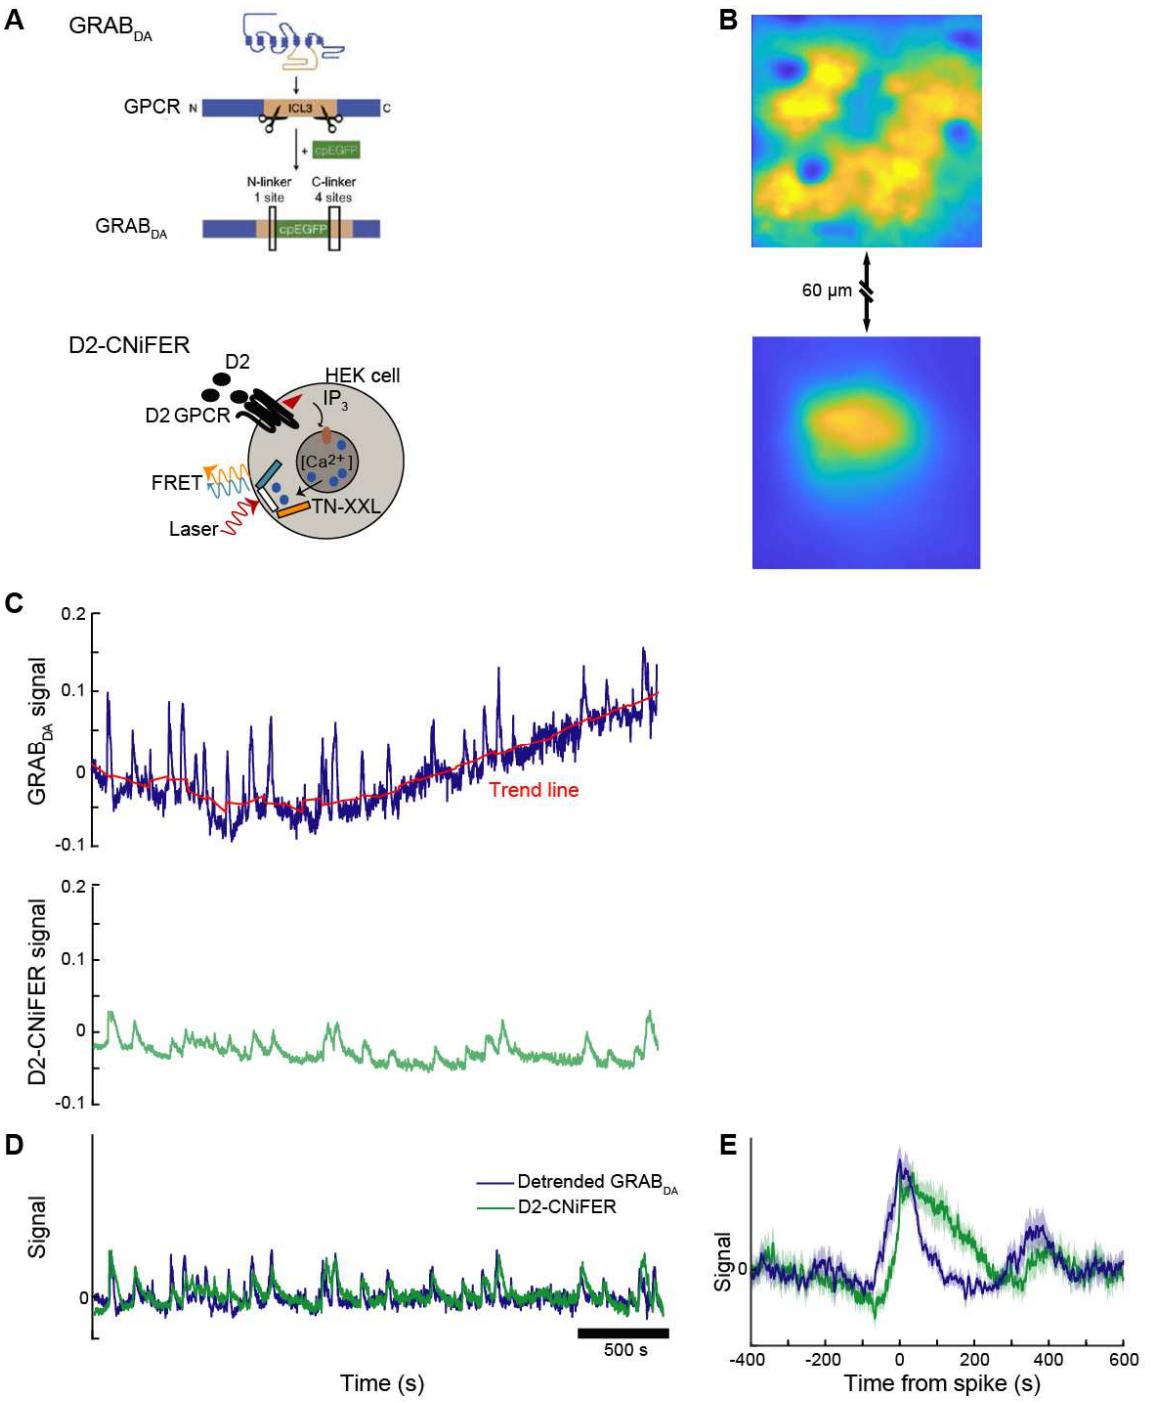

Supplemental Figure S4. Foo, Lozada, LI, Wang, Slesinger & Kleinfeld

Figure 8. Compendium of population averaged changes in $[NA]_{ex}$ over the course of our volitional reinforcement paradigm.

A. Tonic changes in $[NA]_{ex}$ show a significant increase relative to naive animals after feedback training (Day 3; $p = 0.03$). For Day 2 and Day 4, $p = 0.39$ and 0.38, respectively.

B. The rate of transient increases in $[NA]_{ex}$ was significantly higher relative to naive animals with feedback training (Day 3; $p = 0.02$), For Day 2 and Day 4, $p = 0.40$ and 0.38, respectively.

C. Distribution of amplitudes and widths for transient increases in $[NA]_{ex}$for naïve mice (Day 1; top), trained mice when feedback is on (Day 3; middle), and when feedback reinforcement is inactivated (Day 4; Bottom). The amplitudes with feedback training were not significantly different compared to naive animals (Day 3; $p = 0.49$). Interestingly, transient amplitudes were significantly smaller when feedback was turned off (Day 4; $p = 0.005$) compared to naïve animals. Transient widths were similarly unaffected by feedback training (Day 3, $p = 0.15$), but were significantly shorter after feedback was removed (Day 4, $p = 0.0003$) compared to those for naive animals. Rewarded transients had average amplitudes of 0.026 ± 0.001 and average widths of 25.2 ± 1.6 s.

D. Standard boxplot showing the lag of $[NA]_{ex}$ transients behind licking (left) and running (right). The lag of licking behind the transients was -14.0 ± 7.0 s (Day 1), $+14.4 \pm 14.1$ s (Day 2), -9.7 ± 8.4 s (Day 3), and -12.8 ± 10.7 s (Day 4). A two-sample t-test with respect to naive animals yields $p = 0.04$ (Day 2), $p = 0.35$ (Day 3), $p = 0.46$ (Day 4). The lag of running behind the transients was -4.9 ± 5.3 s (Day 1), $+11.9 \pm 19.3$ s (Day 2), -1.5 ± 2.0 s (Day 3), and -14.7 ± 11.1 s (Day 4). Two-sample t-tests with respect to naïve animals yield $p = 0.44$ (Day 2), $p = 0.58$ (Day 3), $p = 0.47$ (Day 4).

E. Standard boxplots showing the variance explained by an optimal linear filter that predicts licking (left) or running (right) from the measured $[NA]_{ex}$. A different filter was calculated for each trial and animal. The variance explained in licking was 0.04 ± 0.02 (Day 1), 0.09 ± 0.02 (Day 2), 0.12 ± 0.04 (Day 3), and 0.10 ± 0.05 (Day 4). The variance explained in running was 0.10 ± 0.02 (Day 1), 0.08 ± 0.02 (Day 2), 0.18 ± 0.06 (Day 3), and 0.14 ± 0.06 (Day 4). The variance explained by the linear filter for licking was significantly higher or days with reward ($p = 0.04$ and 0.03 for Day 2 and Day 3 respectively) compared to the naïve animal; the variance explained by the linear filter for running did not differ significantly over feedback training ($p = 0.40$ and 0.27 for Day 2 and Day 3 respectively).

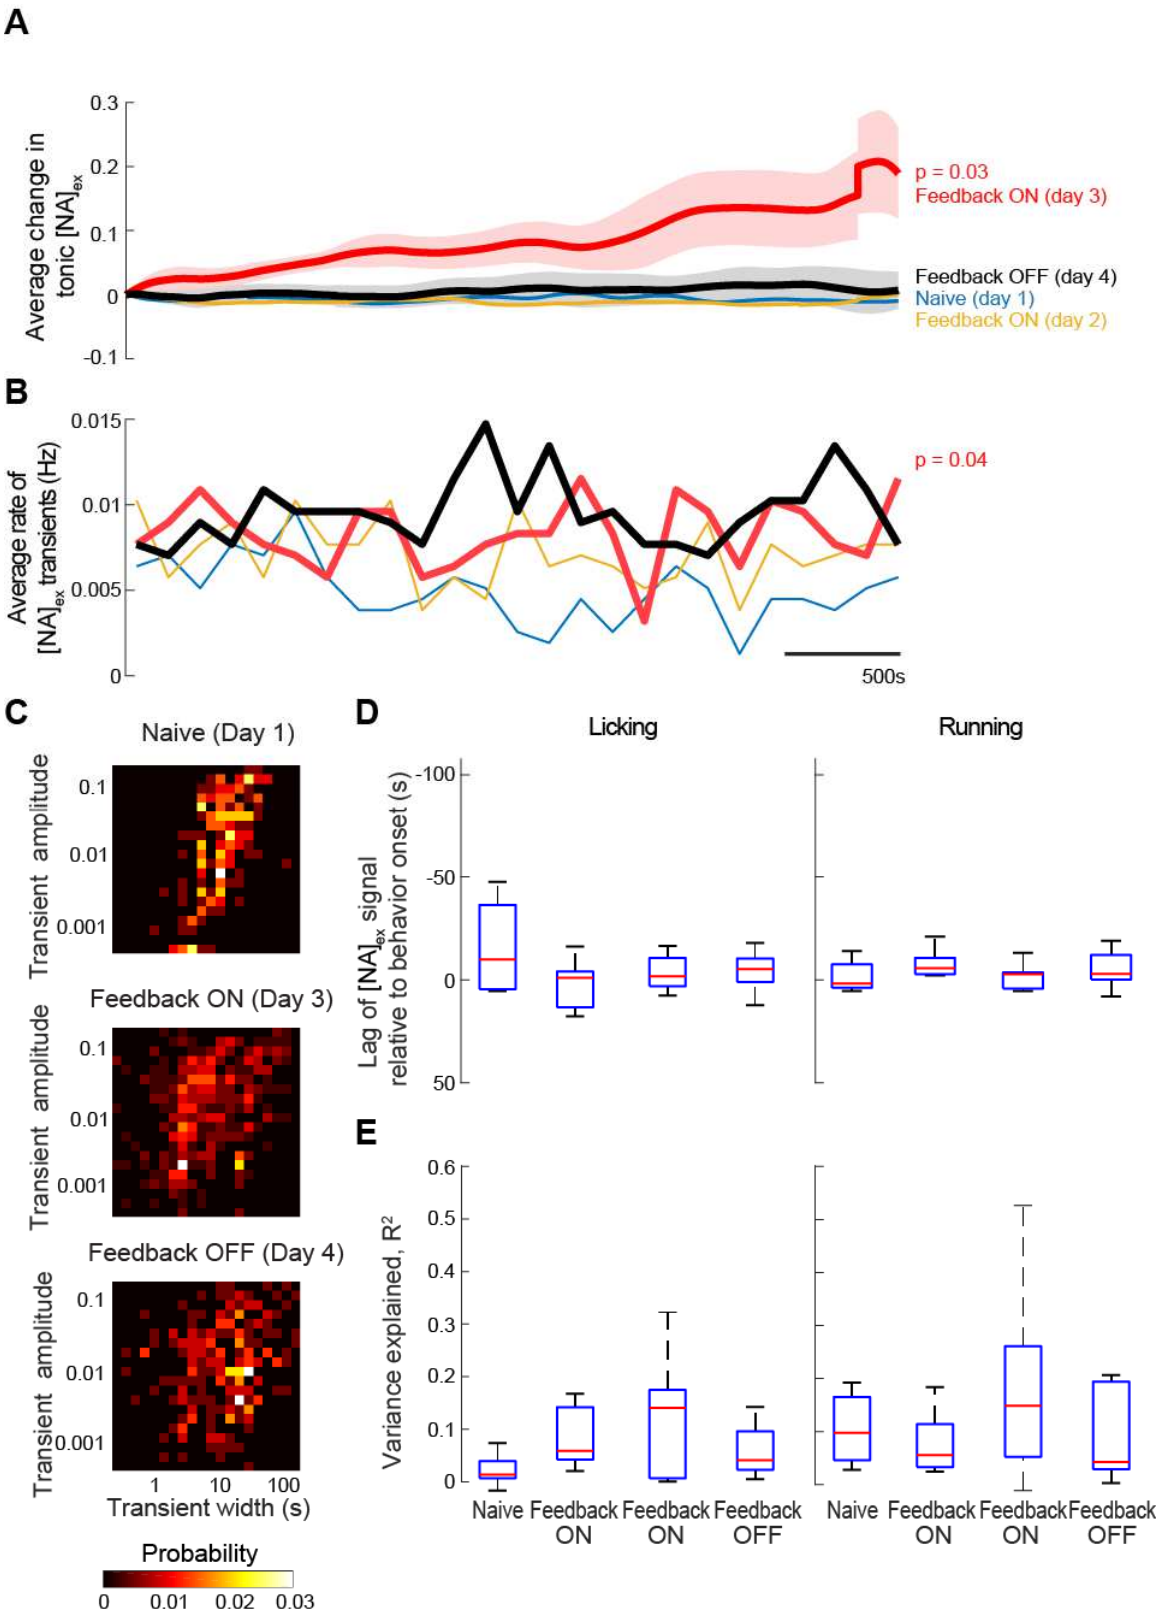
A
Average change in tonic [NA]ex
0.3
0.2
0.1
0
-0.1
p = 0.03
Feedback ON (day 3)
Feedback OFF (day 4)
Naive (day 1)
Feedback ON (day 2)
B
Average rate of [NA]ex transients (Hz)
0.015
0.01
0.005
0
p = 0.04
500s
C
Naive (Day 1)
Feedback ON (Day 3)
Feedback OFF (Day 4)
Transient amplitude
0.1
0.01
0.001
1
10
100
Transient width (s)
Probability
0
0.01
0.02
0.03
D
Licking
Running
Lag of [NA]ex signal relative to behavior onset (s)
-100
-50
0
50
E
Variance explained, R2
0.6
0.5
0.4
0.3
0.2
0.1
0
Naive
Feedback ON
Feedback ON
Feedback OFF
Naive
Feedback ON
Feedback ON
Feedback OFF

Figure 9. Compendium of population averaged changes in $[ACh]_{ex}$ over the course of our volitional reinforcement paradigm.

A. Tonic changes in $[ACh]_{ex}$ did not shift significantly after feedback training relative to naive animals (Day 3; $p = 0.48$). For Day 2 and Day 4, $p = 0.99$ and 0.52, respectively.

B. The rate of transient increases in $[ACh]_{ex}$ was not significantly higher relative to naive animals with feedback training (Day 3; $p = 0.57$), For Day 2 and Day 4, $p = 0.47$ and 0.40, respectively.

C. Distribution of amplitudes and widths for transient increases in $[ACh]_{ex}$ for naïve mice (Day 1; top), trained mice when feedback is on (Day 3; middle), and when feedback reinforcement is inactivated (Day 4; Bottom). The amplitudes with feedback training were not significantly different compared to naive animals (Day 3; $p = 0.06$). Interestingly, transient amplitudes were significantly smaller when feedback was turned off (Day 4; $p = 0.0003$) compared to naïve animals. Transient were significantly shorter after feedback training (Day 3, $p = 0.02$), and remained shorter after feedback was removed (Day 4, $p = 0.004$) compared to those for naive animals. Rewarded transients had average amplitudes of 0.037 ± 0.002 and average widths of 31.1 ± 2.2 s.

D. Standard boxplot showing the lag of $[ACh]_{ex}$ transients behind licking (left) and running (right). The lag of transients behind licking was -23.7 ± 21.8 s (Day 1), -7.5 ± 8.1 s (Day 2), 0.0 ± 3.8 s (Day 3), and -17.0 ± 10.8 s (Day 4). A two-sample t-test with respect to naive animals yields $p = 0.5$ (Day 2), $p = 0.36$ (Day 3), $p = 0.81$ (Day 4). The lag of transients behind running was -18.2 ± 11.2 s (Day 1), -12.8 ± 5.9 s (Day 2), -8.4 ± 4.1 s (Day 3), and -13.2 ± 6.5 s (Day 4). A two-sample t-test with respect to naive animals yields $p = 0.7$ (Day 2), $p = 0.47$ (Day 3), $p = 0.73$ (Day 4).

E. Standard boxplot showing the variance explained by an optimal linear filter that predicts licking (left), and running (right), from the measured $[ACh]_{ex}$. A different filter was calculated for each trial and animal. The variance explained for licking was 0.17 ± 0.09 (Day 1), 0.13 ± 0.06 (Day 2), 0.18 ± 0.06 (Day 3), and 0.12 ± 0.05 (Day 4). This difference was not significant across animals. The variance explained for running was 0.19 ± 0.09 (Day 1), 0.13 ± 0.05 (Day 2), 0.16 ± 0.04 (Day 3), and 0.14 ± 0.03 (Day 4). As with licking, differences were not significant across animals.

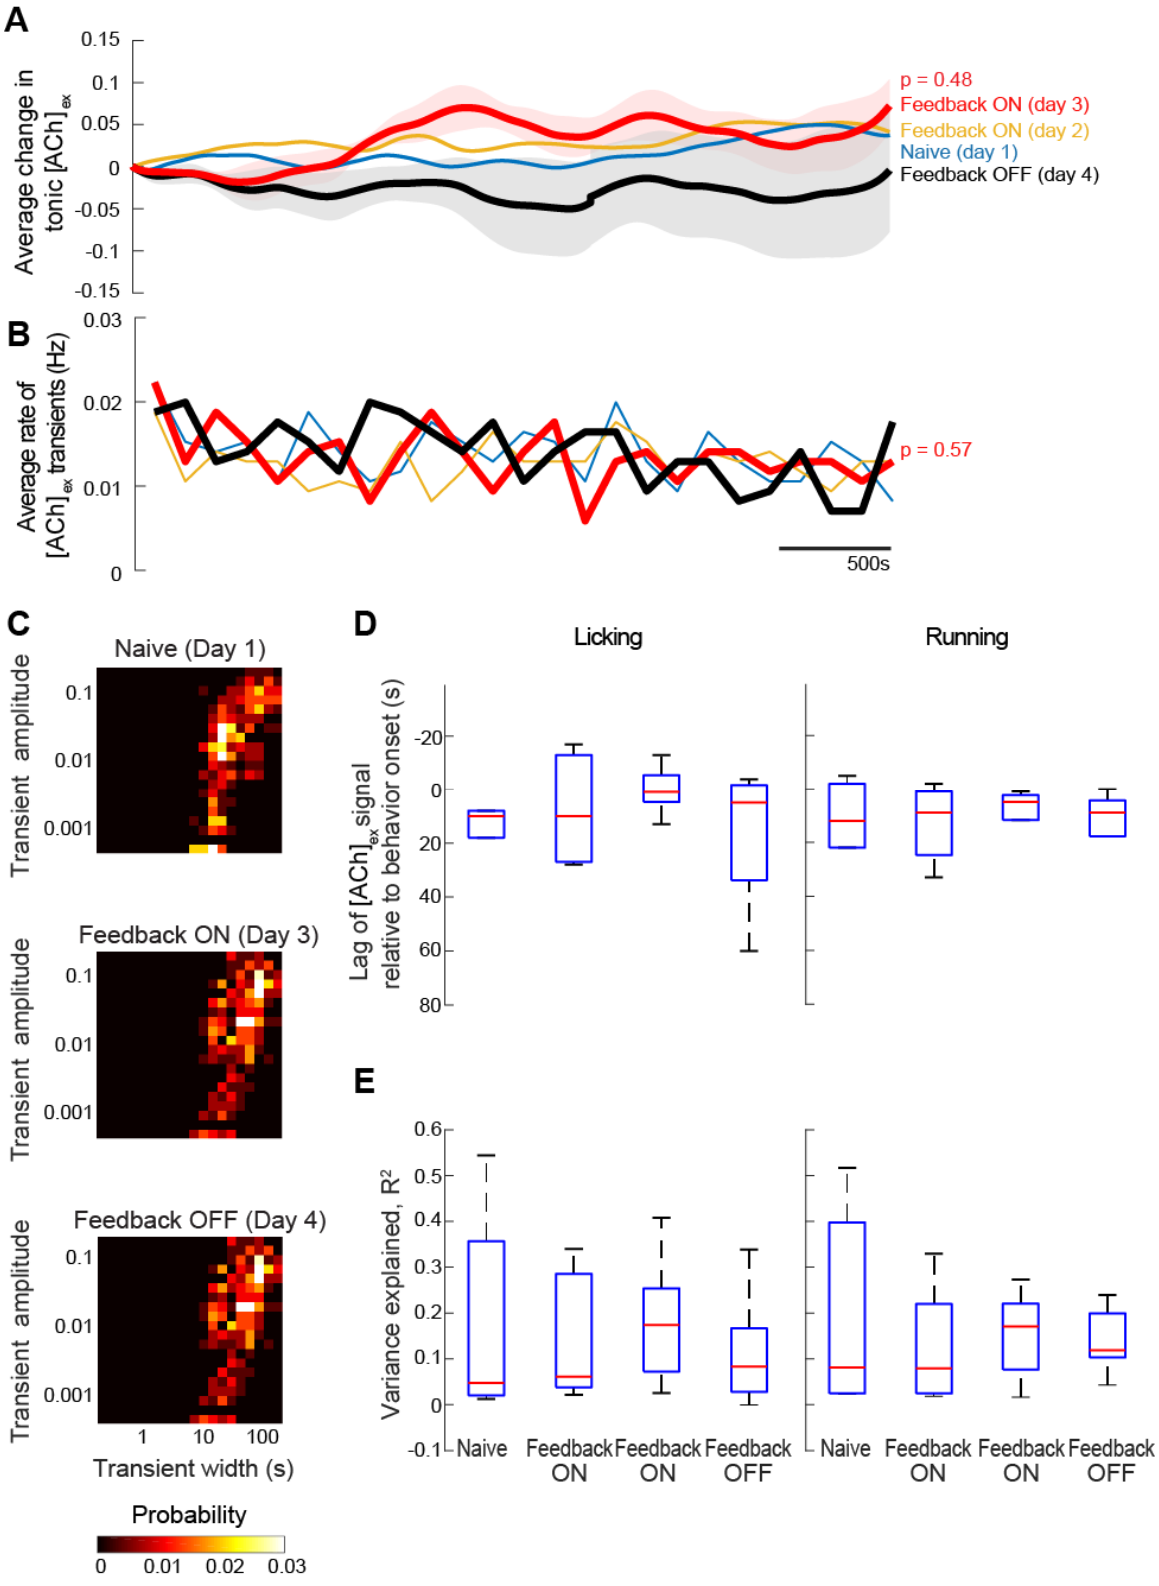
A
Average change in tonic [ACh]ex
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
p = 0.48
Feedback ON (day 3)
Feedback ON (day 2)
Naive (day 1)
Feedback OFF (day 4)
B
Average rate of [ACh]ex transients (Hz)
0.03
0.02
0.01
0
p = 0.57
500s
C
Naive (Day 1)
Feedback ON (Day 3)
Feedback OFF (Day 4)
Transient amplitude
0.1
0.01
0.001
1
10
100
Transient width (s)
Probability
0
0.01
0.02
0.03
D
Licking
Running
Lag of [ACh]ex signal relative to behavior onset (s)
-20
0
20
40
60
80
E
Variance explained, R2
0.6
0.5
0.4
0.3
0.2
0.1
0
-0.1
Naive
Feedback ON
Feedback ON
Feedback OFF
Naive
Feedback ON
Feedback ON
Feedback OFF

Figure A1. Pupil dynamics during the course of our volitional reinforcement paradigm.

A. [NA] over 2 training epochs, one early (left), and one late (right). Threshold for reward shown in green, and liquid reward epochs in grey.

B. Pupil diameter over the same epochs as **A**. Tonic [NA], calculated as in the main text, overlayed in blue and purple respectively.

C. Average running speed during each training session vs average pupil diameter during the same training session. Each gray point is a training session, and the red line is a least squares linear regression. The linear regression fits the data well, with an R^2 value of 0.77. Both pupil diameter and running speed increase with more training.

D. Transient triggered response of pupil for sessions shown in A. Day 2 in blue, Day 6 in purple, and Day 1 (naïve animal) in black. Transient onset in red. Pupil dilation shifted from preceding noradrenaline transients to lagging noradrenaline transients over training.

E. Tonic [NA] over feedback training. Feedback OFF sessions are colored in black (Days 1, 5, and 7).

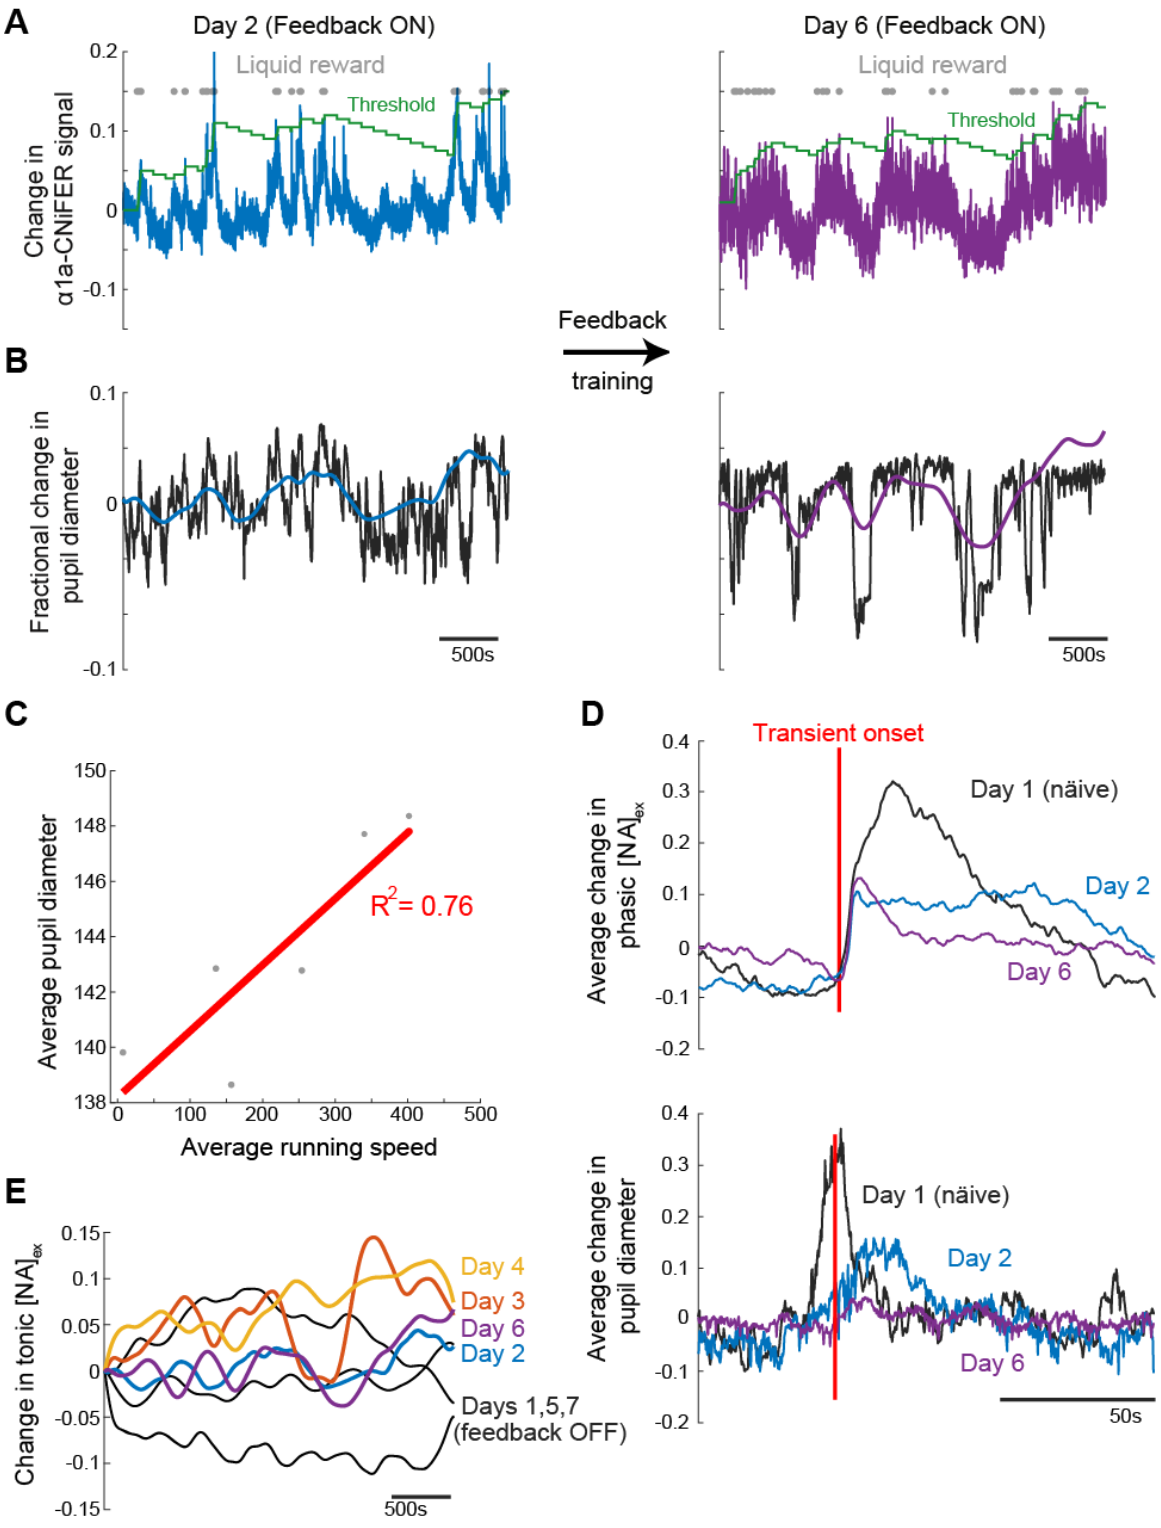
A
Day 2 (Feedback ON)
Day 6 (Feedback ON)
Liquid reward
Threshold
Change in α1a-CNiFER signal
Feedback
training
B
Fractional change in pupil diameter
500s
C
Average pupil diameter
R2= 0.76
Average running speed
D
Transient onset
Day 1 (näive)
Day 2
Day 6
Average change in phasic [NA]ex
Average change in pupil diameter
50s
E
Change in tonic [NA]ex
Day 4
Day 3
Day 6
Day 2
Days 1,5,7
(feedback OFF)
500s

Figure A2. Vasoactive response to optogenetic stimulation of noradrenergic neurons in locus coeruleus.

A. Sample a1a-CNiFER and vessel response to stimulation. The a1a-CNiFER response is shown on the top. The diameters of the two highlighted vessels are shown on the bottom. Stimulation onset is marked with black dots.

B. Stimulus triggered response of a1a-CNiFER (top), and vessel diameter (bottom) to data taken from animal in **A**. Averaging was done over 15 stimulation epochs. Stimulus onset is marked by the red line.

C. Field of view of data taken in **A** and **B**. Data was taken with an arbitrary linescan, with a line scanned over each vessel and a box scanned over the CNiFER implant.

D. Data from an animal with a unique response to optogenetic stimulation. On top, the a1a-CNiFER response; on the bottom, the vessel response.

E. Stimulus triggered response of a1a-CNiFER (top), and vessel diameter (bottom) to data taken from animal in **D**. Averaging was done over 12 stimulation epochs. Stimulus onset is marked by the red line. Noradrenaline transients were consistently evoked across all animals, however vessel response was mixed.

F. Compendium responses to optogenetic stimulation. Top, the a1a-CNiFER response for each animal (N = 5). Middle, vessels that dilated in response to stimulation (N = 4). Bottom, vessels that constricted in response to stimulation (N = 4). In comparison, 14 vessels did not respond to optogenetic stimulation. Red line marks stimulus onset, and grey line marks average response over all vessels.

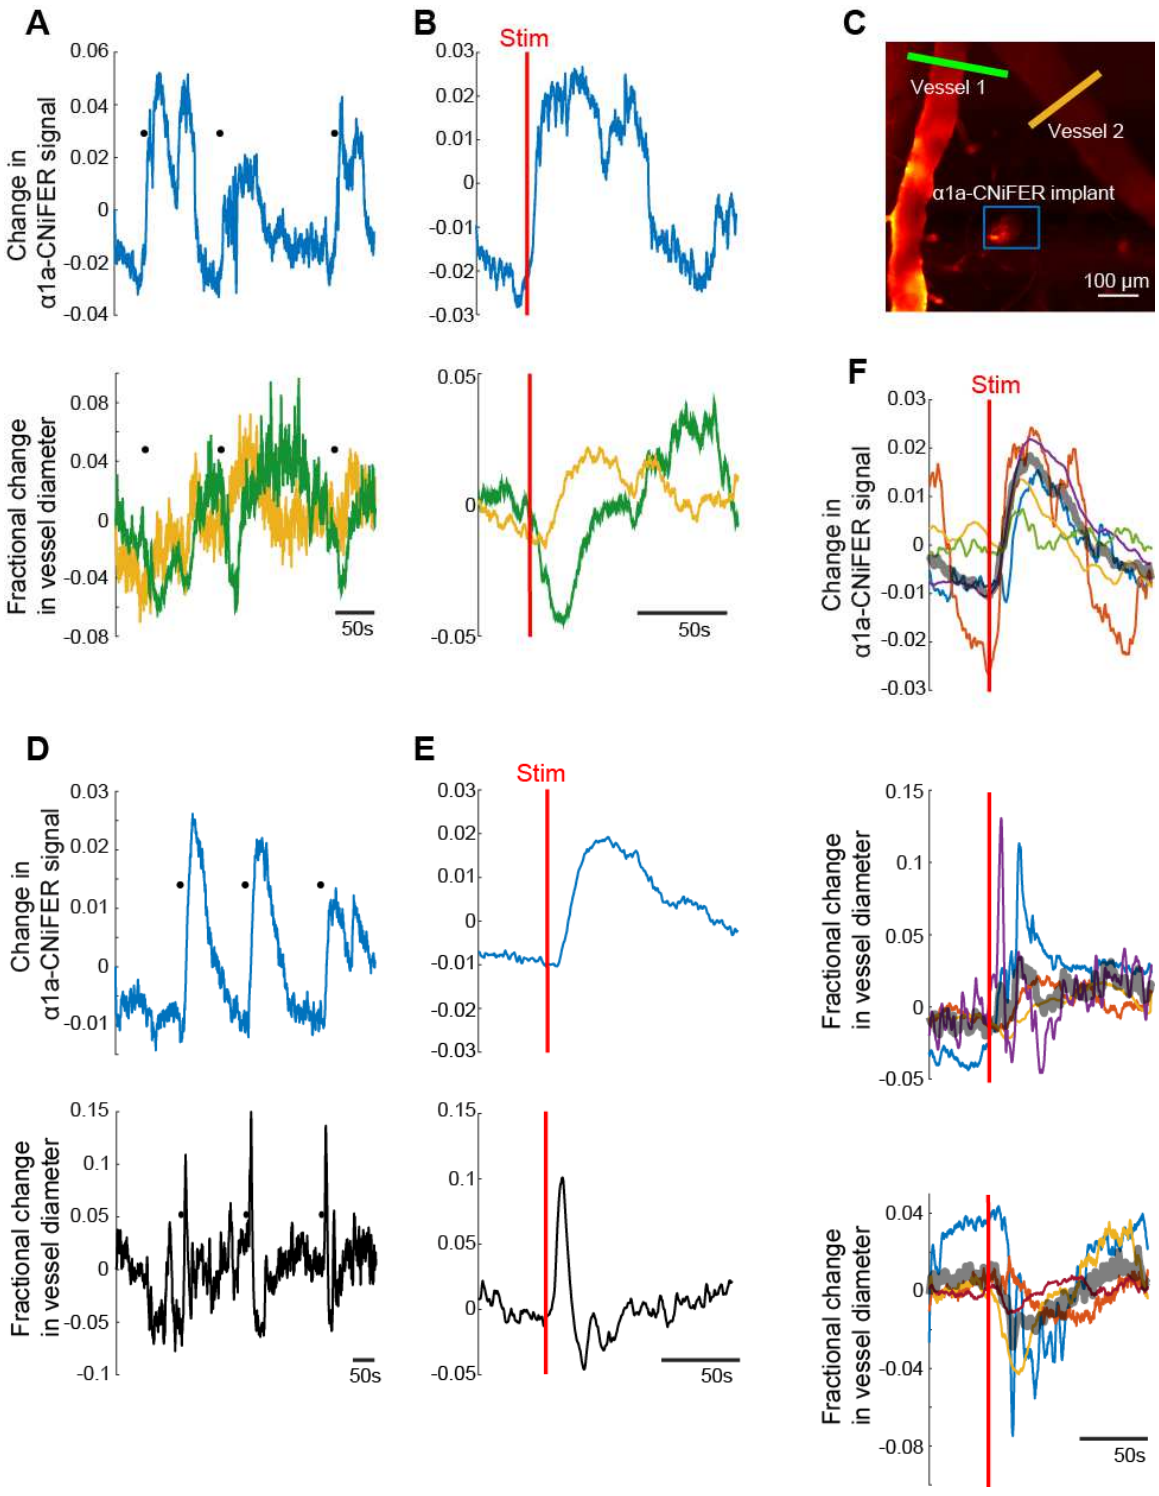
A
B
C
Stim
Vessel 1
Vessel 2
α1a-CNiFER implant
100 µm
Change in α1a-CNiFER signal
Fractional change in vessel diameter
50s
F
D
E

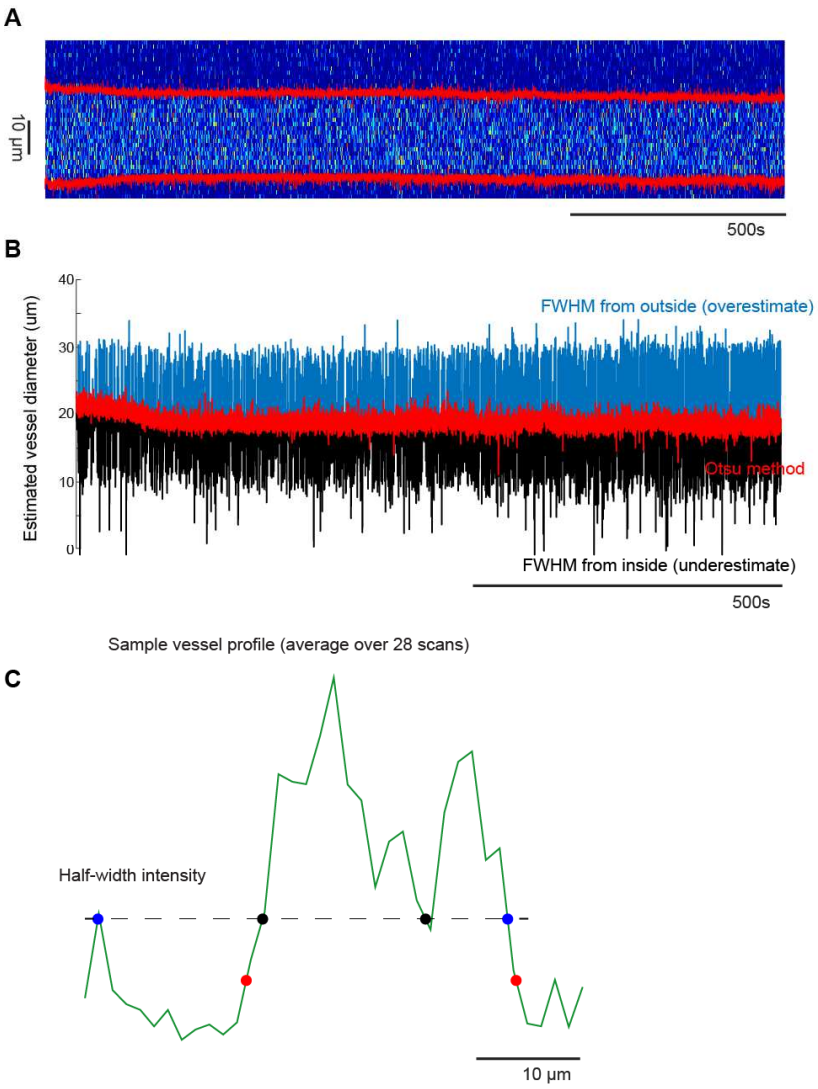

Figure A3. Otsu's method to calculate vessel diameter from arbitrary linescan data.

A. Sample vessel scan data taken from a cortical window using a 2-photon microscope. Vessels were filled with TRITC dye. Vessel edge points as calculated with Otsu's method shown in red.

B. Comparison of FWHM method and Otsu's method. Calculated diameter using the FWHM method measuring from the outside inwards is shown in blue and inside outwards is shown in black. Diameter calculated with Otsu's method is shown in red.

C. Sample vessel profile from which diameter measurements are calculated (green). Half-maximum intensity is denoted by the dotted line. Calculated vessel endpoints for the FWHM method measuring from the outside and inside are shown in blue and black respectively. Endpoints calculated from Otsu's method denoted in red.

References

1. Bekar, L.K., Wei, H.S., and Nedergaard, M. (2012). The locus coeruleus-norepinephrine network optimizes coupling of cerebral blood volume with oxygen demand. J. Cereb. Blood Flow Metab. *32*, 2135–2145.

2. Bekar, L.K., He, W., and Nedergaard, M. (2008). Locus coeruleus alpha-adrenergic-mediated activation of cortical astrocytes in vivo. Cereb. Cortex *18*, 2789–2795.

3. Cohen, Z., Molinatti, G., and Hamel, E. (1997). Astroglial and vascular interactions of noradrenaline terminals in the rat cerebral cortex. J. Cereb. Blood Flow Metab. *17*, 894–904.

.

4. Toussay, X., Basu, K., Lacoste, B., and Hamel, E. (2013). Locus coeruleus stimulation recruits a broad cortical neuronal network and increases cortical perfusion. J. Neurosci. *33*, 3390–3401.

5. Piascik, M.T., and Perez, D.M. (2001). Alpha1-adrenergic receptors: new insights and directions. J. Pharmacol. Exp. Ther. *298*, 403–410.

6. Chistiakov, D.A., Ashwell, K.W., Orekhov, A.N., and Bobryshev, Y.V. (2015). Innervation of the arterial wall and its modification in atherosclerosis. Auton. Neurosci. *193*, 7–11.

7. Hauss-Wegrzyniak, B., Lynch, M.A., Vraniak, P.D., and Wenk, G.L. (2002). Chronic brain inflammation results in cell loss in the entorhinal cortex and impaired LTP in perforant path-granule cell synapses. Exp. Neurol. *176*, 336–341.

8. Xu, H.-T., Pan, F., Yang, G., and Gan, W.-B. (2007). Choice of cranial window type for in vivo imaging affects dendritic spine turnover in the cortex. Nat. Neurosci. *10*, 549–551.

9. Forbes, H.S. (1941). The pial circulation of normal, non-anesthetized animals, part 1. Description of a method of observation. Journal of Pharmacology and Experimental Therapeutics *71*, 325–330.

10. Berridge, C.W., and Waterhouse, B.D. (2003). The locus coeruleus-noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. Brain Res. Brain Res. Rev. *42*, 33–84.

11. Staddon, J.E.R., and Cerutti, D.T. (2003). Operant conditioning. Annu. Rev. Psychol. *54*, 115–144.

12. Reimer, J., Froudarakis, E., Cadwell, C.R., Yatsenko, D., Denfield, G.H., and Tolias, A.S. (2014). Pupil fluctuations track fast switching of cortical states during quiet wakefulness. Neuron *84*, 355–362.

13. Shindou, T., Shindou, M., Watanabe, S., and Wickens, J. (2019). A silent eligibility trace enables dopamine-dependent synaptic plasticity for reinforcement learning in the mouse striatum. Eur. J. Neurosci. *49*, 726–736.

14. Lavín, C., San Martín, R., and Rosales Jubal, E. (2013). Pupil dilation signals uncertainty and surprise in a learning gambling task. Front. Behav. Neurosci. *7*, 218.

15. Preuschoff, K., 't Hart, B.M., and Einhäuser, W. (2011). Pupil dilation signals surprise: evidence for noradrenaline's role in decision making. Front. Neurosci. *5*, 115.

16. McGinley, M.J., David, S.V., and McCormick, D.A. (2015). Cortical membrane potential signature of optimal states for sensory signal detection. Neuron *87*, 179–192.

17. Ariano, M.A., and Sibley, D.R. (1994). Dopamine receptor distribution in the rat CNS: elucidation using anti-peptide antisera directed against D1A and D3 subtypes. Brain Res. *649*, 95–110.

18. Descarries, L., Lemay, B., Doucet, G., and Berger, B. (1987). Regional and laminar density of the dopamine innervation in adult rat cerebral cortex. Neuroscience *21*, 807–824.

19. Bassant, M.H., Ennouri, K., and Lamour, Y. (1990). Effects of iontophoretically applied monoamines on somatosensory cortical neurons of unanesthetized rats. Neuroscience *39*, 431–439.

20. Jacob, S.N., and Nienborg, H. (2018). Monoaminergic neuromodulation of sensory processing. Front. Neural Circuits *12*, 51.

21. Lee, C.R., and Margolis, D.J. (2016). Pupil Dynamics Reflect Behavioral Choice and Learning in a Go/NoGo Tactile Decision-Making Task in Mice. Front. Behav. Neurosci. *10*, 200.

22. Reimer, J., McGinley, M.J., Liu, Y., Rodenkirch, C., Wang, Q., McCormick, D.A., and Tolias, A.S. (2016). Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. Nat. Commun. *7*, 13289.

23. Coddington, L.T., and Dudman, J.T. (2018). The timing of action determines reward prediction signals in identified midbrain dopamine neurons. Nat. Neurosci. *21*, 1563–1573.

24. Floresco, S.B., West, A.R., Ash, B., Moore, H., and Grace, A.A. (2003). Afferent modulation of dopamine neuron firing differentially regulates tonic and phasic dopamine transmission. Nat. Neurosci. *6*, 968–973.

25. Dodson, P.D., Dreyer, J.K., Jennings, K.A., Syed, E.C.J., Wade-Martins, R., Cragg, S.J., Bolam, J.P., and Magill, P.J. (2016). Representation of spontaneous movement by dopaminergic neurons is cell-type selective and disrupted in parkinsonism. Proc Natl Acad Sci USA *113*, E2180-8.

26. Aberman, J.E., and Salamone, J.D. (1999). Nucleus accumbens dopamine depletions make rats more sensitive to high ratio requirements but do not impair primary food reinforcement. Neuroscience *92*, 545–552.

27. da Silva, J.A., Tecuapetla, F., Paixão, V., and Costa, R.M. (2018). Dopamine neuron activity before action initiation gates and invigorates future movements. Nature *554*, 244–248.

28. Berke, J.D. (2018). What does dopamine mean? Nat. Neurosci. *21*, 787–793.

29. Niv, Y., Daw, N.D., Joel, D., and Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. Psychopharmacology (Berl) *191*, 507–520.

30. Beeler, J.A., Daw, N., Frazier, C.R.M., and Zhuang, X. (2010). Tonic dopamine modulates exploitation of reward learning. Front. Behav. Neurosci. *4*, 170.

31. Cinotti, F., Fresno, V., Aklil, N., Coutureau, E., Girard, B., Marchand, A.R., and Khamassi, M. (2019). Dopamine blockade impairs the exploration-exploitation trade-off in rats. Sci. Rep. *9*, 6770.

32. Zhuang, X., Oosting, R.S., Jones, S.R., Gainetdinov, R.R., Miller, G.W., Caron, M.G., and Hen, R. (2001). Hyperactivity and impaired response habituation in hyperdopaminergic mice. Proc Natl Acad Sci USA *98*, 1982–1987.

33. Costa, V.D., Tran, V.L., Turchi, J., and Averbeck, B.B. (2014). Dopamine modulates novelty seeking behavior during decision making. Behav. Neurosci. *128*, 556–566.

34. Dulawa, S.C., Grandy, D.K., Low, M.J., Paulus, M.P., and Geyer, M.A. (1999). Dopamine D4 receptor-knock-out mice exhibit reduced exploration of novel stimuli. J. Neurosci. *19*, 9550–9556.

35. Shyr, M.H., Tsai, T.H., Yang, C.H., Chen, H.M., Ng, H.F., and Tan, P.P. (1997). Propofol anesthesia increases dopamine and serotonin activities at the

somatosensory cortex in rats: a microdialysis study. Anesth. Analg. *84*, 1344–1348.

36. Aransay, A., Rodríguez-López, C., García-Amado, M., Clascá, F., and Prensa, L. (2015). Long-range projection neurons of the mouse ventral tegmental area: a single-cell axon tracing analysis. Front. Neuroanat. *9*, 59.

37. Gurevich, E.V., Robertson, R.T., and Joyce, J.N. (2001). Thalamo-cortical afferents control transient expression of the dopamine D(3) receptor in the rat somatosensory cortex. Cereb. Cortex *11*, 691–701.

38. Yu, Q., Liu, Y.-Z., Zhu, Y.-B., Wang, Y.-Y., Li, Q., and Yin, D.-M. (2019). Genetic labeling reveals temporal and spatial expression pattern of D2 dopamine receptor in rat forebrain. Brain Struct. Funct. *224*, 1035–1049.

39. Bouthenet, M.L., Martres, M.P., Sales, N., and Schwartz, J.C. (1987). A detailed mapping of dopamine D-2 receptors in rat central nervous system by autoradiography with [125I]iodosulpride. Neuroscience *20*, 117–155.

40. Björklund, A., and Dunnett, S.B. (2007). Dopamine neuron systems in the brain: an update. Trends Neurosci. *30*, 194–202.

41. Howe, M.W., and Dombeck, D.A. (2016). Rapid signalling in distinct dopaminergic axons during locomotion and reward. Nature *535*, 505–510.

42. Agnati, L.F., Zoli, M., Strömberg, I., and Fuxe, K. (1995). Intercellular communication in the brain: wiring versus volume transmission. Neuroscience *69*, 711–726.

43. Sulzer, D., Cragg, S.J., and Rice, M.E. (2016). Striatal dopamine neurotransmission: regulation of release and uptake. Basal Ganglia *6*, 123–148.

44. Rice, M.E., and Cragg, S.J. (2008). Dopamine spillover after quantal release: rethinking dopamine transmission in the nigrostriatal pathway. Brain Res. Rev. *58*, 303–313.

45. Picciotto, M.R., Higley, M.J., and Mineur, Y.S. (2012). Acetylcholine as a neuromodulator: cholinergic signaling shapes nervous system function and behavior. Neuron *76*, 116–129.

46. Jorm, C.M., and Stamford, J.A. (1993). Actions of the hypnotic anaesthetic, dexmedetomidine, on noradrenaline release and cell firing in rat locus coeruleus slices. Br. J. Anaesth. *71*, 447–449.

47. Aston-Jones, G., and Cohen, J.D. (2005). Adaptive gain and the role of the locus coeruleus-norepinephrine system in optimal performance. J. Comp. Neurol. *493*, 99–110.

48. Coddington, L.T., and Dudman, J.T. (2019). Learning from Action: Reconsidering Movement Signaling in Midbrain Dopamine Neuron Activity. Neuron *104*, 63–77.

49. Hyland, B.I., Reynolds, J.N.J., Hay, J., Perk, C.G., and Miller, R. (2002). Firing modes of midbrain dopamine cells in the freely moving rat. Neuroscience *114*, 475–492.

50. Mitra, P., and Bokil, H. (2007). Observed Brain Dynamics | Semantic Scholar. undefined.

51. Percival, D.B., and Walden, A.T. (1993). Spectral Analysis for Physical Applications (Cambridge: Cambridge University Press).

52. Aljadeff, J., Lansdell, B.J., Fairhall, A.L., and Kleinfeld, D. (2016). Analysis of neuronal spike trains, deconstructed. Neuron *91*, 221–259.

53. Tsai, P.S., and Kleinfeld, D. (2009). In Vivo Two-Photon Laser Scanning Microscopy with Concurrent Plasma-Mediated Ablation Principles and Hardware Realization. In In Vivo Optical Imaging of Brain Function Frontiers in Neuroscience., R. D. Frostig, R. D. Frostig, R. D. Frostig, and R. D. Frostig, eds. (Boca Raton (FL): CRC Press/Taylor & Francis).

54. Svoboda, K., Denk, W., Kleinfeld, D., and Tank, D.W. (1997). In vivo dendritic calcium dynamics in neocortical pyramidal neurons. Nature *385*, 161–165.

55. Shih, A.Y., Mateo, C., Drew, P.J., Tsai, P.S., and Kleinfeld, D. (2012). A polished and reinforced thinned-skull window for long-term imaging of the mouse brain. J. Vis. Exp.

56. Sun, F., Zeng, J., Jing, M., Zhou, J., Feng, J., Owen, S.F., Luo, Y., Li, F., Wang, H., Yamaguchi, T., *et al.* (2018). A genetically encoded fluorescent sensor enables rapid and specific detection of dopamine in flies, fish, and mice. Cell *174*, 481-496.e19.

57. Muller, A., Joseph, V., Slesinger, P.A., and Kleinfeld, D. (2014). Cell-based reporters reveal in vivo dynamics of dopamine and norepinephrine release in murine cortex. Nat. Methods *11*, 1245–1252.

58. Swamy, B.E.K., and Venton, B.J. (2007). Carbon nanotube-modified microelectrodes for simultaneous detection of dopamine and serotonin in vivo. Analyst *132*, 876–884.

59. Chefer, V.I., Thompson, A.C., Zapata, A., and Shippenberg, T.S. (2009). Overview of brain microdialysis. Curr. Protoc. Neurosci. *Chapter 7*, Unit7.1.

60. Neely, R.M., Koralek, A.C., Athalye, V.R., Costa, R.M., and Carmena, J.M. (2018). Volitional modulation of primary visual cortex activity requires the basal ganglia. Neuron *97*, 1356-1368.e4.

61. Fetz, E.E. (1969). Operant conditioning of cortical unit activity. Science *163*, 955–958.

62. Carmena, J.M., Lebedev, M.A., Crist, R.E., O'Doherty, J.E., Santucci, D.M., Dimitrov, D.F., Patil, P.G., Henriquez, C.S., and Nicolelis, M.A.L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. PLoS Biol. *1*, E42.

63. Mohebi, A., Pettibone, J.R., Hamid, A.A., Wong, J.-M.T., Vinson, L.T., Patriarchi, T., Tian, L., Kennedy, R.T., and Berke, J.D. (2019). Dissociable dopamine dynamics for learning and motivation. Nature *570*, 65–70.

64. Threlfell, S., Lalic, T., Platt, N.J., Jennings, K.A., Deisseroth, K., and Cragg, S.J. (2012). Striatal dopamine release is triggered by synchronized activity in cholinergic interneurons. Neuron *75*, 58–64.

65. Koob, G.F., Riley, S.J., Smith, S.C., and Robbins, T.W. (1978). Effects of 6-hydroxydopamine lesions of the nucleus accumbens septi and olfactory tubercle on feeding, locomotor activity, and amphetamine anorexia in the rat. J. Comp. Physiol. Psychol. *92*, 917–927.

66. Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. Science *275*, 1593–1599.

67. Langston, J.W., Ballard, P., Tetrud, J.W., and Irwin, I. (1983). Chronic Parkinsonism in humans due to a product of meperidine-analog synthesis. Science *219*, 979–980.

68. Albin, R.L., Young, A.B., and Penney, J.B. (1989). The functional anatomy of basal ganglia disorders. Trends Neurosci. *12*, 366–375.

69. Carlsson, A. (1959). The occurrence, distribution and physiological role of catecholamines in the nervous system. Pharmacol. Rev. *11*, 490–493.

70. Patriarchi, T., Cho, J.R., Merten, K., Howe, M.W., Marley, A., Xiong, W.-H., Folk, R.W., Broussard, G.J., Liang, R., Jang, M.J., *et al.* (2018). Ultrafast neuronal imaging of dopamine dynamics with designed genetically encoded sensors. Science *360*.