UNIVERSITY OF CALIFORNIA SAN DIEGO

Learning and plasticity for a brain-computer interface task with two-photon calcium imaging

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

in

Neuroscience

with a Specialization in Computational Neurosciences

by

Akinori Mitani

Committee in charge:

Professor Takaki Komiyama, Chair
Professor Matthew Banghart
Professor Timothy Gentner
Professor Christina Gremel
Professor Nicholas Spitzer

2018

The Dissertation of Akinori Mitani is approved, and is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____

Chair

University of California San Diego

2018

DEDICATION

To my parents and brother.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

The first and foremost, I would like to express my deepest gratitude to Dr. Takaki Komiyama for welcoming me as a graduate student in his lab, being a scientific advisor, providing financial support, and inspiring me through 6 years. I truly appreciate that he let me explore myself while providing me the guidance exactly when I needed it.

I would also like to thank my thesis committee members, Dr. Nicholas Spitzer, Dr. Timothy Gentner, Dr. Christina Gremel, and Dr. Matthew Banghart for their advice and support. The discussion with them motivated me and helped me understand the project more deeply.

I was extremely lucky to have Mingyuan Dong, a talented visiting student, to help me with the experiments for a year. He was dedicated to science and conducted all the experiments very carefully while he was here. He became a co-author of the paper, and kindly allowed me to include the paper in this dissertation.

I would also like to thank other lab members; Anna Kim, Keelin O'Neil, Lucia Hall, Tara Loveland and Oscar Arroyo for technical assistance. Andy Peters, Monica Chu, Jeff Dahlen, Wankun Li, Chi Ren, and Bethanny Danskin for being graduate students at the same time to share the hardships. Hiroshi Makino, Jun Lee, Simon Chen, Eunjung Hwang, Haixin Liu, Nathen Hedrick, An Wu, Ryoma Hattori, Zhongming Lu and Assaf Ramot for being post-docs and role models for me. They taught me how to do experiments and data analysis, and they helped my projects through stimulating conversation and providing me of both mental and scientific support.

My appreciation extends to the program administrator and coordinator, Erin Gilbert and Linh Vandermar, for all their hard work to organize events, run the program smoothly, and help every student. When I had questions about course requirements and paperwork, they were always there to help me and made sure I was doing it right.

The Neuroscience Graduate Program has given me a sense of belonging. All the official and unofficial events, workshops and seminars brought us together and created a friendly and supportive atmosphere. I am very thankful and proud to be part of the program.

My tuition for the first two years and my stipend for the first five years were funded by the Nakajima foundation. I appreciate their generosity and dedication to science and education. I could fully focus on my own research because of their support.

Finally but most importantly, I would like to thank all the animals handled in the research. This research can not be done without their sacrifice and contributions.

Chapter 2 is material previously published. Mitani A, Dong M, and Komiyama T. Brain-Computer Interface with Inhibitory Neurons Reveals Subtype-Specific Strategies. Current Biology, 2018. The dissertation author was the primary investigator and author of this material.

Chapter 4 is material currently being prepared for submission for publication. Mitani A, and Komiyama T. Real-time image processing for two-photon calcium imaging. The dissertation author was the primary investigator and author of this material.

VITA

2012 Bachelor of Medicine
The University of Tokyo, Tokyo, Japan

2018 Doctor of Philosophy in Neuroscience
with a Specialization in Computational Neuroscience
University of California San Diego, La Jolla, CA

PUBLICATIONS

Mitani, A., Sasaki, R., Oizumi, M., and Uka, T. (2013). A Leaky-Integrator Model as a Control Mechanism Underlying Flexible Decision Making during Task Switching. PLoS One 8, e59670.

Jacob, V.,* Mitani, A.,* Toyoizumi, T., and Fox, K. (2017). Whisker row deprivation affects the flow of sensory information through rat barrel cortex. J. Neurophysiol. 117, 4–17.
* The authors contributed equally

Mitani, A., Dong, M., and Komiyama, T. (2018). Brain-Computer Interface with Inhibitory Neurons Reveals Subtype-Specific Strategies. Curr. Biol. 28, 77–83.e4.

ABSTRACT OF THE DISSERTATION

Learning and plasticity for a brain-computer interface task with two-photon calcium imaging

by

Akinori Mitani

Doctor of Philosophy in Neuroscience with a Specialization in Computational Neurosciences

University of California San Diego, 2018

Professor Takaki Komiyama, Chair

Direct communication with the brain by brain-computer interfaces (BCI) has been one of the goals of neuroscience due to their potential therapeutic applications. It has been shown that animals, including humans, can learn to control the movement of prosthetic devices with BCI. The control, however, is far from smooth and precise control of the limbs of healthy individuals, and there is a need to further improve the performance of BCI.

One of the limitations of the previous studies is that they have been mainly targeting somatic activity of excitatory neurons, while different cell types possesses different functions in cortical computations and likely different capacities to control BCI. Here, we made a first step in addressing this issue by tracking the plastic changes of three major types of cortical inhibitory neurons (INs) during a BCI task using two-photon calcium imaging. Mice were rewarded when the activity of the positive target neuron (N+) exceeded that of the negative

target neuron (N-) beyond a set threshold. Mice improved the performance with all subtypes using subtype-specific strategies. When parvalbumin (PV)-expressing INs were targeted, the activity of N- decreased. However, targeting of somatostatin (SOM)- and vasoactive intestinal peptide (VIP)-expressing INs led to an increase of the N+ activity. These results demonstrate that INs can be individually modulated in a subtype-specific manner, and highlight the versatility of neural circuits in adapting to new demands by using cell-type specific strategies.

Another potential limitation is that the mechanisms underlying BCI learning, and how the plasticity of the neurons contributes to learning has not been elucidated yet. In a motor learning task, it has been shown that dendritic spine dynamics and clustering are associated with learning, and nonlinear summation of excitatory inputs among nearby dendritic spines determine how the neuron responds. Does the spatial distribution of spines affect BCI learning? Does BCI training induce further spine formation and elimination? Answering those questions will help understanding the mechanisms of BCI learning. As a first step, here, I applied a modified neural feedback system to spine imaging and examined whether mice can learn to modulate spine-specific activity.

These are the first BCI tasks with lateral motion artifact correction. The implemented algorithms were presented with performance analysis.

Chapter 1. Introduction

In our daily life, we interact with the environment by moving around using muscles and receiving inputs through various sensory systems. Even though we can think inside the brain, brain itself cannot directly interact with the outside world, and our outputs hugely rely on volitional muscle contraction. This becomes more apparent when this information channel is limited. For example, patients with late stage amyotrophic lateral sclerosis have difficulty in communicating with others because of the paralysis, even when their brains are still functional, and they can contemplate. A potential way to overcome such limitation is to communicate directly with the brain by recording activity of and/or stimulating neurons. The technology is called either brain-machine interface, brain-computer interface (BCI), or neural prosthetics, almost interchangeably, depending on the field of study. Not only for therapeutic applications, people have also dreamed of such technology extending abilities of healthy individuals and enabling a new way of communication.

In research, BCI studies have seen an increase in popularity due to their potential for direct neuroprosthetic applications for amputees and disabled individuals. Supporting this promise, animals—including humans—can learn even arbitrary mapping between the activity of cortical neurons and movement of prosthetic devices. In 1969, Fetz and colleagues showed that monkey can increase firing rate of a single neuron to receive food reward (Fetz, 1969). It was shown for the first time that an animal can modulate the activity of a neuron which is directly associated with a task. With the advancement of the technology, it has been shown that the animals can learn to perform a more complex task by modifying neural activity. Valliste and colleagues succeeded in training monkeys to control a robot arm using the recorded neural activity for self-feeding (Velliste et al., 2008). It has been also shown that it

can be applied to human patients with tetraplegia (Collinger et al., 2013; Hochberg et al., 2012), and interestingly, instead of controlling a robot arm, the interface can be used to bypass neural damage and directly activate muscles using neural activity (Moritz et al., 2008).

Yet, the performance of neuroprosthetic device control has been nowhere near that of limb control in healthy individuals, presenting a dire need to improve the performance. One potential method to improve task performance is to increase the number of recorded neurons. Starting from tungsten electrodes to record mainly from one neuron (Hubel, 1957), the development of tetrodes (Gray et al., 1995), multielectrode arrays with silicon technology (e.g. Utah array, (Maynard et al., 1997)), and active CMOS probes (e.g. Neuropixels, (Jun et al., 2017)), has increased the number of simultaneously recorded neurons dramatically (reviewed in (Steinmetz et al., 2018)). Combining the activity information of more neurons, the intention of the subject could potentially be inferred more accurately, achieving smoother and more complex control of the output.

Alternatively, taking biological mechanisms into account is a potential way to improve BCI performance. Most of the existing studies have focused on excitatory neurons, partly because they are abundant and easy to record from. Along with excitatory neurons, there are inhibitory neurons of various cell types in the brain. Recent development of genetically encoded calcium indicators enabled us to record activity of genetically targeted neural population. Combined with molecular techniques, the indicator can be expressed by a specific type of neurons. Inhibitory neurons may show different forms of plasticity, and designing a decoder which handles each cell type differently may improve the performance. Another potentially related biological feature of the neurons is clustering of dendritic spines. Excitatory neurons receive inputs at dendritic spines along the dendrite, and the spatial

distribution of inputs may contribute to the plasticity of the cell and it can influence BCI-task learning. Understanding the effect of functional spine clustering along the dendrite may contribute to BCI development if we can take connectivity pattern into account. Plasticity of inhibitory neurons in BCI task and the effect of spine clustering in the dendrites will be further discussed in the following sections.

Studies of BCI have also been contributing to understand how our brain works, especially in terms of learning and plasticity. A decoder associates neural activity with the output of BCI. We can use a fixed decoder and examine how the brain learn to make use of it. Such approach would tell us about brain plasticity and what neural activity pattern can be achieved. Later this can be incorporated into decoder design, to make use of brain plasticity to increase training speed and achieve higher performance (Shenoy and Carmena, 2014).

Can animals make use of any decoder? A frequently asked question is whether animals can learn completely arbitrary neural activity patterns with BCI-task training, or there is a limited pool of neural activity patterns where they can choose from. Supporting the idea that the animal can independently condition multiple neurons involved in the BCI task, it has been shown that animals perform well in a task with arbitrary neural activity mapping. Jerosiewicz showed that when the learned association was rotated in a subset of the neurons, those rotated neurons contribute less to cursor movements than the non-rotated neurons after relearning, showing that animals can learn to rely on useful neurons (Jarosiewicz et al., 2008). Ganguly and Carmena showed that a stable neural pattern emerges in long-term BCI task with a fixed decoder, and similar performance was achieved with a shuffled decoder compared to a biomimetic decoder, suggesting that how the neural activity is used in a BCI task does not have to be similar to how it is used in a comparable task using their arms, and animals can

learn to control individual neurons in a BCI task (Ganguly and Carmena, 2009). Later, they also showed that modulation of indirect neurons decreases during BCI-task training after the performance reached plateau. This indicates that while the initial learning exploit natural neural activity repertoire associated with movements, there is a mechanism to specifically assign credit to direct neurons and decrease indirect neuron modulation (Ganguly et al., 2011).

Potentially, animals may have various movements they can try, and choose one which evokes specific neural activity pattern leading to rewards. If movement repertoire is broad enough, there may be a movement associated with the specific activity pattern leading to the reward. Against this hypothesis, most BCI studies reported that the animals did not show overt movements, and some studies reported lack of muscle activity even with electromyograms (Carmena et al., 2003). It is still possible that both neural activity repertoire associated with movements and the repertoire explored and exploited during BCI training reflect the constraints of the neural network, and thus share the same property and structure. While the underlying mechanisms of how animals exploit the neural activity patterns associated with regular movements in a BCI task is not known, it has been shown that single neuron targeting is achieved by exploiting natural movement repertoire (Hwang 2013). In addition, BCI perturbation within intrinsic manifold is easier to learn than outside-manifold perturbation (Sadtler, Batista, 2014), suggesting that neural activity repertoire of movements is utilized during BCI learning.

In addition to monkeys and humans, people have recently started to study BCI in rodents using two-photon calcium imaging based on the advancement of molecular tools and imaging technologies. Clancy and colleagues first showed that calcium activity of neurons can

be differentially trained (Clancy et al., 2014). In their task, two ensembles of neurons were selected in an imaging field, and reward and auditory feedback was given based on the difference of the average activity of the two ensembles. They showed that the modulation of indirect neurons depends on the distance of them to the direct neurons in the ensembles. Unravelling the spatial distribution of recorded neurons is a benefit of two-photon calcium imaging, and its relationship to BCI training was first reported here. On the contrary, Hira and colleagues showed that when activity of a non-movement related neuron is associated with water delivery, the modulation of indirect neurons was not distance-dependent, but it was dependent on the activity correlation between the direct and indirect neurons (Hira et al., 2014). The difference between the studies may be partly attributed to the difference in the task, especially that the mice in Hira's study were pretrained in a lever-pull task and they targeted non-movement-related neurons in the BCI task.

Does the BCI readout have to be in the motor cortex? The BCI studies have been mainly targeted the motor cortex. This is partly because the first application of the BCI would be to substitute disabled motor output of tetraplegic patients, and the motor cortex encodes motor output and motor plans, which makes it an ideal candidate for decoding movement-related information. However, the potential of BCI does not have to be limited to substituting motor output. Neely and colleagues showed that the visual cortex can be trained as well, showing the possibility of BCI training based on abstract cognition (Neely et al., 2018). The visual cortex mainly receives external bottom-up inputs, but the design of the task indicates that the volitional modulation in the visual cortex reflects top-down feedback signal. BCI based on feedback signal from higher cognitive areas will expand the application of the technology in the future.

When the goal of a BCI task is to control a robot arm or a cursor on the screen, the animal constantly receives visual feedback during the task. However, it is not clearly shown that whether such constant feedback is necessary to learn a BCI task, or animals can learn to modulate the activity level simply based on reward timing. Fetz's study (1969) showed that auditory feedback makes learning quicker (Fetz, 1969). Clancy and colleagues showed that it is necessary to have auditory feedback for mice to learn the task (Clancy et al., 2014), while the task used in Hira et al does not involve sensory feedback. To further investigate how the sensory feedback is processed and used in BCI task learning, it is beneficial to substitute actual sensory stimulus with neural stimulation to elucidate the pathway and the necessary information in the feedback. Prsa and colleagues integrated optogenetic stimulation in the sensory cortex to a BCI task (Prsa et al., 2017). This approach would give us more insights about how the sensory feedback is used in learning a new BCI task by exploring different stimulation patterns and areas. It may further help develop bidirectional BCI with direct feedback to the brain, which may contribute to easier learning and smooth control of the prosthetics in the future.

Here, I have reviewed BCI studies and remaining questions. I will further focus on the relationship between BCI training and biological mechanisms, such as subtype-specific plasticity of inhibitory neurons and functional clustering of dendritic spines.

Function/anatomy of inhibitory neuron

One potential limitation of the previous studies is that they did not distinguish diverse cell types in the neocortex, or that they focused only on the excitatory neurons. In the cortex,

in addition to excitatory neurons, there are inhibitory neurons which release GABA at axon terminals and produce inhibitory postsynaptic potential on the synaptic targets. Different cell types possess distinct functions in cortical computations (Hangya et al., 2014) and likely distinct capacities to control brain-computer interfaces. Taking such biological characteristics of cell types into account is a potential way to improve BCI performance. First, I discuss functional and anatomical properties of different types of inhibitory neurons in the cortex, showing that they are good candidates to target in a BCI task in a subtype-specific way.

Inhibitory neurons can be further classified into various subtypes, according to the morphology, genetic expression patterns, and connectivity patterns. It has been reported that they can be classified into PV-positive, SOM-positive, and CCK-positive inhibitory interneurons depending on genetic expression patterns (Kawaguchi and Kondo, 2002). They can also be classified with their electrophisiological properties, e.g. fast-spiking and LTS neurons in layer 4 (Beierlein et al., 2003). How molecular, morphological, and electrophisiological classifications correspond to each other is further reviewed in (Rudy et al., 2011). Here, we focus on PV-positive, SOM-positive, and VIP-positive inhibitory neurons as discussed previously (Pfeffer et al., 2013).

Different classes of inhibitory neurons have different functions, as it has been shown in the visual cortex that they have different tuning properties. Hofer and colleagues showed that PV neurons in mouse visual cortex respond to visual stimuli non-specifically, and have a broad tuning for stimulus orientation/direction (Hofer et al., 2011). However, it has also been shown that this phenomenon is species-specific and that in ferrets, PV neurons show similar orientation selectivity to excitatory neurons, while SOM neurons have broader tuning curve (Wilson et al., 2017). These results show that the function of those inhibitory neurons may be

different among species and dependent on how the cortex is organized, e.g. columnar structure in ferret visual cortex and salt-and-pepper structure in mouse visual cortex. Still, it is clear that those inhibitory neurons respond to stimulus in a subtype-specific manner and that they are potentially playing role in functional computation in the cortex.

In addition to the tuning property of the inhibitory neurons, the correlation structure of the activity shows subtype-specific patterns. Karnani et al showed that correlation is high between SOM neurons and between VIP neurons, but not between SOM and VIP neurons in the mouse primary visual cortex (Karnani et al., 2016a). This indicates that each subtype may form its own functional network and that the network has a subtype-specific role.

Such different functionality may come from different connectivity patterns (reviewed in (Fino et al., 2013)). In rat visual cortex, fast-spiking interneurons and excitatory neurons form reciprocal connections, and a pair with reciprocal connection shares excitatory input patterns (Yoshimura and Callaway, 2005). Reciprocal connections between excitatory and PV neurons were also reported in layer 5 (Otsuka and Kawaguchi, 2009). PV and SOM neurons inhibit surrounding excitatory neurons densely, with high probability of connection (Fino and Yuste, 2011; Packer and Yuste, 2011). Karnani et al showed that VIP neurons inhibit PV and SOM neurons, and disinhibiting excitatory neurons (Karnani et al., 2016b). This disinhibitory circuit is known to be recruited in reward-dependent manner, suggesting a role in learning (Pi et al., 2013). VIP neurons also directly inhibit excitatory neurons, and direct inhibition and disynaptic disinhibition through SOM neurons are competing with each other (Garcia-Junco-Clemente et al., 2017). The type-specific inhibition patterns were thoroughly examined and it was reported that VIP mainly inhibits SOM, SOM inhibits PV, VIP and excitatory neurons, and PV inhibits PV and excitatory neurons (Pfeffer et al., 2013).

How are the subtypes different in terms of plasticity? There are multiple potential locations of plasticity. Both synapses onto inhibitory neurons (presynaptic) and inhibitory synapses from the neuron, where the inhibitory neurons connect to and inhibit the post-synaptic neurons (postsynaptic), can be plastic. In vitro, multiple studies reported presynaptic and postsynaptic plasticity (LTP, LTD, and STDP) in multiple types of inhibitory neurons with various stimulus configuration (Huang et al., 2013; Lu et al., 2007; Sarihi et al., 2008). The studies are further reviewed in (Kullmann and Lamsa, 2011; Kullmann et al., 2012; Lamsa et al., 2010).

In vivo, NPY+ neurons have dendritic spines where they receive excitatory inputs, and it was shown that removal of visual input induced reduction in the number of dendritic spines on the inhibitory neurons (Keck et al., 2011). They also showed that the number of inhibitory axon boutons decreased after visual input removal, which indicates the postsynaptic plasticity. Chen and colleagues showed that the number of axonal boutons of SOM neurons in the motor cortex decreased immediately after the initiation of the training of a motor task, and that of PV neurons gradually increased (Chen et al., 2015). These studies showed that such structural changes in presynaptic and postsynaptic happen on the inhibitory neurons in vivo and that the inhibitory neurons may play a role in learning and adaptation. Recently, it was shown that inhibitory neurons in the visual cortex show type-specific changes in selectivity and interaction with excitatory neurons after learning a visual discrimination task (Khan et al., 2018). This supports the idea of subtype-specific functions of inhibitory neurons in learning.

As has been discussed here, it has been shown that inhibitory neurons have subtype specific functional and anatomical properties, and it has been shown that inhibitory circuit is plastic and can contribute to learning and/or adaptation. It makes inhibitory neurons

promising candidates to target in BCI tasks. They can potentially show subtype-specific learning capacity different from excitatory neurons. How do inhibitory neurons behave when targeted in BCI tasks? Are they different from excitatory neurons? These questions were further addressed **Chapter 2.** Briefly, we tracked the plastic changes of three major types of cortical inhibitory neurons (INs) during a neuron-pair operant conditioning task using two-photon imaging. In each animal, one of the IN subtypes was expressing GCaMP6f, and two neurons in an imaging field was randomly selected to be a positive target neuron (N+) and a negative target neuron (N-). Each image from the microscope was processed with a custom software in real time, and mice were rewarded when the activity of N+ exceeded that of N− beyond a set threshold. Mice of all subtypes improved performance, but the strategies were subtype specific. When parvalbumin (PV)-expressing INs were targeted, the activity of N− decreased. However, targeting of somatostatin (SOM)- and vasoactive intestinal peptide (VIP)-expressing INs led to an increase of the N+ activity. These results demonstrate that INs can be individually modulated in a subtype-specific manner and highlight the versatility of neural circuits in adapting to new demands by using cell-type-specific strategies.

Clustered spines

It has been shown that animals can learn to modulate activity of individual neurons. When activity of a neuron increases by learning, is it due to increased activity of presynaptic excitatory neurons, or are the excitatory synapses onto the neuron potentiated? Such underlying mechanisms of BCI learning still have to be elucidated. In a motor learning task, it has been shown that dendritic spine formation and elimination are associated with learning

(Xu et al., 2009), and long-lasting spine clusters are formed through repetitive activation during learning (Fu et al., 2012). The dendrites receive excitatory inputs at dendritic spines, and nonlinear summation of inputs among nearby spines influences how the neuron responds. Does the spatial distribution of spines affect BCI learning? Does BCI training induce further spine formation and elimination? If so, where do they occur? Answering those questions will help understanding both the mechanisms of BCI learning and the plasticity of the brain.

Recording the activity of dendrite spines became possible with the development of imaging technique (Chen et al., 2011). They patched a cell to prevent it from spiking because action potentials propagate along the dendrites and induce calcium influx to the spines. Recently, it has been shown that there is a linear relationship between the signal from backpropagating action potential to calcium signal in the spine, and it can be subtracted linearly with the use of a newly developed calcium sensing protein (Chen et al., 2013). With this, it has been showed that spatial clustering of similarly tuned inputs predicts orientation tuning of the neuron in the primary visual cortex of ferrets (Wilson et al., 2016). In the mouse visual cortex, it has been shown that presynaptic neurons with higher receptive field correlation tend to form synapses spatially closer on the dendrite (Iacaruso et al., 2017). These results suggest that functionally related neurons form a spine cluster on the dendrite.

How does the spatial location of the synaptic input affect firing of the postsynaptic cell? First, EPSP produced at the synapse propagates to soma while being attenuated and filtered. When the cell receives inputs from multiple presynaptic neurons, the EPSPs will be overlapped and the summation of them determines whether the cell fires or not. If the dendrite is completely passive, the combined EPSP is linear summation of individual EPSPs. However, the dendrite has active component, such as voltage-dependent channels, and have

non-linear effect (London and Häusser, 2005). An apparent example is called ‘dendritic spikes’. When a section of a dendrite receives multiple inputs, the membrane potential of the dendrite surpasses the threshold of the voltage-gated channels. Opening of the channels further increases the membrane potential, and it appears as if the dendrite fires. That membrane potential is conducted through the dendrite, and the postsynaptic neuron is more likely to fire than expected from the linear summation of the individual EPSPs. For this to happen, it is important that the multiple inputs occur simultaneously within a physically confined section of the dendrite, requiring a functionally related spatial cluster of the spines on the dendrite. The area induces dendritic spikes works as a subunit for coincidence detection. In vivo, this has been reported to be beneficial in computation to enhance stimulus selectivity (Smith et al., 2013).

A remaining question is how the clusters are formed. While there is much to be investigated, autocrine BDNF signaling has been shown to be important in spine plasticity, and it can influence the plasticity of nearby spines (Harward et al., 2016; Hedrick et al., 2016). Such mechanisms might be contributing to reinforcing nearby spines with similar functional property, enhancing synaptic clusters.

I hypothesized that if we target spine-specific activity, the learning rate would be affected by how the target spines are spatially distributed along the dendrite, and the activity of non-target spines would change according to the spatial location of them relative to the target spines. I applied a modified neural feedback system to spine imaging and examined whether the mice can learn to modulate spine-specific activity (Chapter 3). The spine-specific activity should reflect the activity of presynaptic excitatory neurons, which can be other neurons in the motor cortex or neurons in the other cortical areas with long-range projection

to the motor cortex. Therefore, it was expected that mice could modulate spine-specific activity because it has been shown that excitatory neurons in the motor cortex can be modulated in a similar task (Clancy et al., 2014). However, we did not observe performance increase within a session or between sessions. While there are many potential reasons why the mice could not learn, but a few potential reasons are that the task may have been too difficult due to imperfect dendrite activity subtraction, synaptic failures, and/or larger noise associated with imaging with higher zoom. Due to this issue, we could not examine the relationship between spine clustering, spatial distribution of spines and BCI learning.

Implementation of real-time image processing

In Chapter 4, I examined different algorithms implemented for real-time image processing for motion artifact correction and calcium signal baseline estimation. Motion artifact correction has been an important step for post-hoc analysis of two-photon calcium imaging, but it has not been used in real-time, closed-loop experiments. Using either hill-climbing method with dense search and image pyramid method, or DFT-based template matching method of OpenCV, I showed that motion artifact correction of one image can be done in a few milliseconds, shorter than previously reported near-real-time algorithms (~20 ms/frame). I also introduced a new method of baseline estimation using kernel density estimate, and compared it with a percentile method and a robust mean method.

Conclusion

Brain computer interface has been studied and it has been reported that human tetraplegic patients can control robot arms with brain activity. However, the control of the robot arm is far from smooth and precise control of biological arms of healthy individuals. Current BCIs require training and learning how to use it, which means that it requires the plasticity of the brain. Better understanding and making proper use of the plasticity of the brain may lead to a BCI easy to learn without extensive training. A potential way to improve the performance of BCI is to take biological properties into account, e.g. subtype-specific plasticity, and/or spatial connectivity patterns on the dendrites.

First, we investigated a BCI task performance when inhibitory neurons were targeted. PV, SOM, and VIP neurons were targeted. We showed that animals can learn to do the task utilizing subtype-specific strategies. This will help us further understand the plasticity of the brain and develop a better BCI in the future.

Next, I investigated a BCI task targeting dendritic spine specific activity. Spine specific activity should reflect somatic activity of presynaptic neurons, and we hypothesized that mice can learn to modulate spine specific activity. However, we did not see learning-related changes in spine-specific activity. It can be due to technical difficulty in recording spine-specific activity, and/or synaptic failures.

Chapter 2. Brain-computer interface with inhibitory neurons reveals subtype-specific strategies

Summary

Brain-computer interfaces have seen an increase in popularity due to their potential for direct neuroprosthetic applications to amputees and disabled individuals. Supporting this promise, animals - including humans - can learn even arbitrary mapping between the activity of cortical neurons and movement of prosthetic devices (Ganguly and Carmena, 2009; Jarosiewicz et al., 2008; Moritz et al., 2008; Sadtler et al., 2014). However, the performance of neuroprosthetic device control has been nowhere near the limb control in healthy individuals, presenting a dire need to improve the performance. One potential limitation is the fact that previous work has not distinguished diverse cell types in the neocortex, even though different cell types possess distinct functions in cortical computations (Chen et al., 2015; Pfeffer et al., 2013; Pi et al., 2013) and likely distinct capacities to control brain-computer interfaces. Here, we made a first step in addressing this issue by tracking the plastic changes of three major types of cortical inhibitory neurons (INs) during a neuron-pair operant conditioning task using two-photon imaging of IN subtypes expressing GCaMP6f. Mice were rewarded when the activity of the positive target neuron (N+) exceeded that of the negative target neuron (N-) beyond a set threshold. Mice improved the performance with all subtypes, but the strategies were subtype-specific. When parvalbumin (PV)-expressing INs were targeted, the activity of N- decreased. However, targeting of somatostatin (SOM)- and vasoactive intestinal peptide (VIP)-expressing INs led to an increase of the N+ activity. These results demonstrate that INs can be individually modulated in a subtype-specific manner, and

highlight the versatility of neural circuits in adapting to new demands by using cell-type specific strategies.

Results and discussion

Water-restricted mice expressing GCaMP6f in PV-, SOM-, or VIP-expressing INs were trained in a neuron-pair operant conditioning task with two-photon calcium imaging (modified from (Clancy et al., 2014)). Briefly, two neurons in layer 2/3 of the primary motor cortex were randomly selected out of those that met predetermined activity criteria (Methods) and designated as N+ and N-. In some of the imaging fields (14 / 36), there were only two labeled neurons that met the activity criteria. The distance between N+ and N- was 191 ± 119 μm (Mean ± S.D., n = 36 pairs). The calcium signal of the targeted neurons was not saturated (Figure 2.5A-C). Mice were rewarded when the calcium activity of N+ exceeded that of N- by a set threshold (Figure 2.1A). The reward contingency based on the difference between N+ and N- activity ensured that mice could not solve the task by simply activating all neurons in the area simultaneously. During the trial, the difference of the calcium signal of the two neurons was transformed to create a dynamically frequency-modulated auditory feedback. After each reward, the activity of the targeted neurons had to return to baseline, which resumed the auditory feedback and initiated the next trial. The same neurons were targeted over 4-6 sessions, one session per day, with the same reward contingency. We note that a previous study with a similar approach that targeted neurons of unidentified cell types showed that auditory feedback was essential for the learning of the task (Clancy et al., 2014). However, in the current study we did not explicitly test the necessity of the auditory feedback.

For each of the three major IN types, mice were able to improve the performance over sessions, significantly increasing the reward frequency (PV-INs, $p = 0.011$, $n = 48$ sessions, 10 imaging fields in 5 mice, using fitlme in Matlab hereafter unless otherwise stated; SOM-INs, $p = 0.001$, 50 sessions, 10 imaging fields in 5 mice; VIP-INs, $p = 0.017$, 82 sessions, 16 imaging fields in 8 mice) (Figure 2.1B-D). Immunostaining showed a high degree of overlap between GCaMP6f-expressing and PV-expressing neurons in PV animals, and little overlap between GCaMP6f-expressing and PV-expressing neurons in SOM and VIP animals (Figure 2.1E-G). In a subset of the animals, we imaged the same neurons for an additional 1-3 contingency degradation sessions in which rewards were provided without regard to the activity of the targeted neurons. In these sessions, the targeted neurons reached the reward threshold significantly less frequently (the difference between the contingency degradation sessions and the last two training sessions, 2.49 ± 0.64 / min (estimate ± S.E., $p < 0.001$, $n =$ 12 and 14 sessions, 7 imaging fields in 4 mice), 3.07 ± 0.71 / min ($p < 0.001$, $n = 20$ and 20 sessions, 10 imaging fields in 5 mice), and 1.08 ± 0.38 / min ($p = 0.006$, $n = 28$ and 30 sessions, 15 imaging fields in 7 mice) in PV, SOM, and VIP animals, respectively). These results indicate that mice are indeed able to modulate the activity of IN subtypes.

We considered three possible strategies by which mice could achieve an increase in reward frequency (Figure 2.2A). For example, the reward frequency could increase with an increase in the activity frequency of N+ (Figure 2.2A, 'N+ increase'). Alternatively, a decrease in the activity of N- (Figure 2.2A, 'N- decrease') or decoupling of activity between N+ and N- (Figure 2.2A, 'decoupling') can improve the reward frequency by increasing the chance that activity in N+ leads to a reward. In the following analysis, we focused on the

frequency of calcium events because the amplitude of the events did not significantly change across sessions (Figure 2.5D-F).

When PV-INs were targeted (Figure 2.2B), the frequency of N+ calcium events did not change significantly, arguing against the 'N+ increase' strategy (Figure 2.2C). In contrast, the frequency of N- calcium events decreased ($p = 0.002$), supporting the 'N- decrease' strategy (Figure 2.2D). The slopes of the changes of N+ activity (Figure 2.2C) and N- activity (Figure 2.2D) were significantly different ($p = 0.001$), indicating the specificity of the decrease of N- activity. Accordingly, the frequency of the N+ calcium events that did not lead to a reward decreased ($p = 0.025$) (Figure 2.2E). The frequency of N- events during the time periods when N+ was inactive also significantly decreased over sessions ($p = 0.011$) (Figure 2.2F). This result argues against the 'decoupling' strategy, which would predict that N- events which originally coincided with N+ activity would move into the periods of N+ inactivity. Furthermore, the correlation coefficient between N+ and N- activity during task period did not decrease across sessions (slope estimate ± SE = 0.0145 ± 0.0089 /session, $p = 0.112$). To test whether the changes of neural activity contribute to the improvement in task performance, we conducted mediation analysis (Falk and Biesanz, 2016). In this analysis, we found that N+ activity and N- activity were positively and negatively correlated with task performance respectively (coefficients of a linear model: 0.57 ± 0.06 ($p < 0.001$) and -0.22 ± 0.06 ($p < 0.001$) (estimate ± S.E.)), and there was a significant mediation effect with N- activity decrease ($p < 0.001$) and not with N+ activity ($p = 0.233$). These results indicate that, when PV-INs were targeted, mice specifically decreased the activity of N- while maintaining N+ activity.

The task improvement in mice with SOM-IN targeting (Figure 2.3A) involved a different strategy. In contrast to PV-INs, N+ event frequency increased in the later sessions in SOM-INs ($p = 0.003$, Figure 2.3B). This suggests that the 'N+ increase' strategy was utilized to perform the task. Conversely, the frequency of N- events did not change (Figure 2.3C), nor did the frequency of N+ events that were not associated with rewards (Figure 2.3D), arguing against the 'N- decrease' hypothesis. The slopes of the activity changes of N+ (Figure 2.3B) and N- (Figure 2.3C) were significantly different ($p = 0.010$). Furthermore, the frequency of N- events during the periods of N+ inactivity did not increase (Figure 2.3E) and the correlation coefficient between N+ and N- activity during task period did not significantly decrease across sessions (slope estimate ± SE = -0.0009 ± 0.0119 /session, $p = 0.939$), discounting the 'decoupling' strategy. In the mediation analysis, N+ and N- activity were positively and negatively correlated with task performance, respectively (coefficients of a linear model: 0.60 ± 0.03 ($p < 0.001$) and -0.06 ± 0.02 ($p = 0.018$) (estimate ± S.E.)), and the mediation effect through N+ activity was significant ($p = 0.003$) and not through N- activity ($p = 0.082$). We conclude that mice improved the task performance with SOM-INs primarily by activating specifically the N+ neuron.

Next, we investigated VIP-INs (Figure 2.3F). As with SOM-INs, in VIP mice the event frequency of N+ significantly increased ($p = 0.044$) (Figure 2.3G). Neither the frequency of N- events nor the N+ event frequency not associated with reward changed significantly (Figure 2.3H-I), arguing against the 'N- decrease' strategy, although the difference in the slopes of activity changes of N+ (Figure 2.3G) and N- (Figure 2.3H) did not reach statistical significance ($p = 0.101$). The N- event frequency within the periods of N+ inactivity did not change (Figure 2.3J) and the correlation coefficient between N+ and N-

activity during task period did not change (slope estimate ± SE = -0.0103 ± 0.0091 /session, p = 0.261), excluding the 'decoupling' strategy. N+ and N- activity were positively and negatively correlated with task performance respectively (coefficients of a linear model: 0.70 ± 0.03 ($p < 0.001$) and -0.10 ± 0.04 ($p = 0.011$) (estimate ± S.E.)), and the mediation effect through N+ activity was significant ($p = 0.040$, Methods). These data demonstrate that VIP mice improved the task performance by increasing N+ activity.

To test whether the difference between cell types was significant, we examined whether cell type had a significant effect on the slope of the linear model (Methods, Figure 2.4A-D). The reward frequency increase was not different among three cell types. The frequency increase of N+ peaks was significantly larger in SOM neurons than in PV and VIP neurons. The decrease of N- peak frequency was greater in PV neurons than in VIP neurons, and the decrease of the frequency of N+ peaks without rewards was specific to PV.

For SOM and VIP mice, we further investigated whether N+/N- activity changed differently from the activity of non-target neurons which met the same activity criteria. We did not perform this analysis for PV mice because 6 out of 10 imaging fields did not have any non-target neuron that met the activity criteria, and the other 4 fields only had 1 non-target neuron. The result shows that N+ activity increase in SOM animals was specific to N+ and significantly greater than the non-target neurons (slope difference from non-target neurons: 1.81 ± 0.73 ($p = 0.014$) (SOM, N+, Figure 2.4E), -0.68 ± 0.77 ($p = 0.376$) (SOM, N-, Figure 2.4F), 0.43 ± 0.41 ($p = 0.304$) (VIP, N+, Figure 2.4E) and -0.12 ± 0.42 ($p = 0.770$) (VIP, N-, Figure 2.4F) (estimate ± S.E)). In addition, we simulated reward frequency increase as if each non-target neuron were used in the task as either N+ or N-, and the actual target neuron was used for the other target (for example, in a simulation in which we used a non-target neuron as

N+, reward frequency was simulated using that non-target neuron as N+ and the actual N-). Figure 2.4G shows that the simulated reward frequency increase with non-target neurons used as N+ is significantly lower than the actual reward frequency increase in SOM neurons, but not in VIP neurons. The reward frequency increase was not affected in either SOM or VIP animals when non-target neurons were used as N- (Figure 2.4H). The results show that the activity increase in SOM animals was specific to N+ leading to improved task performance.

To our knowledge, this is the first study to test the plasticity of individual neurons of molecularly-identified cell types in a brain-computer interface task. We demonstrate that cell type has a profound impact on the way by which performance improvement is achieved. For example, the activity of a SOM-IN could be increased without activating a second SOM-IN, similarly to a previous study that did not identify the cell type of the targeted neurons (and thus most of the targeted neurons were presumably excitatory) (Clancy et al., 2014). However, we found no evidence that the activity of a PV-IN could be increased without also activating a second PV-IN. Instead, the activity of a PV-IN could be reduced without inactivating a second PV-IN. The performance of brain-machine interfaces may improve in the future if such cell-type specific constraints on plasticity are considered (Shenoy and Carmena, 2014).

The differences in baseline activity levels may have partially contributed to the difference in strategies among subtypes. For example, if the baseline activity level of PV neurons is higher, it might be more difficult to increase N+ activity in PV-INs than in other subtypes. In addition, due to different calcium buffering in each cell, the relationship between spikes and GCaMP6f signals may be different from cell to cell, leaving the possibility for PV neurons to require more spikes to cause a calcium event.

Alternatively, the cell-type specific strategies may be partially explained by the differences in the levels of activity correlation within each cell type (pairwise correlation coefficients of pre-task activity between candidate neurons were: PV, 0.61 ± 0.84 (n_{pair} = 18); SOM, 0.14 ± 0.03 (n_{pair} = 40); VIP, 0.31 ± 0.01 (n_{pair} = 394), mean ± S.E.M., $p<0.001$ for all pairwise comparisons after removing the effect of event rate (Figure 2.6)). Nevertheless, we argue that the difference in correlation is an important reflection of their intrinsic properties.

The ability to improve task performance with PV-INs is particularly striking, given the high activity correlation between PV-INs. It has been shown that neural feedback tasks based on the difference of two neural ensembles is harder to learn if the activity of the two ensembles is correlated (Clancy et al., 2014). However, the animals could perform the task with a specific strategy.

What could be the potential mechanisms underlying cell-type specificity of strategies? At the cellular level, the performance of this task could be mediated by specific plasticity of N+ and/or N- neurons, such as plasticity of intrinsic excitability, and synaptic plasticity of inhibitory and excitatory synapses onto these neurons. There are likely differences among cell types in their ability for these plasticity mechanisms (reviewed in (Kullmann et al., 2012)). For example, repetitive correlated spiking induced spike-timing-dependent plasticity (STDP) in low-threshold spiking (LTS) interneurons (putative non-PV INs (Yavorska and Wehr, 2016)), while it induced long-term depression in fast spiking interneurons (putative PV-INs (Yavorska and Wehr, 2016)) (Lu et al., 2007). If excitatory synapses onto PV neurons are less likely to be potentiated when the neurons are targeted in the operant conditioning task, it can explain why the N- decrease strategy was employed with PV-INs. However, it has been shown that long-term potentiation can be induced in PV-INs with theta burst stimulation

(Sarihi et al., 2008), and synaptic plasticity is sensitive to the neuromodulatory state of the circuit (Huang et al., 2013) and behavioral context (Donato et al., 2013; Kuhlman et al., 2013). Therefore it remains unclear how the capacities for synaptic plasticity may differ across cell types in the intact brain during learning.

Another potential mechanism for the improved performance in the current task is by modulating the activity of neurons presynaptic to the targeted neurons. Different subtypes of inhibitory neurons receive inputs from different populations of excitatory neurons. PV-INs receive dense inputs from nearby excitatory neurons (Hofer et al., 2011), which suggests that nearby PV-INs share similar excitatory inputs and thus a specific increase of excitatory inputs to one PV-IN may be difficult. On the other hand, SOM- and VIP-INs receive excitatory inputs from largely non-overlapping populations (Karnani et al., 2016a). Differences in the strategy for the operant conditioning task may originate from the differences in the characteristics of the neurons providing excitatory inputs to the target neurons.

In addition, excitatory and inhibitory neurons form highly interconnected networks. In general, it is thought that VIP-INs inhibit SOM-INs, SOM-INs inhibit excitatory, PV- and VIP-INs, PV-INs inhibit PV-INs and excitatory neurons (Pfeffer et al., 2013), and excitatory neurons project to all four types. These connectivity patterns provide many possible pathways that could mediate the plasticity observed in the current study. For example, for SOM-INs, SOM → PV → excitatory → SOM provides a potential positive feedback loop, possibly underlying the increase of N+ activity in SOM-INs. Furthermore, the SOM → PV inhibition could underlie the decrease of N- activity in PV-INs. Future experiments are required to test these specific possibilities.

Finally, the activity of the inhibitory neurons can be associated with movements and sensory stimulus in a subtype-specific manner. A study showed that monkeys perform a brain-machine interface task by exploring and exploiting neural patterns associated with natural movements (Hwang et al., 2013). Through our visual observations, we did not identify overt behavioral strategies during task performance, similar to a previous study (Clancy et al., 2014). Nevertheless, examining the relationship between existing neural activity patterns and behavioral variables ("intrinsic neural manifold" and "intrinsic behavioral manifold" (Jazayeri and Afraz, 2017)) before training and how they change through training in each subtype will be of future interest. Furthermore, studies using neural feedback task with two-photon calcium imaging have reported that auditory (Clancy et al., 2014) or artificial sensory (Prsa et al., 2017) feedback was necessary for the successful learning of the task, while similar learning only with rewards as a feedback has also been reported (Hira et al., 2014). Future experiments can be aimed at examining whether sensory feedback was necessary to learn to modulate inhibitory neurons, and if so how the dependency is different among cell types. Lastly, we note that our contingency degradation experiments suggest that the behavioral performance was goal-directed. However, the lack of auditory feedback in the degradation experiments leaves room for other interpretations, such as that the mice might have been in a completely different behavioral state without the feedback. It will be of future interest to investigate the relationships between targeted cell types and behavioral strategies, dependence on sensory feedback, and whether the behavioral performance is goal-directed.

Acknowledgements

We thank A.N. Kim, K. O'Neil, L. Hall, T. Loveland and O. Arroyo for technical assistance, the Janelia GENIE Project for GCaMP6f, and members of the Komiyama lab, especially R. Hattori, N. Hedrick, E. Hwang, and H. Liu for comments and discussions. This work was supported by grants from NIH (R01 NS091010A, R01 EY025349, R01 DC014690, U01 NS094342, and P30EY022589), Pew Charitable Trusts, David & Lucile Packard Foundation, McKnight Foundation, New York Stem Cell Foundation, Kavli Institute for Brain & Mind, and NSF (1734940) to T.K. A.M. is supported by the Nakajima foundation.

Author contributions

T.K. and A.M. conceived the project and wrote the paper. Immunostaining was performed by M.D. and A.M. All other experiments were performed by A.M. and analyzed by A.M. and T.K.

Methods

Animals

All procedures were in accordance with protocols approved by UCSD Institutional Animal Care and Use Committee and guidelines of the US National Institutes of Health. All animals before water restriction were group housed, and during water restriction they were singly housed or group housed when all littermates were under water restriction. They were housed in disposable plastic cages with standard bedding in a room on a reversed light cycle (12h/12h). Experiments were typically performed during the dark period. With mice acquired from Jackson Laboratories, we generated double transgenic mice, PV-Cre

[JAX:017320](Hippenmeyer et al., 2005)::lsl-GCaMP6f [JAX:024105](Madisen et al., 2015), SOM-Cre [JAX:013044](Taniguchi et al., 2011)::lsl-GCaMP6f, VIP-Cre [JAX:010908](Taniguchi et al., 2011)::lsl-GCaMP6f, and triple transgenic mice VIP-Cre::ZtTA::TITL-GCaMP6f [JAX:024107](Madisen et al., 2015) to express the calcium indicator GCaMP6f in a specific subtype of inhibitory neurons. Each line was maintained in the original background and crossed to generate the double or triple transgenic mice for experiments. Double transgenic and triple transgenic mice expressing GCaMP6f in VIP-INs had similar expression levels and the results were combined. Average body weight at the beginning of the first session was 21.5 ± 1.3 g (PV, n = 9), 21.8 ± 4.1 g (SOM, n = 10), and 20.2 ± 2.9 g (VIP, n = 16) (mean ± S.D.).

Immunostaining and cell counting

30 μm-thick motor cortex coronal sections were prepared with a microtome (Thermo Fisher) and blocked with 10% normal goat serum, 1% bovine serum albumin and 0.3% Triton-X100 in PBS for an hour at room temperature. Immunostaining was then performed with 24-hr primary antibody incubation at 4 °C (anti-PV (rabbit, ab11427, abcam), 1:1000; anti-GFP (chicken, GFP-1020, Aves), 1:1000, diluted in PBS with 3% normal goat serum) and 2-hr secondary antibody incubation at room temperature (anti-chicken DyLight 488 (goat, SA5-10070, Thermo Fisher), 1:1000; anti-rabbit Alexa 594 (donkey, A-21207, Thermo Fisher), 1:1000, diluted in PBS). Stained sections were mounted with CC/Mount mounting medium (Diagnostic BioSystems) and imaged with Apotome.2 (ZEISS). PV and GCaMP6f quantification was performed manually using ImageJ. Representative sections (3 for each animal) were chosen, and in each section, one rectangle area (880 μm × 680 μm) was selected

for counting. We were not able to achieve reliable labeling with SOM and VIP antibodies (T-4103 rabbit anti-Somatostatin-14 antibody (Peninsula Laboratories International, Inc.), 366004 guinea pig anti-somatostatin-28 antibody (Synaptic Systems) and 20077 rabbit anti-VIP antibody (Immunostar), all in a variety of conditions).

Surgery

Surgical procedures were performed as previously described (Peters et al., 2014). Adult mice (6 weeks or older, male and female) were anesthetized with isoflurane and injected with Baytril (10 mg/kg), dexamethasone (2 mg/kg) and buprenorphine (0.1 mg/kg) subcutaneously to prevent infection, inflammation and discomfort. A custom head-plate was glued and cemented to the skull. Craniotomy (~3 mm) was performed over the right caudal forelimb area (300 μm anterior and 1,500 μm lateral from the bregma). Experiments were performed at least 7 days after surgery.

Imaging

Imaging was conducted with a commercial two-photon microscope (Bscope, Thorlabs) running Scanimage using a 16x objective (NIKON) with excitation at 925 nm (Ti-Sa laser, Newport). Imaging was conducted in awake animals. For calcium imaging, images (512×512 pixels covering 472×508 μm) were recorded continuously at 28 Hz. Before each session, images at 3 different planes (10 μm step in z-axis) were acquired as a reference. When ROIs were pre-defined in the previous session, it was warped with enhanced correlation coefficient maximization algorithm (Evangelidis and Psarakis, 2008) using the reference image. ROIs were not updated during imaging. During the task, each frame was

transferred to another Matlab instance, and lateral motion was corrected by locally maximizing correlation coefficient to the reference image with gradient descent and processed for neural feedback. It was implemented with a custom Matlab executable written in C++. Slow lateral and vertical drifts were detected by comparing correlation coefficients of the average of 200 motion corrected frames to the reference images and automatically corrected throughout imaging sessions.

Behavioral task

Before starting the neuron-pair operant conditioning task, mice first underwent 1 – 4 lick sessions when they received water drops at random intervals. These sessions allowed mice to get acclimated to the environment, and the data from calcium imaging in these sessions were used to select the target neurons. ROIs were manually selected to include the soma of neurons with at least one apparent activity event during the session. For real-time neuron-pair operant conditioning, time-varying baseline of fluorescence signal of each ROI is defined as the 8th percentile of the Savitzky–Golay filtered (order = 2, length = 9) fluorescence intensity in the preceding 200 frames. The baseline was used to calculate $\Delta F/F$, which was then filtered by taking 15-frame moving average. Fluorescence in the background pixels (defined as all the pixels more than 2 pixels away from any ROI and also more than 32 pixels away from the border of the reference image) was averaged and the background $\Delta F/F$ was subtracted from $\Delta F/F$ of each ROI to compensate for neuropil contamination. Among candidate ROIs whose $\Delta F/F$ exceeded 0.4 at a rate between 1.5 / min to 15 / min, two were randomly selected as target ROIs (N+ and N-) for neuron-pair operant conditioning task. A threshold for the task was determined so that the simulated reward frequency using the

calcium activity during lick sessions is approximately 4 / min (threshold was 0.52 ± 0.27, mean ± S.D.). Mice were then trained to modulate the activity of the target neurons with auditory feedback (adapted from (Clancy et al., 2014)). The 15-frame rolling average of the difference of ΔF/F of N+ and N- was exponentially transformed to auditory frequency between 1 kHz and 16 kHz. 4 kHz indicates the same ΔF/F between N+ and N-, and 16 kHz indicates it reached the threshold. The delay caused by image acquisition and processing was 130 ± 33 ms (mean ± S.D., n = 40000 frames from one session). Mice were given a water reward (~10 μl) when the auditory frequency reached 16 kHz, and the sound was kept at 16 kHz for 850 ms to indicate a reward delivery (reward tone). This is followed by inter-trial intervals (ITIs) without auditory feedback, which continued until the difference of ΔF/F became smaller than 0.3 times the reward threshold for the first time after 850 ms. In neuron-pair operant conditioning sessions, imaging was performed for 6 min first without auditory feedback which served as the pre-task period. After the pre-task period, the task with auditory feedback was initiated. Behavioral sessions continued up to 50 mins or until the mice appeared disengaged from the task, e.g. not responding to available water rewards, and it was manually determined by the experimenter. During the task, the animals were monitored online by the experimenter using IR camera, and most animals showed little to no overt movements during successful trials. After 4-6 sessions of operant conditioning, a subset of mice underwent 1-3 post-operant-conditioning contingency degradation sessions. In these sessions, reward and reward tone were given at the same intervals as in the last two days of neuron-pair operant conditioning sessions without auditory feedback. In some cases the image quality degraded over the sessions, in which case the experiments were terminated and data were excluded. Only the experiments that had at least 5 imaging sessions including both neuron-

pair operant conditioning and post-operant-conditioning contingency degradation sessions were included in the analysis. Some mice were trained consecutively with multiple imaging fields.

Quantification and statistical analysis

Data analysis

Data analysis was performed using Matlab R2014a.

Task-induced reward frequency:

Instead of choosing threshold based on pre-task activity of each session as in (Clancy et al., 2014), we used the same threshold over multiple sessions. Task-induced reward frequency change was calculated by taking the reward frequency between 3 and 9 mins after initiation of the task, and subtracting the simulated reward frequency during the pre-task period. We chose this time windows because it was when mice were considered to be engaged in the task in most sessions, and the main results were not sensitive to the choice of the analysis time window (Figure 2.7). The simulated performance during the pre-task period did not change over time (Figure 2.8).

Task-induced neural activity:

To delineate the strategies underlying the performance improvement in neuron-pair operant conditioning, in post hoc analysis, we computed the frequency of N+ peaks, N- peaks, N+ peaks not associated with rewards, and N- peaks during N+ inactivity. Calcium signal was considered as a peak if $\Delta F/F$ is the largest within the surrounding 1 sec window and larger than the fluorescence baseline (corrected as below) by the reward threshold used for the

neuron-pair operant conditioning task. Peak frequency was calculated during the time periods when the sound feedback was provided and the first 850 ms after that the sound feedback turned off after each reward. 850 ms was added to accommodate for the time that the fluorescence reaches a peak after exceeding the threshold. N+ peaks were considered associated with a reward if criteria reward was delivered within 1 sec window centered at the N+ peak. N+ inactivity was when the baseline-corrected ΔF/F was below the reward threshold.

ΔF/F baseline correction:

The distribution of ΔF/F tends to be skewed due to positive Ca events, and thus the mean of the ΔF/F distribution is not representative of the baseline of ΔF/F. To estimate the baseline of ΔF/F, which corresponds to the mean of ΔF/F during the inactive periods, we computed the mean of a truncated distribution defined recursively as below: First, we calculated the mean and the S.D. of the whole distribution during the pre-task period, and used mean ± 2 S.D. as the boundaries to make a truncated distribution. Next, we calculated the mean and the S.D. of the truncated distribution, and used mean ± 2.274 S.D. as the new boundaries. This step was repeated 30 times. The coefficient for S.D. was twice the reciprocal of the S.D. of Gaussian distribution ($\sigma = 1$) truncated at ± 2, which was chosen so that this method converges in one step when applied to a Gaussian distribution. For the baseline-corrected post hoc analysis described above, the estimated baseline was subtracted from the original ΔF/F.

Linear models

Task-induced frequency changes:

Task-induced reward frequency changes from all the sessions were fit by a multiple linear regression model. The model formula is:

y~1+i_session+(1|imaging field)+(i_session-1|imaging field),

where (1|imaging field) and (i_session-1|imaging field) indicate a random effect constant and a random effect slope term for each imaging field, i_session is a discrete variable representing the session number, imaging field is a categorical variable representing the identity of imaging fields, and y is reward or neural activity frequency adjusted to pre-task period. fitlme function of Matlab was used to fit the model and to test the coefficient of the session number against 0.

To test the decrease of frequency of reward threshold crossing in the contingency degradation sessions, reward frequency changes from the last 2 neuron-pair operant conditioning sessions and simulated reward frequency changes from the 1-3 contingency degradation sessions were fit by a multiple linear regression model with additional task category term:

y~1+task category+(1|imaging field),

where task category is a categorical variable indicating whether the session was a neuron-pair operant conditioning task or a contingency degradation session. The model was fit and the coefficient of the task category was tested against 0.

Comparison between cell types:

To examine the difference between two cell types, the following model was fit to the results from the two cell types:

y~1+i_session+i_session*cell type+(1|imaging field)+(i_session-1|imaging field).

When the coefficient of the interaction term is significantly different from 0, it means that the two types are significantly different.

Comparison between target types:

To examine the difference due to target types (N+, N- and non-target neurons), the following model was fit to the activity of all the neurons of the two target types:

y~1+i_session+i_session*target type+(1|imaging field)+(i_session-1|imaging field).

y indicates peak frequency (Figure 2.4E-F) and simulated reward frequency if each non-target neuron was used as N+ with the actual target neuron N- (Figure 2.4G), and if each non-target neuron was used as N- with the actual target neuron N+ (Figure 2.4H). Note that activity of individual neurons was used without averaging. When the coefficient of the interaction term is significantly different from 0, it means that the two types are significantly different.

Mediation analysis:

For mediation analysis, the following three linear models were used:

y ~ 1 + N+ + N- + (1 | imaging field) + (N+ | imaging field) + (N- | imaging field),

N+ ~ 1 + i_session + (1 | imaging field) + (i_session - 1 | imaging field),

N- ~ 1 + i_session + (1 | imaging field) + (i_session - 1 | imaging field).

N+ and N- indicates N+ and N- activity frequency.

Correlation coefficient:

Correlation coefficient of ΔF/F during pre-task period on the first session between neurons which met activity criteria was compared across cell types. To remove the effect of activity frequency, the following linear model was used:

Correlation coefficient ~ 1+mean activity frequency+cell type,

Where mean activity frequency is the mean of the activity frequency of the two cells of a pair, and cell type is a categorical variable representing a cell type of the pair.

Statistics

Linear models with random effect were fit to data using fitlme function in Matlab R2014, and the significance of the coefficient is calculated by fitlme using t-test. For mediation analysis, t-statistics from fitlme was used with Mediation Analysis P-value calculator with normal approximation (Falk and Biesanz, 2016). The experiments and the analysis were not blinded. Animal assignment was not randomized. Statistical sample-size estimation was not performed.

Code availability

Code for online image processing is available at https://github.com/amitani/2pNFB.

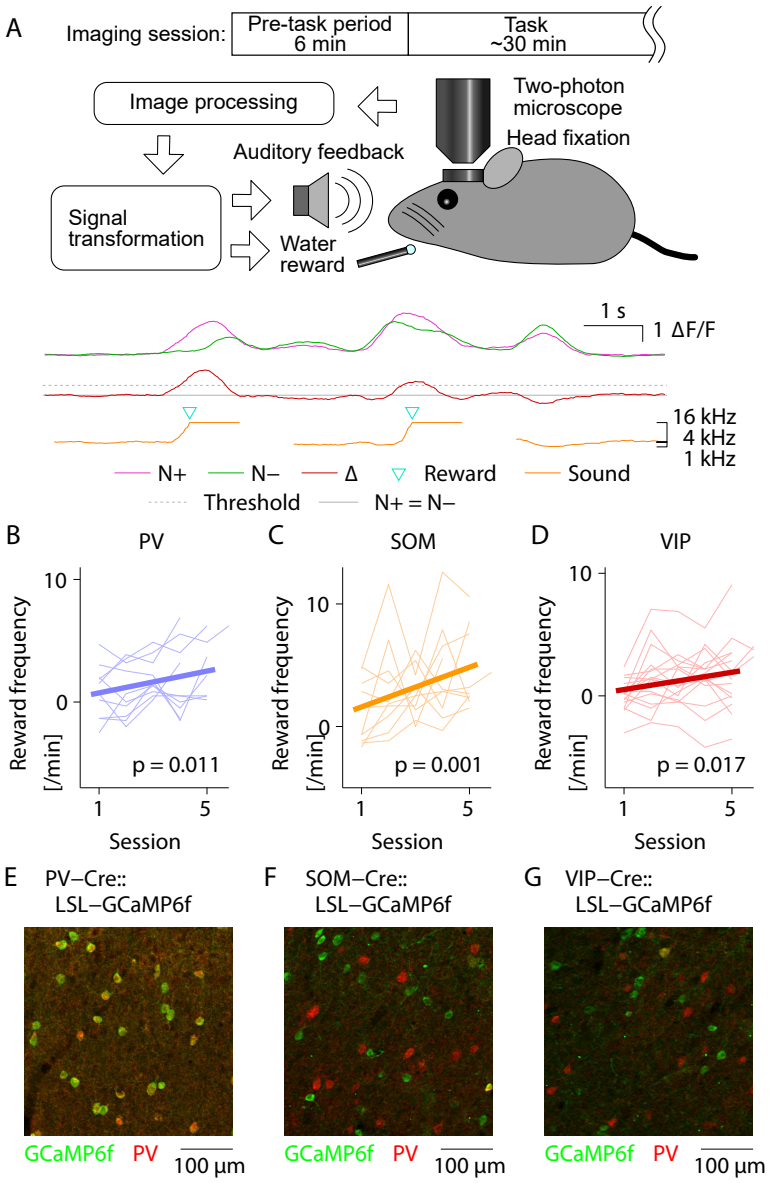

Figure 2.1 Mice improved performance in a neuron-pair operant conditioning task with IN subtypes.

(**A**) Schematic of task. Δ, N+ - N-. (**B**-**D**) Reward frequency, adjusted for chance-level performance during pre-task period in each session, across 4-6 sessions in PV (**B**, 0.44 ± 0.17 (slope estimate ± S.E.), n = 48 sessions, 10 fields in 5 mice, using fitlme in Matlab), SOM (**C**, 0.81 ± 0.24, 50 sessions, 10 fields in 5 mice), and VIP mice (**D**, 0.35 ± 0.14, 82 sessions, 16 fields in 8 mice). Thin lines: each imaging field; thick lines: linear fit. (**E**) Example PV (red) and GCaMP6f (green) labeling in PV amimals. In total, among 309 GCaMP6f-positive cells and 289 PV-antibody stained cells, 279 were double positive. (**F**) Same as (**E**) in SOM animals. Among 196 GCaMP6f-positive cells and 176 PV-antibody stained cells, only 7 were double positive. (**G**) Same as (**E**) in VIP animals. Among 164 GCaMP6f-positive cells and 202 PV-antibody stained cells, only 4 were double positive.

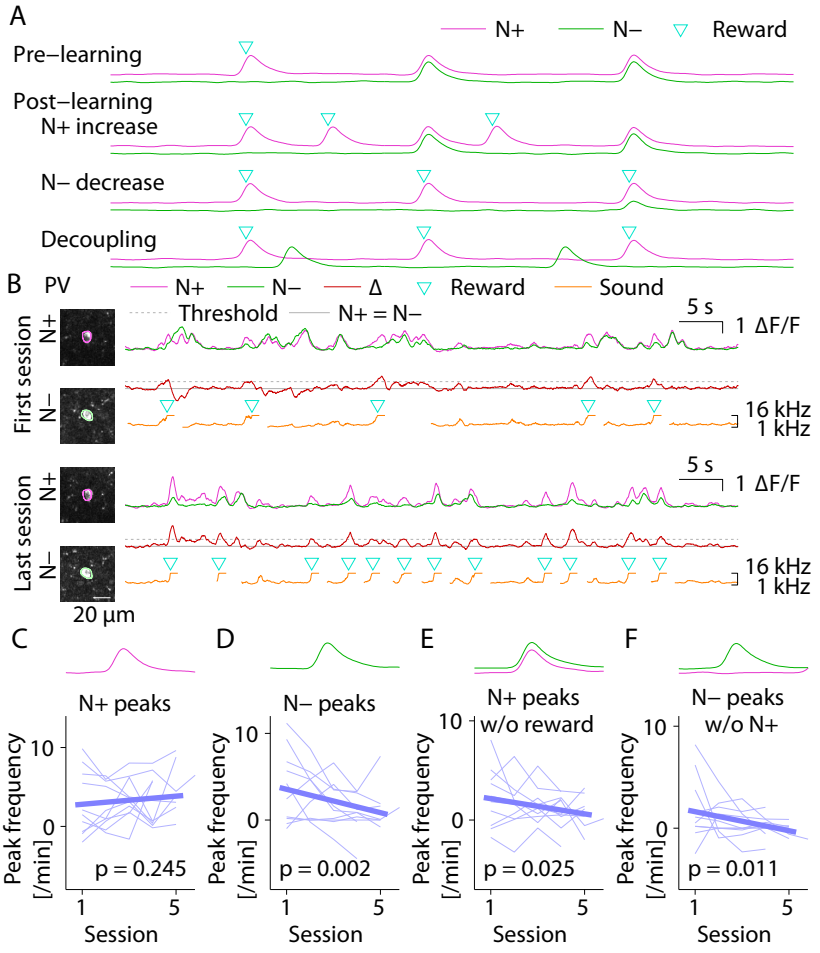

Figure 2.2 Neural activity modulation during PV-IN targeting.

(**A**) Schematic of three potential strategies to achieve the task. (**B**) Example activity traces of N+ and N-, activity difference (Δ), auditory feedback and reward timing in the first (top) and the last (bottom) session of a PV mouse. (**C**-**F**) Daily changes of the frequency of N+ peaks (**C**, 0.26 ± 0.2 (slope estimate ± S.E.), n = 48 sessions, 10 fields in 5 mice, using fitlme in Matlab), N- peaks (**D**, -0.68 ± 0.22), N+ peaks which were not associated with rewards (**E**, -0.39 ± 0.17), and N- peaks during the periods of N+ inactivity (**F**, -0.47 ± 0.18) during PV-IN targeting, relative to pre-task period.

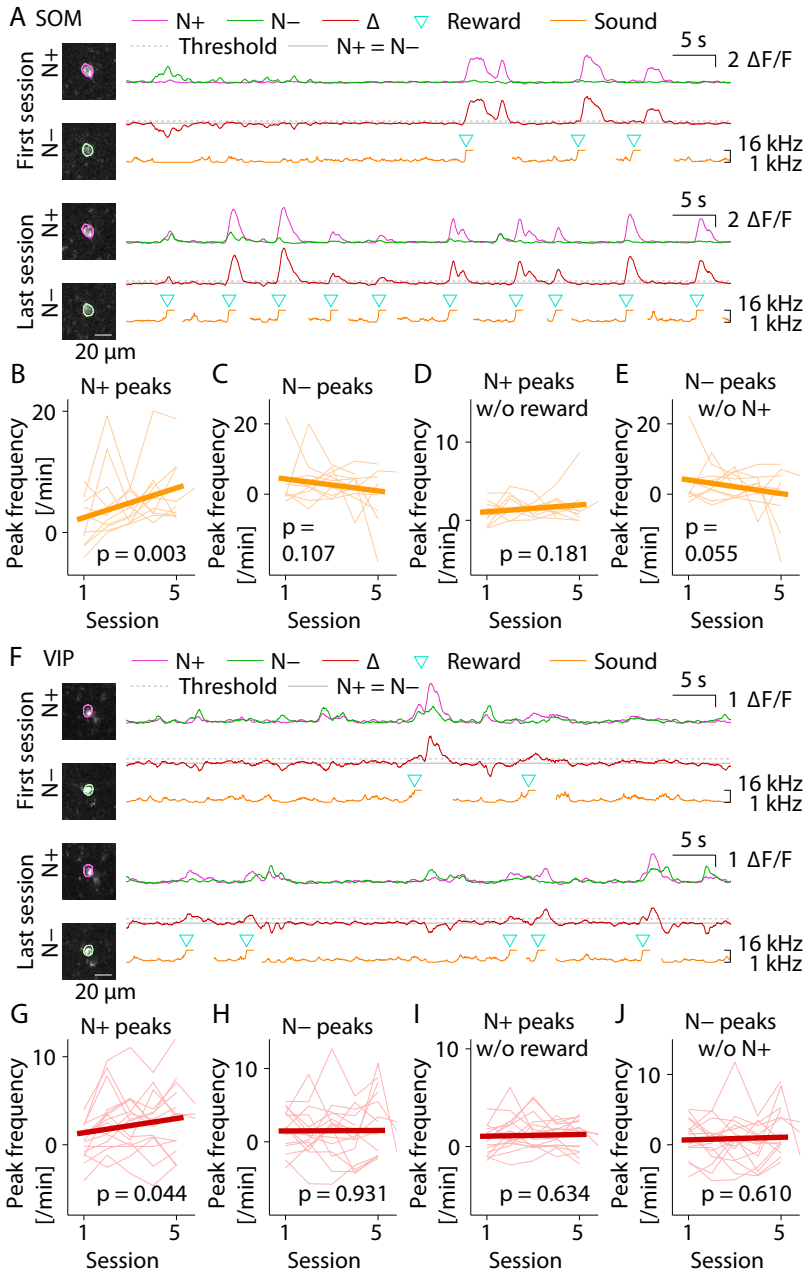

Figure 2.3 Neural activity modulation during SOM-IN and VIP-IN targeting.
(**A**) Example activity traces of N+ and N-, activity difference (Δ), auditory feedback and reward timing in the first (top) and the last (bottom) session of a SOM mouse. (**B-E**) Daily changes of the frequency of N+ peaks (**B**, 1.21 ± 0.39 (slope estimate ± S.E.), n = 50 sessions, 10 fields in 5 mice, using fitlme in Matlab), N- peaks (**C**, 0.84 ± 0.51), N+ peaks which were not associated with rewards (**D**, 0.22 ± 0.16), and N- peaks during the periods of N+ inactivity (**E**, -0.97 ± 0.50) during SOM-IN targeting, relative to pre-task period. (**F-J**) Same as (**A-E**) for VIP-INs (n = 82 sessions, 16 fields in 8 mice). Slope estimates ± S.E are 0.40 ± 0.20 (**G**), 0.02 ± 0.21 (**H**), 0.04 ± 0.09 (**I**), and 0.09 ± 0.18 (**J**).

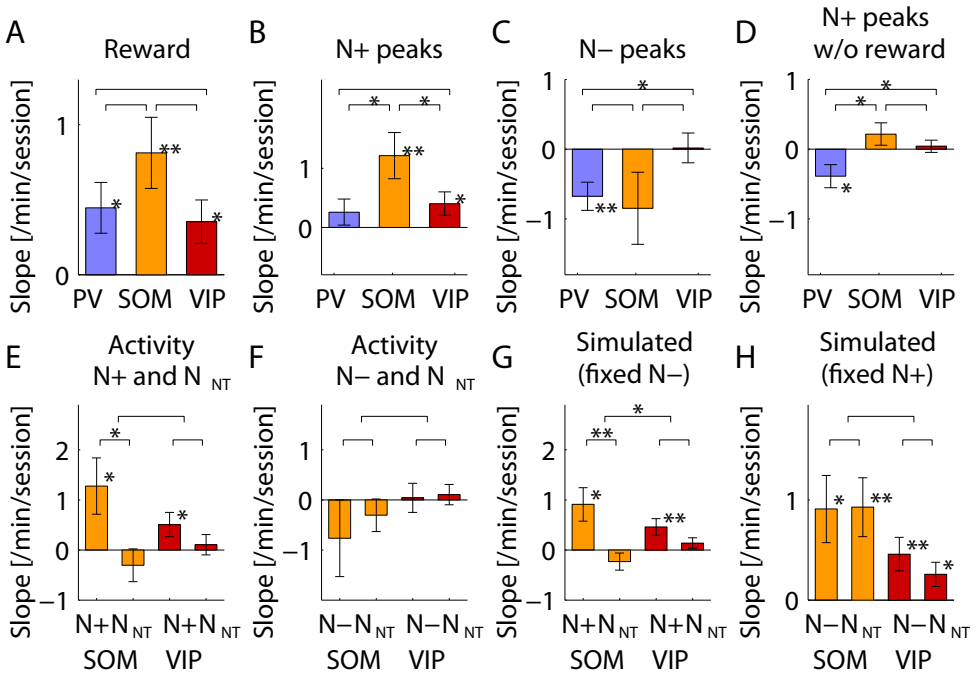

Figure 2.4 Difference between cell types in neural activity modulation.
(**A**-**D**) Difference between cell types in the slopes of the linear model. The frequency of rewards (**A**), N+ peaks (**B**), N- peaks (**C**) and N+ peaks which were not associated with rewards (**D**). (**E**-**F**) Difference between activity increase of target neurons (N+ (**E**) and N- (**F**)) and that of non-target neurons (N_{NT}). (**G**-**H**) Difference between actual reward frequency increase and simulated reward frequency increase if activity of non-target neurons were used instead of N+ (**G**) or N- (**H**). Bar plot and error bars indicate slope estimates ± S.E ((*) $p<0.05$, (**) $p<0.01$). * and ** between bars indicate a significant interaction term of a linear model ((*) $p<0.05$, (**) $p<0.01$). In (**E**-**H**), only imaging fields which had non-target neurons were included in the analysis.

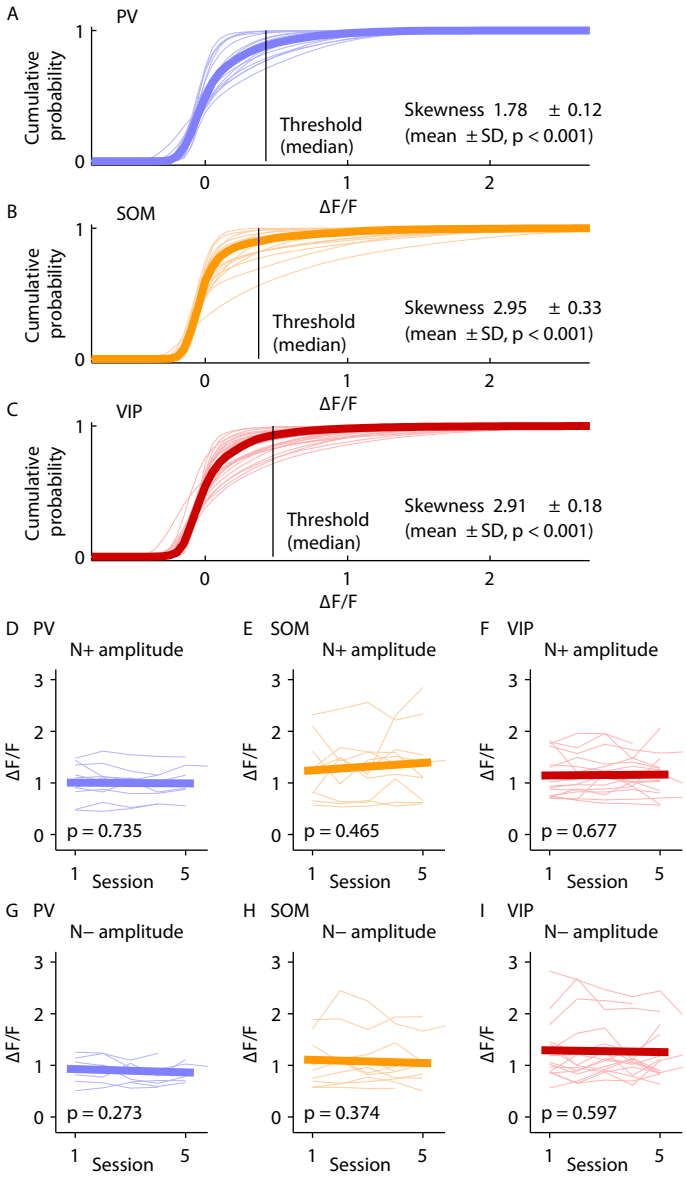

Figure 2.5 GCaMP6f signals did not show obvious signs of saturation.
(A-C) Cumulative probability of ΔF/F of the target neurons during neural feedback task during PV-IN (A), SOM-IN (B) and VIP-IN (C) targeting. Black line indicates the median of the reward thresholds. (D-I) Average peak amplitude of N+ neurons during PV-IN (D), SOM-IN (E) and VIP-IN (F) targeting and N- neurons in PV-IN (G), SOM-IN (H) and VIP-IN (I) targeting.

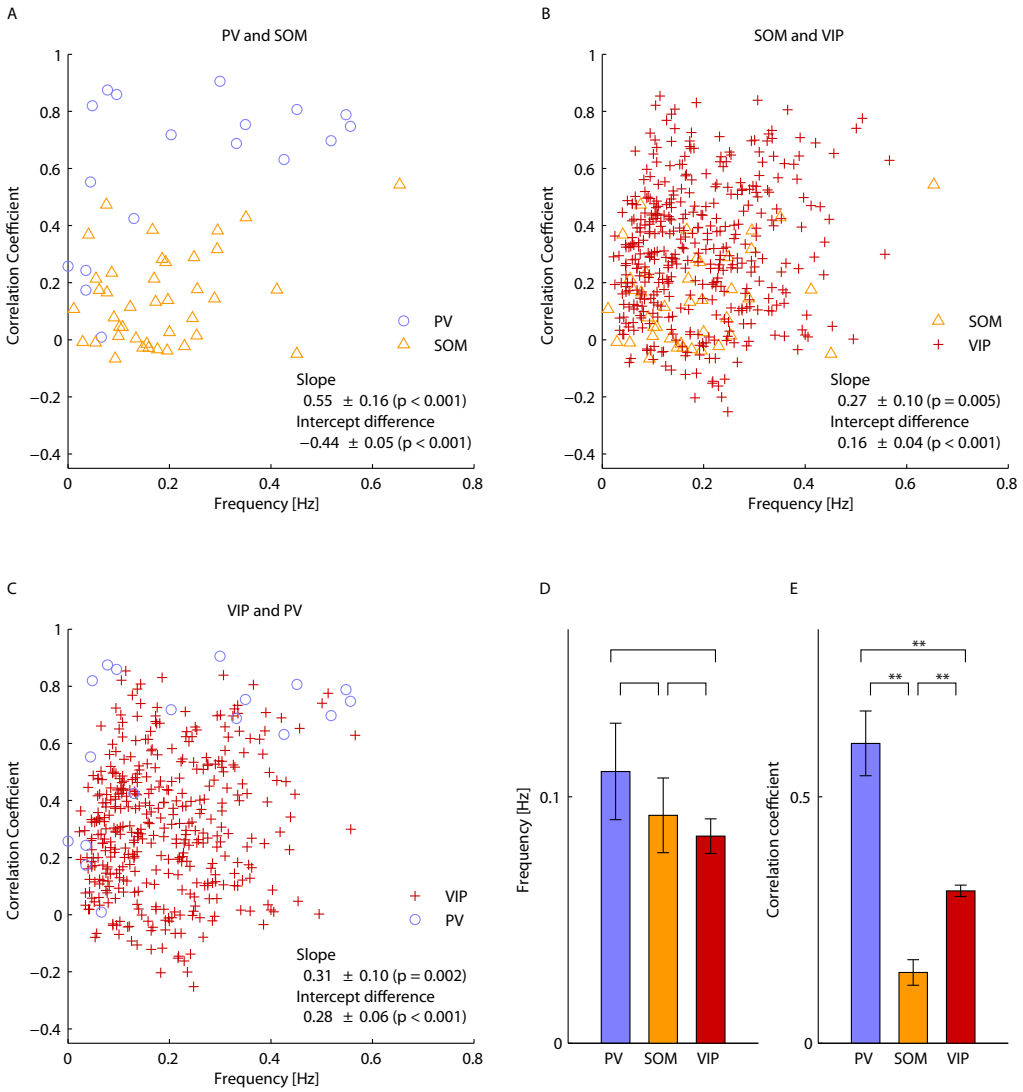

Figure 2.6 Event frequency and correlation coefficient.
(A-C) Average event frequency and pairwise correlation coefficient of ΔF/F of potential target neurons during pre-task period of the first session. Comparison between PV and SOM (A), SOM and VIP (B) and VIP and PV (C). (D) Event frequency of potential target neurons (mean ± S.E.). PV vs. SOM, p = 0.743; SOM vs. VIP, p = 0.588; PV vs. VIP, p = 0.903, Wilcoxon ranksum test. (E) Correlation coefficients of pairs of potential target neurons (mean ± S.E.). PV vs. SOM, p < 0.001; SOM vs. VIP, p < 0.001; PV vs. VIP, p < 0.001, Wilcoxon ranksum test.

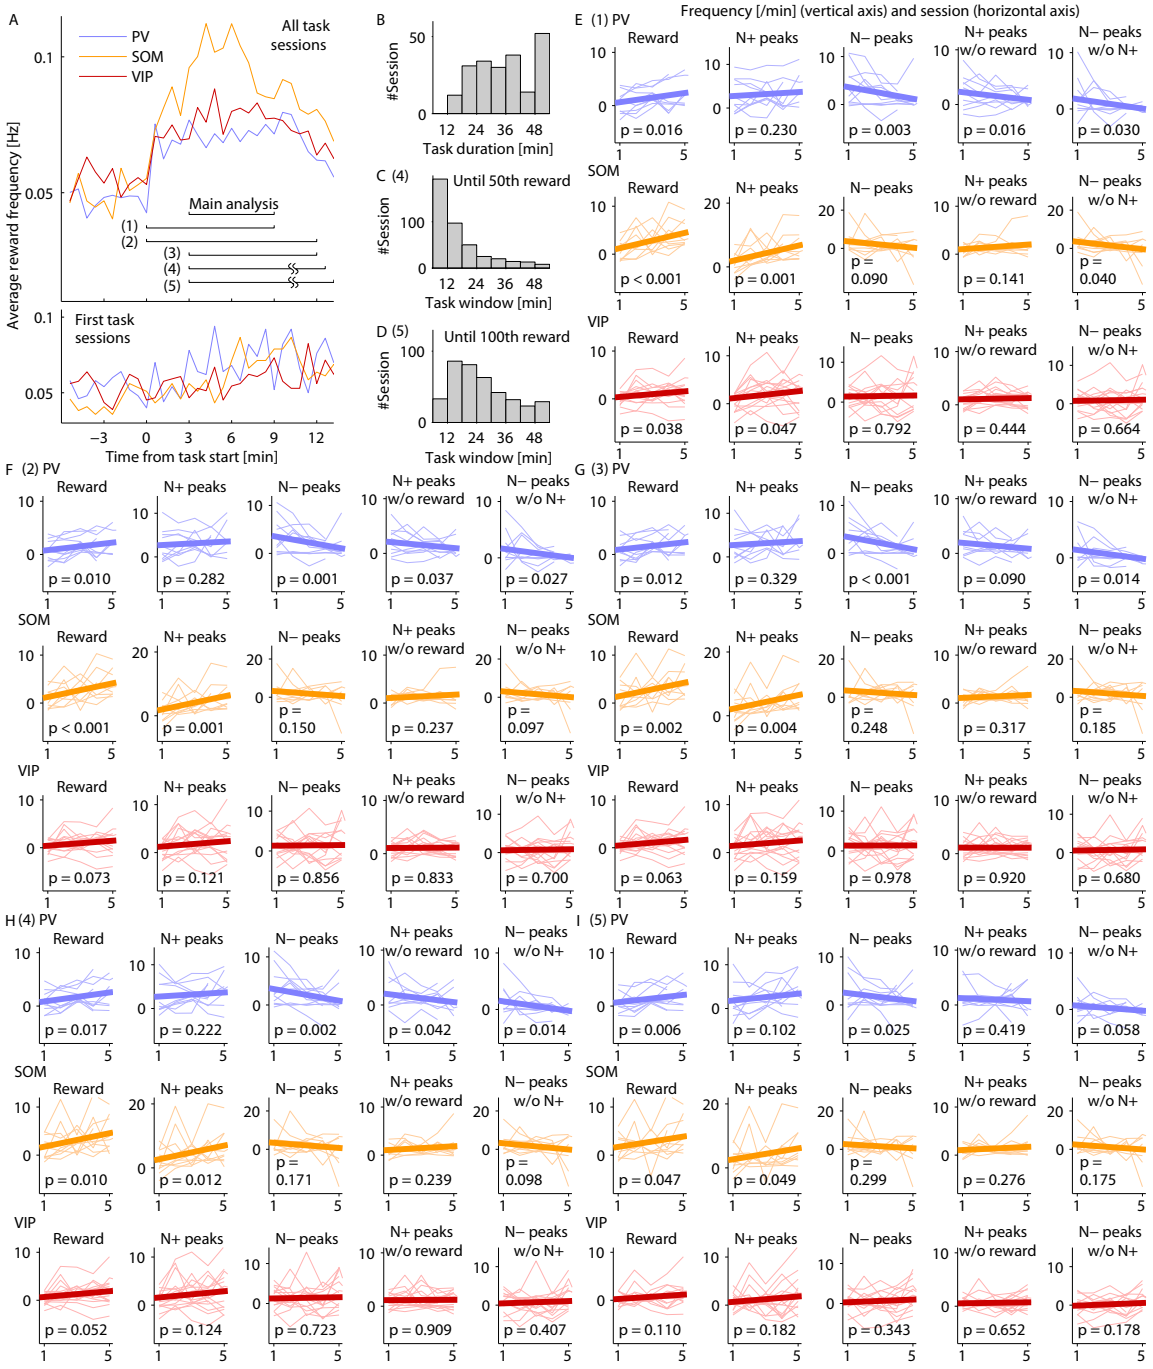

Figure 2.7 Task duration distribution and analysis with various time windows.
(A) Average reward frequency from all task sessions (top) and from the first task sessions (bottom). Horizontal bars indicate various task windows used in (E-I). (B) Task duration distribution. (C-D) Distribution of task window length if it is extended after 9 min until each mouse received the 50th reward (4) (C) or the 100th reward (5) (D). (E-I) Daily change of reward and activity frequency analysis using task time window (1) (E), (2) (F), (3) (G), (4) (H) and (5) (I).

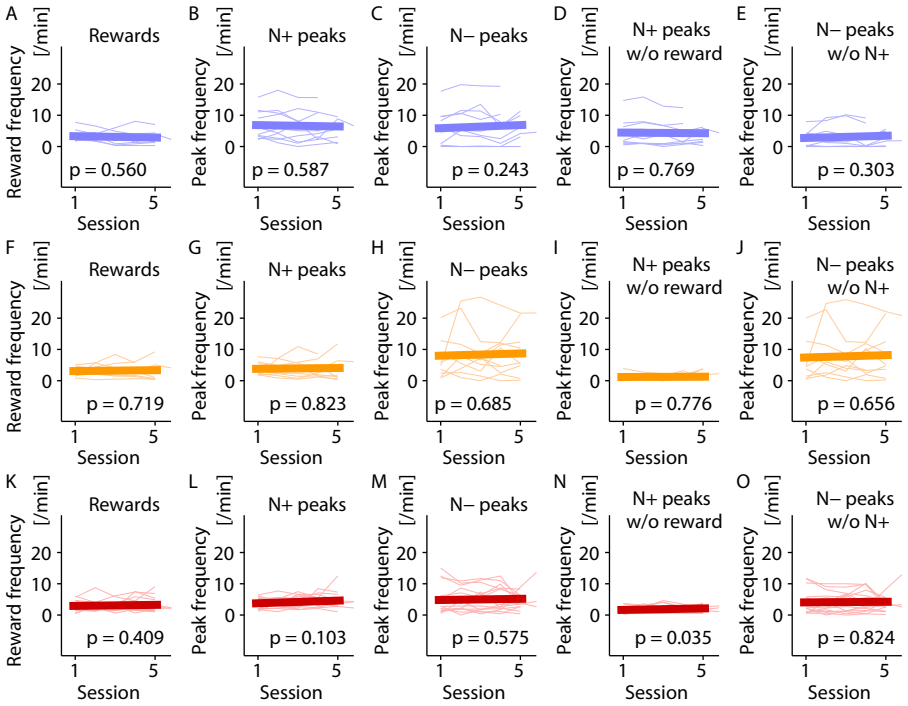

Figure 2.8 Simulated reward and activity frequency in pre-task period did not change. (A-E) Daily changes of the frequency of simulated rewards (A), N+ peaks (B), N- peaks (C), N+ peaks which were not associated with rewards (D), and N- peaks during the period of N+ inactivity (E) during pre-task period of PV-IN targeting. (F-O) Same as (A-E) during SOM-IN targeting (F-J) and VIP-IN targeting (K-O).

Chapter 2 is material previously published. Mitani A, Dong M, and Komiyama T. Brain-Computer Interface with Inhibitory Neurons Reveals Subtype-Specific Strategies. Current Biology, 2018. The dissertation author was the primary investigator and author of this material.

Chapter 3. Brain-computer-interface task with spine-specific activity

Introduction

It has been shown that animals can learn to modulate activity of individual neurons. When activity of a neuron increases by learning, is it due to increased activity of presynaptic excitatory neurons, or are the excitatory synapses onto the neuron potentiated? Such underlying mechanisms of BCI learning still have to be elucidated. In a motor learning task, it has been shown that dendritic spine formation and elimination are associated with learning (Xu et al., 2009), and long-lasting spine clusters are formed through repetitive activation during learning (Fu et al., 2012). The dendrites receive excitatory inputs at dendritic spines, and nonlinear summation of inputs among nearby spines may determine how the neuron respond. Does the spatial distribution of spines affect BCI learning? Does BCI training induce further spine formation and elimination? If so, where do they occur? Answering those questions will help understanding the mechanisms of BCI learning.

When a dendrite is imaged, the spines in the imaging field may not be responsible for learning. Instead of targeting the dendritic activity, I decided to target spine-specific activity. I hypothesized that if we target spine-specific activity, the learning rate would be affected by how the target spines were spatially distributed along the dendrite, and the activity of non-target spines would change according to the spatial location of them relative to the target spines. Here, I applied a modified neural feedback system to spine imaging and examined whether mice can learn to modulate spine-specific activity.

Methods

Animals

All procedures were in accordance with protocols approved by UCSD Institutional Animal Care and Use Committee and guidelines of the US National Institutes of Health. All animals before water restriction were group housed, and during water restriction they were singly housed or group housed when all littermates were under water restriction. They were housed in disposable plastic cages with standard bedding in a room on a reversed light cycle (12h/12h). Experiments were typically performed during the dark period. All the animals were b6 background.

Surgery

Surgical procedures were performed as previously described (Mitani et al., 2018). Adult mice (6 weeks or older, male and female) were anesthetized with isoflurane and injected with Baytril (10 mg/kg), dexamethasone (2 mg/kg) and buprenorphine (0.1 mg/kg) subcutaneously to prevent infection, inflammation and discomfort. A custom head-plate was glued and cemented to the skull. Craniotomy (~3 mm) was performed over the right caudal forelimb area (300 μm anterior and 1,500 μm lateral from the bregma). A mixture of AAV1.Syn.Flex.GCaMP6f (1:5000-10000 final dilution) and AAV1.CMV.PI.Cre (1:2 final dilution) diluted in saline was injected 20-30 nL at 3-5 sites (~250 μm depth, ~500 μm apart). Experiments were performed at least 7 days after surgery.

Behavioral task

Before starting the spine activity operant conditioning task, we selected an imaging field which contains a dendrite with dendritic spines. First, mice underwent 3-5 min lick task.

The data from calcium imaging during the lick task were used to manually select spine and dendrite ROIs. Most of the chosen dendrites showed backpropagating action potentials. The neurofeedback task used here was based on the task described in Chapter 2. Briefly, ΔF/F was calculated as described in Chapter 2 or using kernel density estimate (subchapter 3.2.3). Dendritic ΔF/F was defined as 50 or 75 percentile of all the dendrite ROI ΔF/Fs, and it was subtracted from spine ROI ΔF/Fs to calculate spine-specific activity component of ΔF/F. Half of the spine ROIs were randomly chosen as target spine ROIs. When spine-specific activity exceeded 1 at any target spine ROI, animals were rewarded with a water drop (~10μl). The maximum of the spine-specific activity of the target spine ROIs was exponentially transformed to auditory frequency between 1 kHz and 16 kHz. 4 kHz indicates no spine-specific activity.

Motion correction

The previous study (Peters et al., 2014) used post-hoc motion correction by TurboReg. This is an ImageJ plugin written in Java, and it took ~10 times longer to process a movie than its actual time (512x512 pixels, 30 Hz). This plugin finds dislocation which maximizes mutual information between the template image and the shifted image. Maximizing mutual information is beneficial in a situation with different brightness compared to maximizing correlation coefficient, but the problem is that it requires more computation. To run motion correction in real time, we switched to maximizing correlation coefficient, and made a custom function with C++ to call from Matlab. In this algorithm, it starts from zero dislocation, and tries to find a direction in which a shift by a pixel increases correlation coefficient most. This procedure was repeated until it reaches local maximum. To avoid repeating too many times

when the required dislocation is large image pyramid method was used so that the downsampled image was motion corrected first, and the estimated dislocation was used to initialize the dislocation in this hill-climb method.

Later, I switched to use OpenCV library for template matching. This enables calculating correlation coefficient for all possible displacement correction and choosing the one with the highest correlation coefficient in real-time. It is especially beneficial in this spine imaging with high magnification, because the number of pixels of the displacement from motion can be large, and hill climbing can be unstable.

Kernel density estimate for baseline estimation

We used an 8-percentile method for baseline estimation in the experiments in the previous chapter, and the first half of the experiments in this chapter. However, this method has multiple shortcomings. In calcium imaging, the calcium signal tends to have symmetric noise and sparse positive activity. When the positive activity is sparse, 8 percentile underestimates the center of the baseline. The degree of underestimation is dependent on the sparseness of the activity. With more activity, the estimate becomes closer to the true baseline, which means that the estimate is influenced by the activity level. Lastly, when the baseline increases, the increase of the estimate is slower. On the other hand, when the baseline decreases, the decrease of the estimate is faster. According to the increasing or decreasing trend of the baseline, there is different degree of delay in the estimate of the baseline.

To overcome these, I developed a new algorithm to estimate the baseline using kernel density estimation. Kernel density estimation is a method to estimate kernel density from limited number of samples under the assumption that the kernel density function is smooth.

Kernel density is a probability distribution from which each sample is produced. After the estimation, under the conditions where the baseline distribution have higher peaks than the distribution of intensities during calcium events, the peak of the kernel density provides a good estimate of baseline intensity. To make this peak sharper and higher, we conducted binning and averaging first (20 frames), and subsequently KDE was performed using ksdensity function of Matlab (R2011b, R2014a). It is done by convolving a Gaussian kernel of width optimized to samples. Taking the peak of kernel density of a continuous distribution is comparable to taking the mode of a discrete distribution. Because of this, there is little bias from increased activity, if there is enough baseline period in the window.

The script to process calcium intensity at each frame and estimate baseline is available at (https://github.com/amitani/baseline_kde)

Results

First, we examined how the animals learn in a spine-based BCI task. In this task, animals receive a reward when spine-specific activity of any target spine ROI (Methods) exceeds a set threshold. The reward rate did not increase statistically significantly during the first session (Figure 3.1A (left), $p = 0.21$). To see if there is any activity change specific to target spines, we examined simulated reward rate calculated from non-target spines. The trend of reward rate and simulated reward rate were not significantly different (Figure 3.1A (right), $p = 0.91$). The overall rate was different between actual and simulated reward rate ($p = 0.01$), but not when the simulated reward rate during baseline period was subtracted (Figure 3.1B, $p = 0.99$).

While the performance did not increase in a single session, it is possible that the performance increases across days. This can appear in two ways. First, the performance can be higher from the beginning of the task on day 2. Second, the previous experience in day 1 can help animals to learn the task quicker on day 2. In this case, the changing rate of the performance increase during the session is higher on day 2 than on day 1. However, there was no overall increase in performance from day 1 to day 2 ($p = 0.14$), and there was no increase in the increasing trend during the session from day 1 to day 2 ($p = 0.79$) (Figure 3.2A).

What can potentially prevent animals from learning? As shown in the figure, there was initial increase of activity up to 5 min right after the task initiation (Figure 3.1A). This can be due to state change associated with the beginning of the task, e.g. excitement and attention, or some artifact. To minimize the effect of this, I split the task into 3 blocks. Also, we have been the 8-percentile method for baseline estimation, but it turned out that the spine activity signal is noisy and sometimes this method is not appropriate because it underestimates baseline. To overcome this issue, I started to use a new method based on kernel density estimation. With these two modifications, we repeated the same experiments. However, we did not observe any performance increase during a session (Figure 3.1C, D) and across two sessions (Figure 3.2B).

There was no performance increase observed on day 2, but it is still possible that there is something, e.g. image quality, negatively affecting the performance on day 2 and learning is just compensating for it. To answer this, we conducted contingency degradation tasks to have a control about the effect of imaging quality on day 2. If the task is more difficult due to degraded imaging quality on day 2, and learning is just enough to compensate for the loss from it, we may not see any performance increase on the second day of two consecutive

sessions. If this is the case, the simulated reward rate in contingency degradation session would decrease showing the difficulty of the task on day 2. However, we did not observe performance decrease ($p = 0.07$) or decrease in increasing trend on day 2 ($p = 0.72$) (Figure 3.2C).

From those results, we concluded that mice could not improve the task performance in this specific task. This does not exclude the possibility that they can learn this task in different conditions.

Discussion

Here, we examined the performance of an unprecedented BCI task specifically targeting spine-specific activity. The animals did not show performance increase during the first session nor across two sessions. If the spine specific activity reflects activity of the presynaptic neurons, why could not mice learn the task? It has been shown that they can learn to control optically recorded activity of neurons in the motor cortex(Clancy et al., 2014; Hira et al., 2014). The main difference between directly targeting somatic activity of excitatory neurons and spine specific activity is the probability of detecting an event. Spine is a very small structure, and the amount of GCaMP in there is limited. The image has to be magnified further, and image tend to be dimmer. Therefore, just reading the calcium concentration in the spine can be noisy compared to that in soma. Because of this noise, at least one of false alert rate and miss rate should go up. Another source of noise is dendritic activity subtraction (Chen et al., 2013). If the subtraction is too much, some spine inputs may be masked. On the other hand, if the subtraction is not enough, activity of postsynaptic neuron may contaminate

spine activity measurement. Finally, there is biological source of randomness in this measurement. The synaptic release is probabilistic (Hessler et al., 1993), and there are cases when presynaptic neural firing does not lead to synaptic vesicle release. The spine calcium measurement only detects calcium influx in the spine, and it does not increase without synaptic vesicle release.

From these randomness, the correlation of reward timing and activity of the presynaptic neuron can be assumed to be lower when we target spine specific activity than when we target somatic activity of the presynaptic neuron. This could have prevented the animal from learning this task.

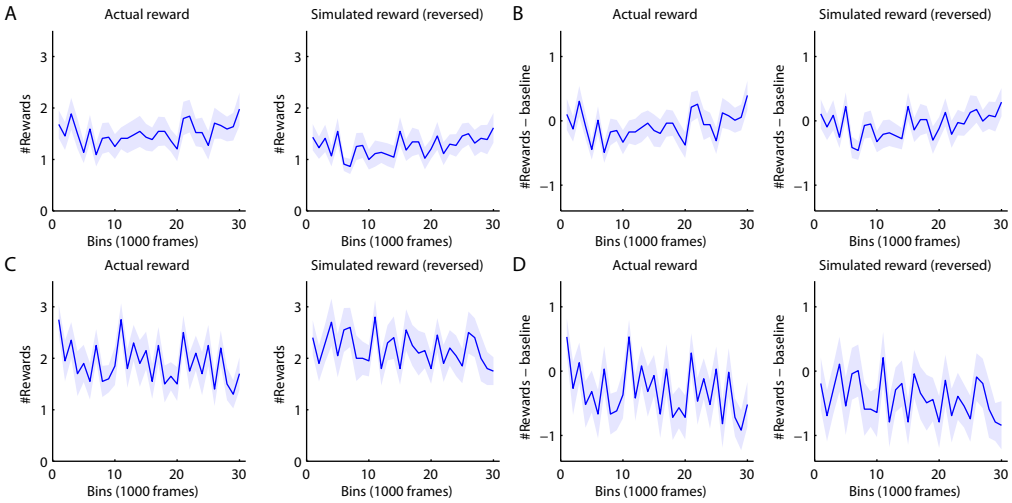

Figure 3.1 Time course of average reward rate in a session
(A) Time course of average reward rate on day 1. (B) Time course of average simulated reward rate on day 2. Simulated reward rate was calculated with non-target spines. (C, D) Same as (A, B) after modifying the protocol with 3 block structure and KDE-based baseline estimation. Shaded area indicates standard error.

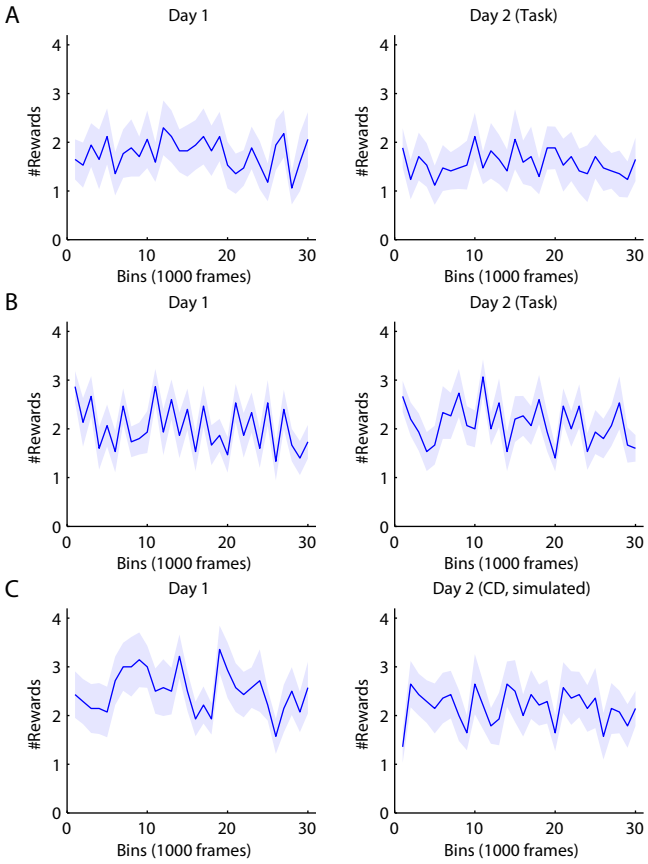

Figure 3.2 Time course of average reward rate on day 1 and 2
(A) Time course of average reward rate on day 1 (left) and day 2 (right). (B) Same as (A) after modifying the protocol with 3 block structure and KDE-based baseline estimation. (C) Time course of average reward rate on day 1 (left) and average simulated reward rate on day 2 (contingency degradation, right). Shaded area indicates standard error.

Chapter 4. Real-time processing of two-photon calcium imaging with lateral motion artifact correction

Abstract

Two-photon calcium imaging has been extensively used to record neural activity in the brain. It has been long used solely with post-hoc analysis, but the recent efforts began to include closed-loop experiments. Closed-loop experiments require fast, real-time image processing. When imaging awake animals, one of the crucial steps of post-hoc image analysis is correction of lateral motion artifacts. In most of the closed-loop experiments, this step has not been implemented and ignored due to technical difficulties. We recently reported the first experiments with real-time processing of calcium imaging that included lateral motion correction. Here, we report the details of the implementation of fast motion correction and present performance analysis. Additionally, we introduce a novel method to estimate baseline calcium signal using kernel density estimate.

Introduction

Two-photon calcium imaging has been widely used to image the activity of neurons in awake behaving animals. Neurons are loaded with a calcium-sensitive dye or, more commonly, made to express a genetically-encoded calcium indicator, such that their fluorescence signal reflects spiking activity of the neurons. The recorded movie often entails motion artifacts caused by body movements of the awake behaving animals. Of motion artifacts, lateral motion can be computationally corrected. As such, lateral motion correction has been a crucial step in processing calcium images from awake behaving animals. After

motion artifact is removed, regions of interest (ROIs) are defined, and average fluorescence intensity of the pixels in each ROI is calculated to report the spiking activity of the cell.

TurboReg (Thévenaz et al., 1998) has been widely used to correct lateral motion artifacts. It utilizes a pyramid approach, and it first constructs an image pyramid of series of downscaled images. The transformation at the final resolution is obtained by optimization using a transformation estimated with a downscaled image as an initial value, and this step is repeated recursively several times.

Two-photon microscopes generally use the excitation light to scan across the sample, and movement during scanning can not only shift but also distort the image, because each pixel or line is scanned at a different time. Several methods based on hidden Markov model (Dombeck et al., 2007) and optical flow (Greenberg and Kerr, 2009) have been reported to correct distortion caused by slow scanning rate. However, with higher scanning rate that has become more common with resonant scanners, within-frame distortion has become negligible.

TurboReg can correct motion artifacts with transformation up to 4 landmarks, but it could be beneficial to align more than 4 landmarks when imaging a larger field of view. Recently, NoRMCorre (Pnevmatikakis and Giovannucci, 2017) has been reported to estimate non-rigid transformation for such application. This method splits imaging field into overlapping patches, estimates translation of each patch, and upsamples the dislocation to obtain translation at each pixel. This method requires each patch to contain enough spatial signal to enable frame-by-frame alignment, and may not be applicable when labelling is sparse or weak. If some patches do not contain enough spatial features, the alignment of the patch may be unstable and affects the registration of nearby pixels. Non-rigid transformation may not be necessary with resonant scanners due to fast scanning speed, even though it may

be problematic with higher zoom that can be affected by small within-frame displacements. In this study, we focused on rigid translation correction.

moco is a fast motion correction algorithm based on discrete Fourier transform and cache-aware upsampling (Dubbs et al., 2016). It minimizes L2 distance between the template image and the corrected image normalized by the area of the overlap from all possible pixel-by-pixel shifts. It was written in Java as an ImageJ plugin and reported to have close to real-time performance in post-hoc analysis.

Discrete Fourier transform-based registration corrects motion artifact up to pixel-by-pixel accuracy. When each ROI only contains a small number of pixels, subpixel registration can potentially improve accuracy of estimating the calcium signal. Registering upsampled image is an approach to achieve subpixel accuracy, but it increases computation for registration. An efficient method has been introduced to only calculate upsampled correlation coefficients around the optimal pixel-by-pixel shift (Guizar-Sicairos et al., 2008). This can be done without fully calculating the inverse discrete Fourier transform of the upsampled images, thus reducing the memory requirement and computation time. However, it has been reported that the overall registration accuracy was lower compared to moco or Turboreg when applied to images from two-photon calcium imaging (Dubbs et al., 2016).

Algorithms discussed above are intensity-based registration algorithms. Alternatively, feature-based registration can be used to correct motion artifact (Aghayee et al., 2017). This can be beneficial when the features are easily recognizable in each frame, but it may fail when signal to noise ratio is low. At the time of the research, the implementation was not available (https://github.com/saghayee/Tracking-based-registration, empty repository, accessed on 7/8/2018) and in this study we focused on intensity-based registration.

Recent studies began to use two-photon calcium imaging in real-time closed-loop experiments (Clancy et al., 2014; Hira et al., 2014; Mitani et al., 2018; Prsa et al., 2017). They require fast, real-time image processing, and lateral motion correction has not been implemented and ignored in most of these studies due to technical difficulties. However, imaging in awake animals necessarily includes motion artifacts, leading to many studies with post-hoc analysis utilizing Turboreg and other image registration algorithms. Mitani et al., (2018) was the first to our knowledge to report real-time processing of calcium imaging incorporating lateral motion correction. This method used hill-climbing method to reduce computation of correlation coefficient between template and shifted images. Here, we report the details and the performance analysis of the implementation of fast motion correction, including some improvements we have made since the original study. In addition to real-time image processing, the method can also be used for faster post-hoc processing.

After motion correction, typically ROIs are identified, and the relative change of the average fluorescence intensity of all the pixels in each ROI is calculated, based on the estimation of the baseline of the average fluorescence intensity. To estimate the baseline, percentile method and robust mean method are widely used, but each has shortcomings. We introduce a novel method to estimate the baseline calcium signal using kernel density estimate and discuss the benefits of the algorithm.

Material and Methods

Experimental methods

Animals

All procedures were in accordance with protocols approved by UCSD Institutional Animal Care and Use Committee and guidelines of the US National Institutes of Health. All animals were group housed in disposable plastic cages with standard bedding in a room on a reversed light cycle (12h/12h). Experiments were typically performed during the dark period. Cross between CaMK2a-tTA [JAX 003010] and tetO-GCaMP6s [JAX 024742] was used for cell body imaging. All the animals were C57bl/6 background.

Surgery

Surgical procedures were performed as previously described (Mitani et al., 2018). Adult mice (6 weeks or older, male and female) were anesthetized with isoflurane and injected with Baytril (10 mg/kg), dexamethasone (2 mg/kg) and buprenorphine (0.1 mg/kg) subcutaneously to prevent infection, inflammation and discomfort. A custom head-plate was glued and cemented to the skull. Craniotomy (~3 mm) was performed over the right caudal forelimb area (300 µm anterior and 1,500 µm lateral from the bregma). A mixture of AAV1.Syn.Flex.GCaMP6f (1:5000-10000 final dilution) and AAV1.CMV.PI.Cre (1:2 final dilution) diluted in saline was injected 20-30 nL at 3-5 sites (~250 µm depth, ~500 µm apart) for dendrite imaging. Experiments were performed at least 7 days after surgery.

Imaging

Imaging was conducted with a commercial two-photon microscope (Bscope, Thorlabs) using a 16x objective (NIKON) with excitation at 925 nm (Ti-Sa laser, Newport). Images were acquired with ScanImage 4 (Vidrio Technologies). Imaging was conducted in awake animals.

Computational methods

Motion correction (Hill-climbing method)

First, the template image was generated from 1000 image frames obtained before the experiments. Using the OpenCV template matching method (explained later), the first 500 images were aligned to the average of the last 500 images, the last 500 images were aligned to the average of the 500 corrected images, and the average of all the 1000 corrected images were used as a template.

Maximum absolute shift (m) in each direction was set to be 1/4 of the width (w) and height (h) of the image. From each edge of the template image, m pixels were cropped to take the central part ((w-2m)*(h-2m) pixels). To correct motion artifact of each image, the objective is to find where in the image best matches this central part of the template and maximizes the correlation coefficient.

Instead of the global maximum, a local maximum can be reached iteratively by hill climbing. Let (x, y) be the current position, where the shift maximizes the correlation coefficient among all the shifts tested up to that point. The correlation coefficients for shifts (x+1, y), (x-1, y), (x, y+1), and (x, y-1) are calculated, and if there is any shift that increases the correlation coefficient, the current position is updated by 1 pixel to maximize the correlation coefficient. This step is repeated until the current position reaches the local maximum. To assess the computational complexity, we use big O notation here to indicate how the running time requirements grow as the input size grows. When an algorithm takes O(n) time for an input of size n, it means that the computation time scales linearly or less with n. Assuming the path is somewhat straight, the above hill climbing method requires O(m)

steps to converge, and each step takes O(wh) time. Therefore, the computational complexity of the algorithm is O(mwh) under the assumption.

The complexity can be further reduced with a pyramid approach. Aligning images downscaled by 2 takes 1/8 of the time, and the corresponding shift in the original image should give a good estimate. It may not give the exact maximum, but the difference should be small. Using this shift as an initial shift constrains the number of expected steps until the convergence to the final target. With a deep enough image pyramid, the computational complexity approaches O(wh). However, it requires spatial features for alignment to be available in all the downscaled images, and practically too deep a pyramid makes the algorithm unstable (Dubbs et al., 2016). We typically use 2 layers of downscaled images, up to the factor of 4. (https://github.com/amitani/mex_tools/tree/master/mexBilinearRegistrator)

Motion correction (OpenCV template matching method)

The objective is the same as the hill-climbing method, but with matchTemplate (https://docs.opencv.org/3.2.0/de/da9/tutorial_template_matching.html) function of OpenCV (https://opencv.org/), it calculates correlation coefficient for every possible shift to reach the global maximum. The function uses discrete Fourier transform internally. To increase the speed of computation, the image can be downscaled first, and then the shift can be multiplied for the original resolution, even though this would reduce the resolution of motion correction unless subpixel registration is applied.
(https://github.com/amitani/mex_tools/tree/master/cvMotionCorrect)

Motion correction (subpixel registration)

In many calcium imaging experiments, ROIs are small (~10 pixels wide), and subpixel registration can improve the accuracy of calcium activity estimation. Subpixel registration has been done with optimization (Thévenaz et al., 1998) or upsampling (Guizar-Sicairos et al., 2008), but each has either speed or accuracy problem (Dubbs et al., 2016). Here, we used a parabola fit approach, which is faster and more suitable for real-time application. Subpixel registration was achieved by finding the peak of the correlation coefficient heatmap in subpixel accuracy using a parabola fit. In the hill-climbing method, it uses custom implementation in C++ using the 5 points in the heatmap (the peak and the adjacent points in 4 directions). In the OpenCV method, the same algorithm but implementation by William Arden (minMaxLocSubPix) was used (originally acquired from http://answers.opencv.org/question/29665/getting-subpixel-with-matchtemplate/, and the copy is in https://github.com/amitani/mex_tools/blob/master/cvMotionCorrect/CV_SubPix.cpp and https://github.com/amitani/mex_tools/blob/master/cvMotionCorrect/CV_SubPix.h).

Baseline estimation with kernel density estimate

Every 20 frames, the baseline of each selected ROI was updated as follows. The preceding 2000 frames were used to estimate the baseline. 2000 frames were first split into 100 bins of 20 consecutive frames, and the average fluorescence of each bin was calculated. Kernel density estimate of the 100 average values were calculated, and the baseline was estimated to be the value at the peak of the distribution. The length of the window can be adjusted depending on the signal to noise ratio, the activity duration, and how quickly and

how often the baseline fluctuates. The script to process calcium intensity at each frame and estimate baseline is available at (https://github.com/amitani/baseline_kde).

Configurations

Intel Core i7 4790, 16GB DDR3, Windows 10, ImageJ 1.52d (Java 1.8.0_66, 64 bit), Matlab R2014a, OpenCV 3.2.0, Qt Creator 4.4.1, Visual Studio 2013

Results and discussions

Motion correction

Especially in a real-time analysis application, time spent for motion correction should be kept short. While direct comparison to previously reported computation time is difficult, partly because the details of the configuration are often not reported, NoRMCorre was reported to correct a 512 × 512 × 2000 movie in 40 sec (rigid transformation) and 117 sec (non-rigid transformation) (Pnevmatikakis and Giovannucci, 2017). moco was reported to correct a 512 × 512 × 2000 movie in 90 sec, while TurboReg took 170 sec (fast) and 298 sec (slow) as reported in the same study (Dubbs et al., 2016). At 30 frames per second of image acquisition, 2000 frames take 66 seconds to capture. Only rigid transformation with NoRMCorre is slightly faster than this, but it still takes more than half of the acquisition time to process a frame. Considering the overhead of other steps, e.g. image transfer between processes and baseline estimation, it is necessary to make the processing time shorter for a real-time analysis.

We compared three implementations of motion correction based on correlation coefficient maximization. Our method is similar to moco, but instead of using the whole overlap, we took the central part of the template image, and tried to maximize the correlation coefficient with the corresponding part of the target image. We first tested a hill-climbing method to find a local maximum. To increase speed, a pyramid approach was used. In this approach, the initial shift for hill-climbing is determined by the alignment of a downscaled image. Theoretically, it becomes faster with a deeper image pyramid, but there was no significant speed increase when the shift was small (Figure 4.1, left column). This is because the default initial shift (no shift) is close enough to the final shift.

A caveat of the hill-climbing method is that it does not perform well when the true final shift is far from the initial shift. As the path becomes longer, it requires more computation, making it slower. Furthermore, if there is another local maximum along the path, the algorithm can converge to the local maximum, not the true final shift. This can be problematic in long experiments when a slow drift was not adjusted properly during the experiment. To simulate this situation where the images are far from the template, we artificially shifted each image frame by 32 pixels in both X and Y axis before motion correction (Figure 4.2). Without an image pyramid, the algorithm almost never converged to the true final shift; instead, it jumped among local maxima as indicated by sudden jumps in the corrected distance (Figure 4.2, top left). With an image pyramid with a few layers, it converged to the same final shift estimated without the additional shift (Figure 4.2, left column). However, it has been reported that an image pyramid method downsampling the image 3 and 4 times gave severe errors (Dubbs et al., 2016). While we did not observe such errors in our cell body imaging data, it implies that the number of layers of an image pyramid

has to be properly adjusted for each experimental condition. Parameter tuning often incorporates trials and errors, which is not suitable for real-time experiments.

To overcome this, we performed dense search to align the most downscaled image of an image pyramid, calculating correlation coefficients for all possible shifts to find the global maximum (Figure 4.2, middle column). With an image pyramid downscaling the image twice, this algorithm could converge to the same final shift as originally estimated without additional shift, different from the hill-climbing algorithm. However, it was implemented by naive exhaustive search, hindering the speed with shallow or no image pyramid.

Alternatively, we used matchTemplate function in OpenCV to search all possible shifts to reach the global maximum. When downscaling was used to reduce computation time, the estimated shift was transformed for the original image resolution with a parabola fit (no hill climbing was used). This was fast enough to apply to a 2× downscaled image, and more accurate than the hill-climbing method with similar speed.

We further examined the performance of the algorithms with sparsely-labelled dendrite imaging data (Figure 4.3). This data is more challenging for motion correction because the motion artifact is severer with a higher zoom, and the signal tends to be weaker in dendrites. The results show that the alignment becomes unstable when the image is downscaled too much (Figure 4.3, bottom row), and the hill-climbing method is unstable without an image pyramid or with a shallow image pyramid (Figure 4.3, left column). To simulate severer motion artifacts, we added random shift up to 8 pixels in each direction at each frame (Figure 4.4). The results further illustrate the speed and the stability of OpenCV template matching method in severer conditions. Interestingly, the estimate becomes noisier with OpenCV template matching method when applied to more downscaled images. This is

because the estimate is affected by how the pixels are split into patches for downscaling. Note that this effect is negligible when the image is only downscaled up to a factor of 2.

Balancing between accuracy and speed, we recommend applying the OpenCV method to a twice downscaled image. Using a regular personal computer, our best method processed a 512 × 512 × 1000 movie in less than 3 seconds. This is a significant improvement from the previous motion correction methods, while maintaining the accuracy.

Baseline estimation

A certain percentile of the fluorescence intensity time series has been commonly used to estimate the baseline. However, this method has multiple shortcomings. In calcium imaging, the calcium signal tends to have symmetric noise and sparse positive activity. Therefore, the percentile that represents the true baseline depends on the activity level. With no activity, the baseline should be the median of the distribution, while it corresponds to a lower percentile with more activity. Typically, a percentile lower than the median is used as estimated baseline. In this case, there is a jitter in timing caused by the changing trend of the baseline. When the baseline is increasing, the increase of the estimate happens later. On the other hand, when the baseline is decreasing, the decrease of the estimate happens earlier. For example, let us consider a situation estimating baseline from 100 values without noise and activity. In this case, the 20th percentile of the intensity values is at the 20th bin when the baseline is constantly increasing, and at the 80th bin when the baseline is constantly decreasing. This creates temporal difference equivalent to 60 bins. According to the increasing or decreasing trend of the baseline, there is a different degree of delay in the estimate of the baseline.

Another popular method for baseline estimation is robust mean. This method calculates mean of the signal while excluding outliers, which are mainly from calcium activity. Outliers are typically defined as values different from the mean more than a set threshold, e.g. 2 standard deviation. An assumption for this method is that the mean is close to the true baseline, which is not the case when the activity level is high and can lead to poor baseline estimation.

To overcome these issues of the common methods, we estimated the baseline using kernel density estimation. Kernel density estimation is a method to estimate kernel density from a limited number of samples under the assumption that the kernel density function is smooth. Kernel density is a probability distribution from which each sample is produced. After the estimation, the peak of the kernel density approximates the center of the baseline. It only assumes that the baseline distribution peaks at the center with symmetric noise, and the peak is higher than the kernel density at fluorescence values during calcium events. To increase the accuracy of the peak estimation, we first reduced noise by binning and averaging (20 frames), and subsequently kernel density estimation was performed using ksdensity function of Matlab (R2014a). It is done by convolving a Gaussian kernel of width optimized to samples. Taking the peak of kernel density of a continuous distribution is comparable to taking the mode of a discrete distribution. Therefore, we hypothesized that there is little bias from increased activity and increased noise. We compared the baseline estimation on two consecutive windows of 2000 frames with and without apparent calcium activity. Furthermore, to simulate increased noise, we compared downsampling and averaging of the bins (Figure 4.5). These results show that the estimates by the kernel density estimation

method were the most robust across different conditions, even with increased activity and increased noise.

Implementation of real-time image processing for closed-loop experiments

We developed a real-time image processing pipeline for two-photon calcium imaging for closed-loop experiments that includes a lateral motion correction with a comparable accuracy to popular post-hoc methods as well as an improved baseline estimation method. In the pipeline, each image is first copied to a memory-mapped file at the time of acquisition by a custom plugin for ScanImage 4 (https://github.com/amitani/mex_tools). A custom Qt GUI application reads each image from the memory-mapped file, corrects for motion artifact with the OpenCV template matching method, and saves the corrected image in another memory-mapped file (https://github.com/amitani/onlineMotionCorrection). Another instance of Matlab reads the corrected image from the memory-mapped file, calculate average pixel intensity of each ROI, estimate baseline (https://github.com/amitani/baseline_kde), and calculate ΔF/F. This information of relative change in fluorescence is further used in a closed-loop experiment.

Conclusions

Here, we discussed our implementation of fast motion correction and baseline estimation algorithms for real-time image processing for two-photon calcium imaging. This is to our knowledge the first reported software pipeline for real-time processing integrated with motion artifact correction for closed-loop experiments. Our implementation of motion

correction is significantly faster than previously reported software packages, while maintaining the accuracy. The pipeline can also be used for fast post-hoc analysis. Indeed many of the recent publications from our laboratory used this or previous versions of the method for post-hoc analysis (Chu et al., 2016, 2017; Hwang et al., 2017; Li et al., 2018; Mitani et al., 2018; Peters et al., 2017).

Funding

This work was supported by grants from NIH (R01 NS091010A, R01 EY025349, R01 DC014690, U01 NS094342, and P30EY022589), Pew Charitable Trusts, David & Lucile Packard Foundation, McKnight Foundation, New York Stem Cell Foundation, Kavli Institute for Brain & Mind, and NSF (1734940) to TK. AM is supported by the Nakajima foundation.

Acknowledgments

We thank K. O'Neil, L. Hall, T. Loveland and O. Arroyo for technical assistance, the GENIE Project at Janelia Research Campus for GCaMP6f, and members of the Komiyama lab for comments and discussions.

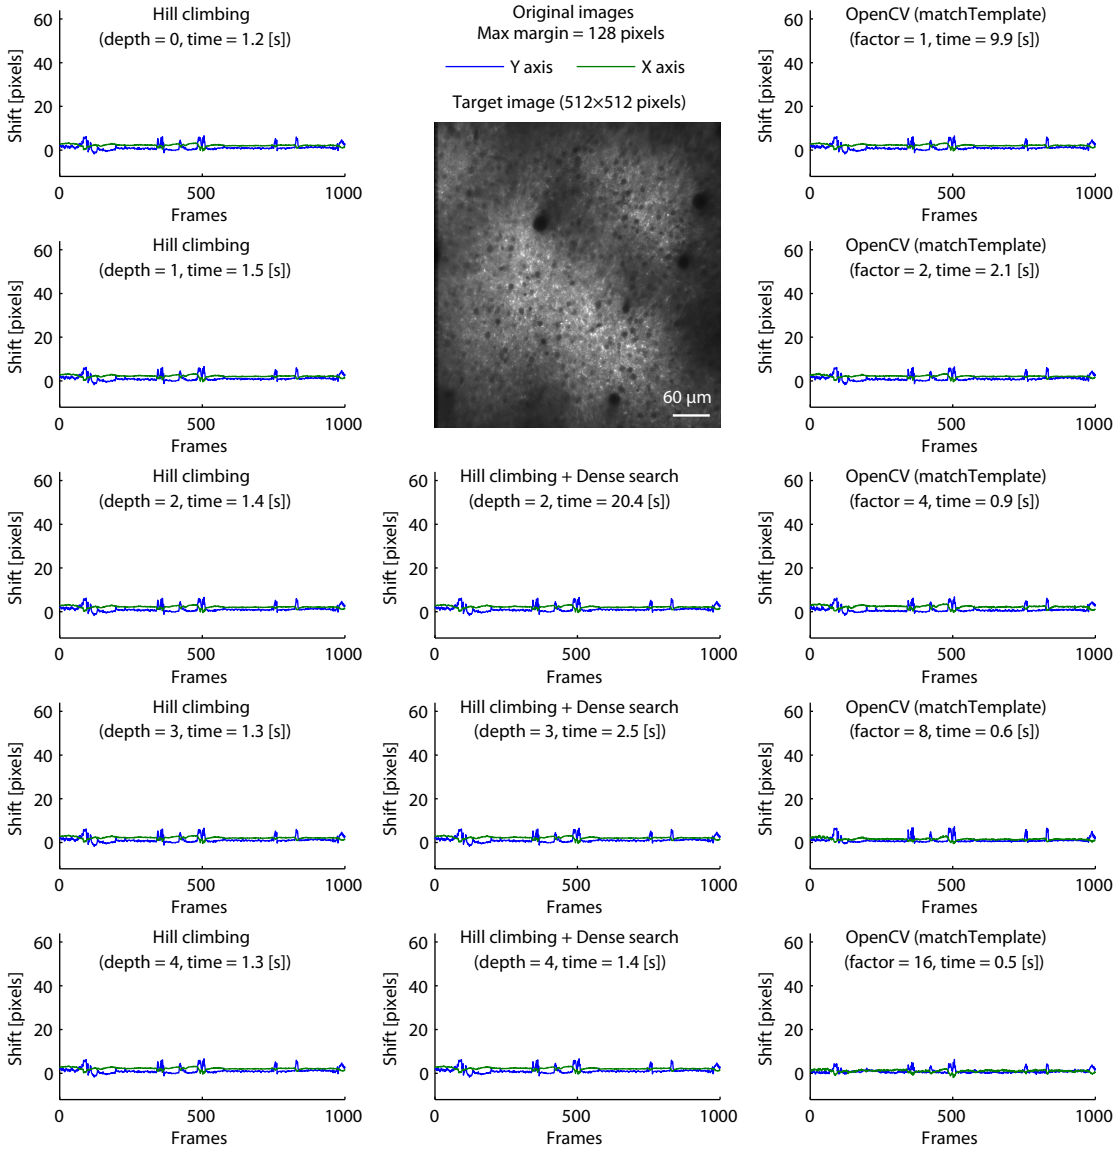

Figure 4.1. Comparison of speed and stability with different algorithms and parameters for neural ensemble imaging.
Estimated shifts (blue: Y axis, green: X axis) with different algorithms and parameters for cell body imaging. Dense search was slowest even at depth = 2 and shallower depth was not included.). Maximum shift for correction was set to be 128 pixels in all conditions.

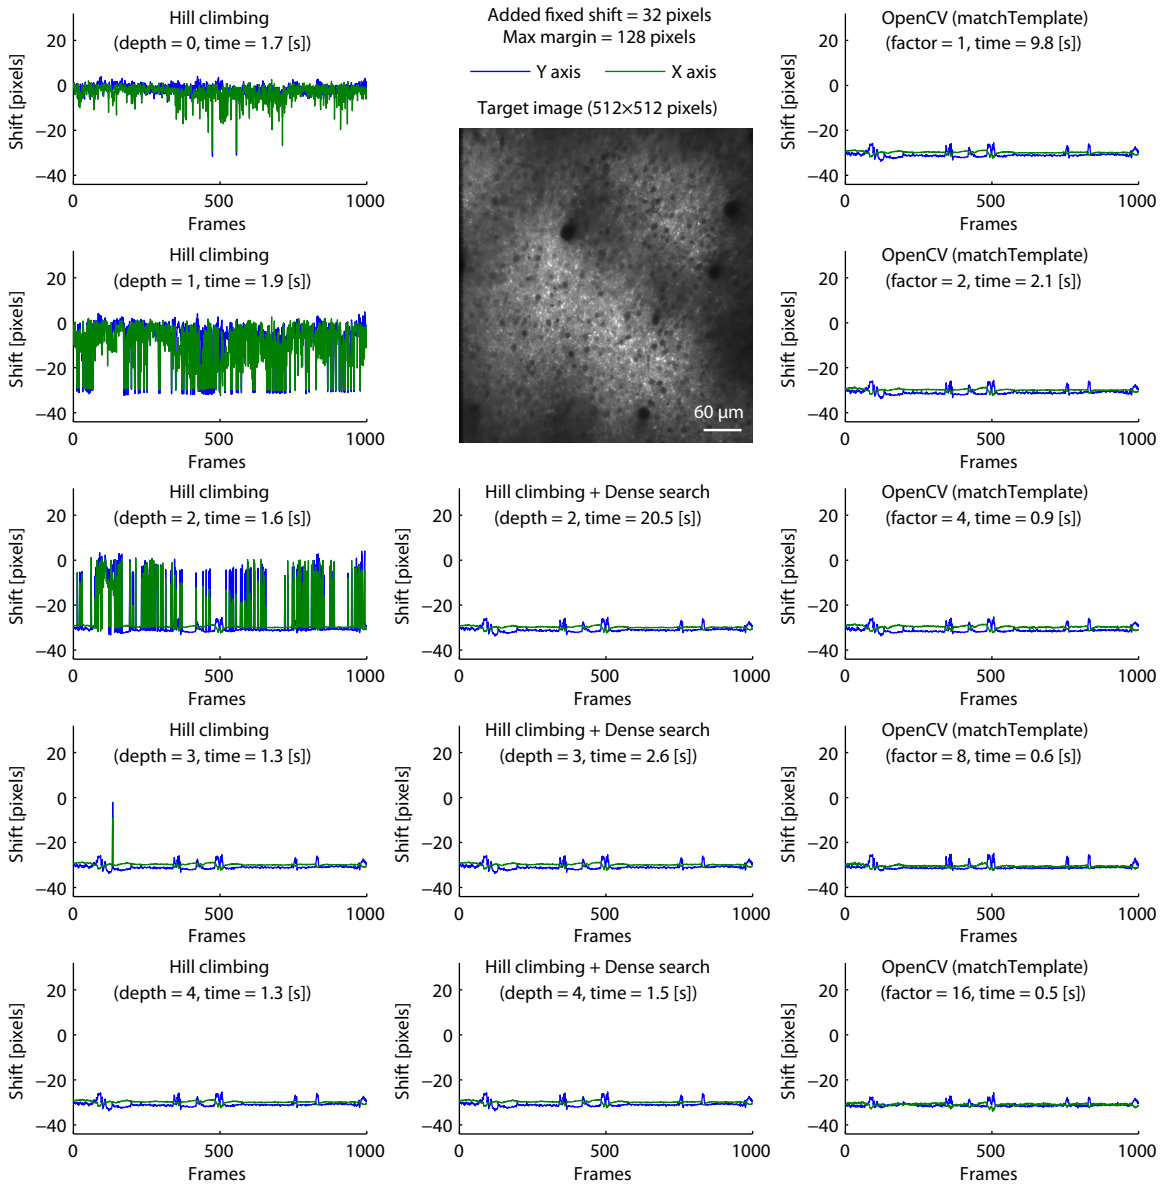

Figure 4.2. Comparison of speed and stability with different algorithms and parameters with artificially added fixed shifts for neural ensemble imaging.
Same as Fig. 1 with fixed shifts (32 pixels in both X and Y axis) added artificially before motion correction.

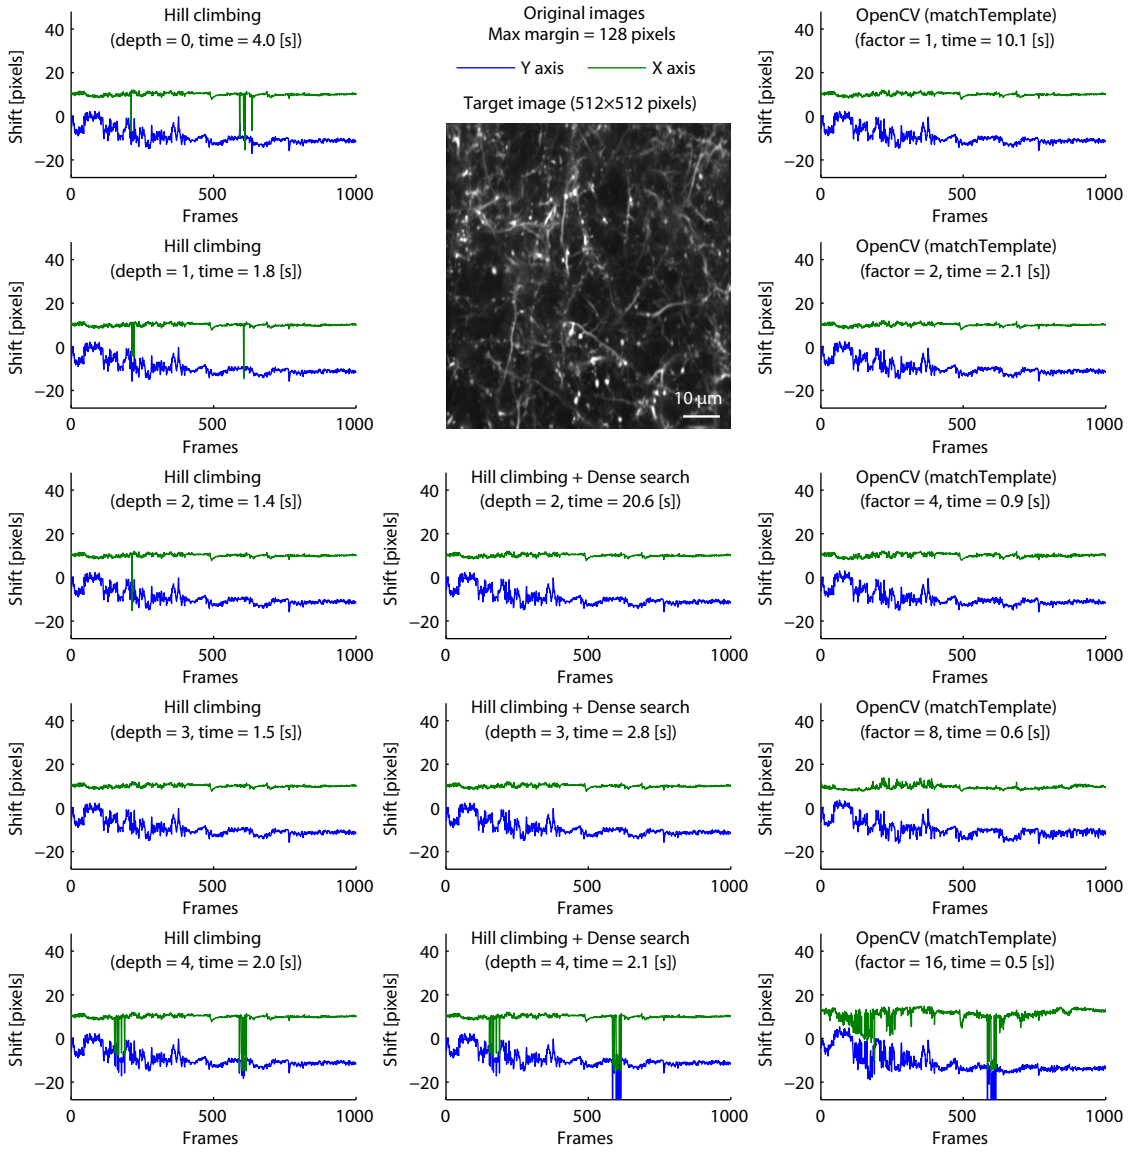

Figure 4.3. Comparison of speed and stability with different algorithms and parameters for dendrite imaging.
Same as Fig. 1 applied to dendrite imaging.

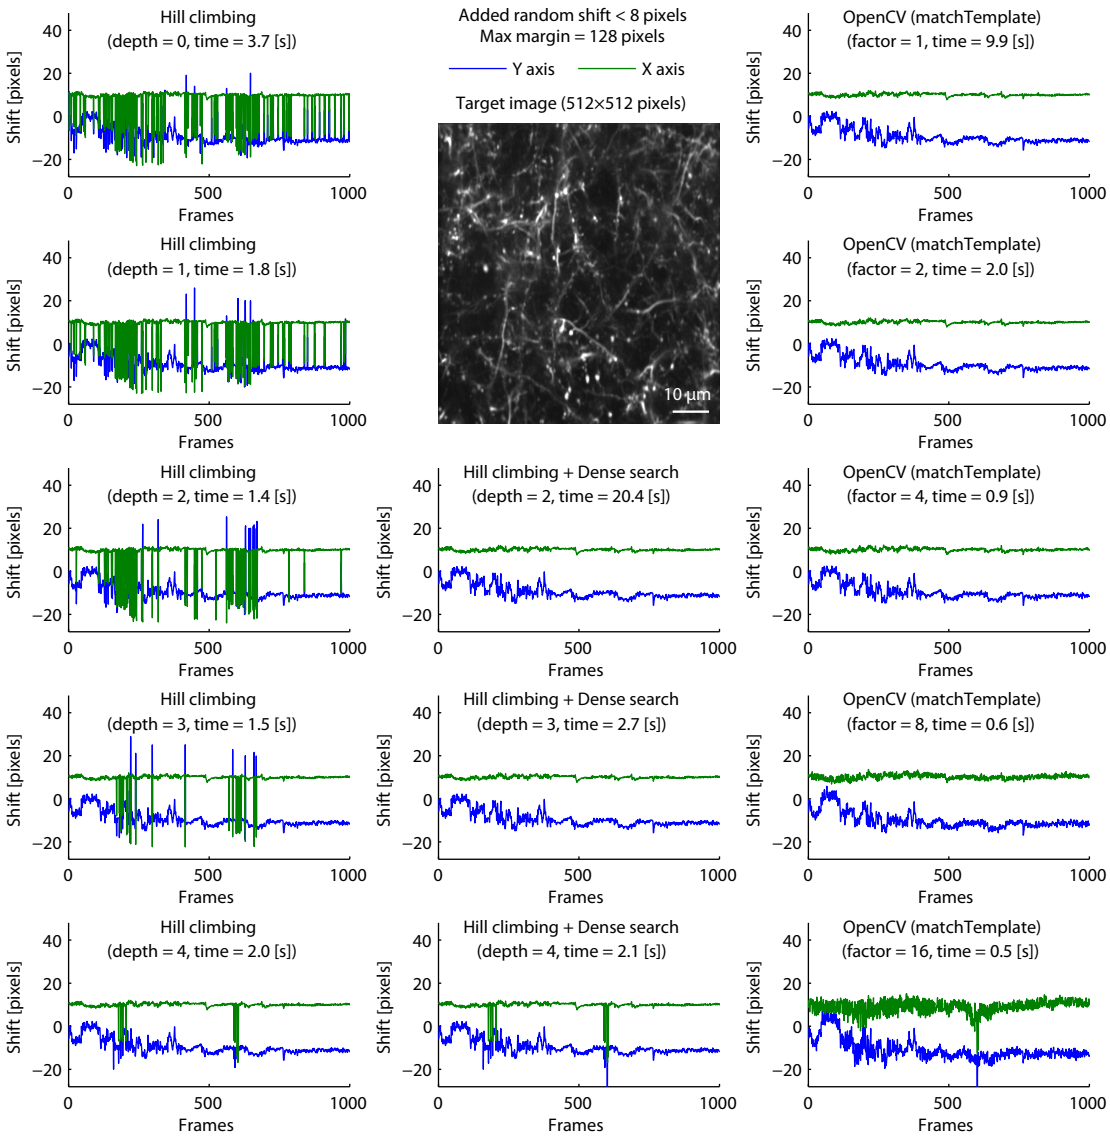

Figure 4.4. Comparison of speed and stability with different algorithms and parameters with artificially added random shifts for dendrite imaging.
Same as Fig. 1 applied to dendrite imaging with random shifts (up to 8 pixels in each direction) added artificially before motion correction. For visualization, these artificially added shifts were subtracted before the corrected distances are plotted.

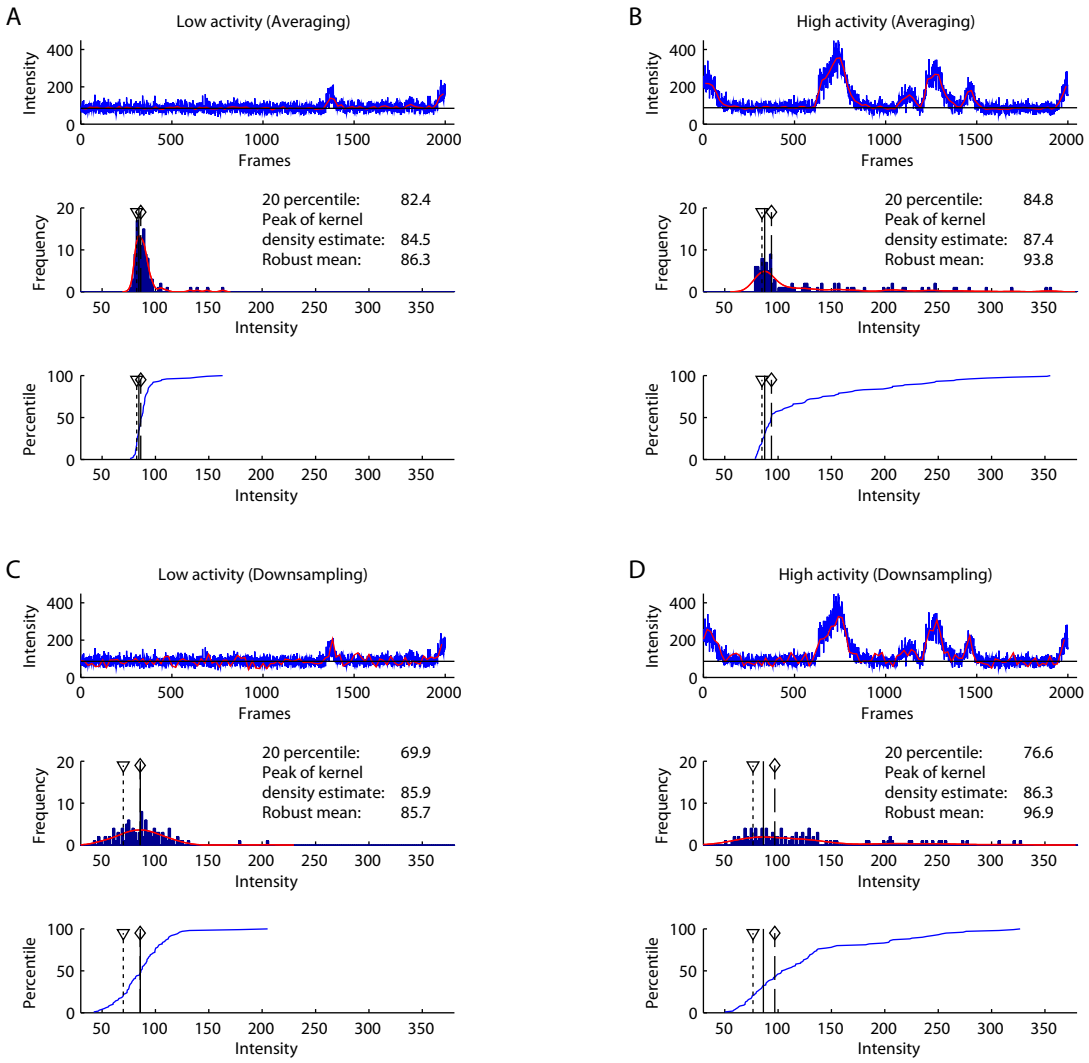

Figure 4.5. Comparison of percentile, robust mean, and kernel density estimate methods for baseline estimation.

(A) Analysis on a window of 2000 frames with little calcium activity. Blue line shows raw signal, and red line shows the signal after binning and averaging every 20 frames, which was used for baseline estimation (Top). Histogram of signal intensity. Red line indicates kernel density estimate. A triangle with dotted line indicates the 20th percentile, a solid line indicates the peak of kernel density estimate, a diamond with a dashed line indicates robust mean (excluding values not within 2 S.D. from the mean) (Middle). Cumulative density function of the intensity distribution (blue line) (Bottom). (B) Same as (A) on the next 2000 frames to (A). This window contains a high level of calcium activity. (C-D) Same as (A) and (B) with downsampling. Red line indicates the signal downsampled by 20 (excluding 19 frames from 20 frames), which is further used for baseline estimation (Top). Note that only the kernel density estimate method identified consistent baseline values in averaged and downsampled data.

Chapter 4 is material currently being prepared for submission for publication. Mitani A, and Komiyama T. Real-time image processing for two-photon calcium imaging. The dissertation author was the primary investigator and author of this material.

Chapter 5. General discussion

In this dissertation, we investigated BCI tasks using two-photon calcium imaging. Until recently, only neural activity of excitatory neurons has been targeted in BCI tasks. This is mainly due to the technical limitations of electrical recordings. However, BCI tasks using two-photon calcium imaging have been recently reported, opening the possibility of targeting genetically specified neural population, or substructure of the neurons. In Chapter 2, we targeted inhibitory neurons and showed that animals can learn to modulate each subtype of inhibitory neurons (PV-positive, SOM-positive, and VIP-positive inhibitory neurons) with a subtype-specific strategy. In Chapter 3, we expressed GCaMP6f in sparsely labelled neurons and targeted spine-specific activity, although we could not train animals to perform the task. In Chapter 4, we discussed the algorithms we examined to achieve real-time image processing for BCI tasks with two-photon calcium imaging.

To our knowledge, it was the first reported study examining the plasticity of inhibitory neurons for BCI task. In the task, two neurons (of the same subtype) in an imaging field was randomly selected to be N+ and N-. When the calcium activity of N+ exceeded that of N- by a set threshold, the animal was rewarded. When a pair of PV neurons are targeted, the activity of N- decreased leading to increased reward rate. When either SOM or VIP neurons are targeted, the activity of N+ increased. This shows that each subtype is modulated differently during BCI training, suggesting that each subtype plays a different role.

A remaining question is where this difference comes from. When activity of a neuron is modulated, there are many potential mechanisms achieving the modulation. For example, to make the target neuron more active, its excitability can increase. The excitatory synapses on to the neuron can be potentiated. Alternatively, the presynaptic excitatory neurons can be

more active, or more correlated with each other being more efficient in activating the target neuron. The presynaptic inhibitory neurons can be less active or be less correlated with excitatory neurons. These can be different among subtypes of inhibitory neurons, and between inhibitory neurons and excitatory neurons.

One of the intriguing hypotheses is based on the connections between inhibitory neuron subtypes. It has been shown that SOM inhibits PV, PV inhibits excitatory neurons, and excitatory neurons excites SOM neurons. Therefore, this can form a positive feedback loop, potentially underlying the increase of N+ activity of SOM neurons. Furthermore, if this increased activity of SOM neurons is inhibiting N- PV neurons when PV neurons are targeted, this loop can explain the results of both SOM neurons and PV neurons. To test this hypothesis, imaging a different type of neurons while targeting one subtype will be beneficial. This can be done by expressing calcium indicators of different colors, or by post-hoc histology to identify subtypes (Khan et al., 2018). Further research regarding the mechanism of subtype specific plasticity is warranted.

Understanding the mechanisms underlying BCI training is also important when excitatory neurons are targeted. It has been shown that corticostriatal plasticity depending on NMDA receptors in the striatum is necessary for BCI learning (Koralek et al., 2012). However, it does not mean that it is the sole mechanism for learning the abstract skills. It has been shown that dopaminergic projection to the cortex is crucial for motor learning, suggesting that reward-dependent plasticity in the cortex is playing a role (Hosp et al., 2011; Molina-Luna et al., 2009). What is the contribution of the plasticity in the cortex in BCI learning? To answer this question, we investigated a BCI task targeting spine-specific activity in the cortex. Excitatory synapses along the dendrite of cortical excitatory neurons form

dendritic spines, and calcium influx to the spines caused by vesicle release from the presynaptic terminal can be captured by two-photon calcium imaging. While we could not train animals to learn the task, it does not necessarily mean that animals can never learn to modulate spine-specific activity.

If animals can learn to modulate spine-specific activity, it may reveal the association between spatial distribution and plasticity of spines. For example, targeting a spine may enhance spine dynamics in the vicinity, or the correlation structure of the spine-specific activity of nearby spines may determine how well the animal can learn to modulate the spine-specific activity of the target spine. Even when these are the case, it is possible that functionally related neurons simply tend to make spines close to each other, and nearby spines receive inputs from neurons whose activity is regulated together. Perturbation experiments will be required to examine whether or not the effect is due to local interaction between nearby spines or how the presynaptic neurons are regulated. One possibility is NMDA receptor knockout, but it has some issues. It has been shown that autocrine BDNF plays an important role in synaptic plasticity (Harward et al., 2016; Hedrick et al., 2016). Sparse knockout of BDNF receptor will prevent the synapses on the neuron from getting potentiated, and it will distinguish whether the spatially correlated plasticity of spines in BCI learning is due to local interaction of the spines or not.

Two-photon calcium imaging enabled targeting a specific subtype and substructure of neurons. It will help further understand the underlying mechanisms of BCI learning and it will also help develop better BCI. For example, cell type of a recorded neuron can be estimated from its spike width/shape, and knowledge about the plasticity and the function of them can

be integrated into a decoder design. This can potentially lead to a BCI design easier to use and quicker to learn without extensive training.

First application of BCI would likely be therapeutic neuroprosthetic devices for paralyzed patients. In the future, it may be further applied to higher cognitive deficits e.g. dementia. For such applications, we need to decode more information, and having access to information encoded in different cell types will be useful. Advancing recording technology to increase the number of neurons and long-term stability is important, but how to decode information from the neural activity is also crucial. We need to understand how the memory is stored and implemented in the brain across different cell types. We have to advance and integrate our knowledge about recording techniques, decoding algorithms, and underlying biological mechanisms to expand BCI applications.

References

Aghayee, S., Winkowski, D.E., Bowen, Z., Marshall, E.E., Harrington, M.J., Kanold, P.O., and Losert, W. (2017). Particle Tracking Facilitates Real Time Capable Motion Correction in 2D or 3D Two-Photon Imaging of Neuronal Activity. Front. Neural Circuits *11*, 56.

Beierlein, M., Gibson, J.R., and Connors, B.W. (2003). Two Dynamically Distinct Inhibitory Networks in Layer 4 of the Neocortex. J Neurophysiol. *90*, 2987–3000.

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.

Chen, S.X., Kim, A.N., Peters, A.J., and Komiyama, T. (2015). Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. Nat. Neurosci. *18*, 1109–1115.

Chen, T.-W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295–300.

Chen, X., Leischner, U., Rochefort, N.L., Nelken, I., and Konnerth, A. (2011). Functional mapping of single spines in cortical neurons in vivo. Nature *475*, 501-505.

Chu, M.W., Li, W.L., and Komiyama, T. (2016). Balancing the Robustness and Efficiency of Odor Representations during Learning. Neuron *92*, 174–186.

Chu, M.W., Li, W.L., and Komiyama, T. (2017). Lack of Pattern Separation in Sensory Inputs to the Olfactory Bulb during Perceptual Learning. eNeuro *4*, ENEURO.0287-17.2017.

Clancy, K.B., Koralek, A.C., Costa, R.M., Feldman, D.E., and Carmena, J.M. (2014). Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. Nat. Neurosci. *17*, 807–809.

Collinger, J.L., Wodlinger, B., Downey, J.E., Wang, W., Tyler-Kabara, E.C., Weber, D.J., McMorland, A.J.C., Velliste, M., Boninger, M.L., and Schwartz, A.B. (2013). High-performance neuroprosthetic control by an individual with tetraplegia. Lancet *381*, 557–564.

Dombeck, D.A., Khabbaz, A.N., Collman, F., Adelman, T.L., and Tank, D.W. (2007). Imaging Large-Scale Neural Activity with Cellular Resolution in Awake, Mobile Mice. Neuron *56*, 43–57.

Donato, F., Rompani, S.B., and Caroni, P. (2013). Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning. Nature *504*, 272-

276.

Dubbs, A., Guevara, J., and Yuste, R. (2016). moco: Fast Motion Correction for Calcium Imaging. Front. Neuroinform. *10*, 6.

Evangelidis, G.D., and Psarakis, E.Z. (2008). Parametric Image Alignment Using Enhanced Correlation Coefficient Maximization. IEEE Trans. Pattern Anal. Mach. Intell. *30*, 1858–1865.

Falk, C.F., and Biesanz, J.C. (2016). Two Cross-Platform Programs for Inferences and Interval Estimation About Indirect Effects in Mediational Models. SAGE Open *6*, 2158244015625445.

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

Fino, E., and Yuste, R. (2011). Dense inhibitory connectivity in neocortex. Neuron *69*, 1188–1203.

Fino, E., Packer, A.M., and Yuste, R. (2013). The Logic of Inhibitory Connectivity in the Neocortex. Neurosci. *19*, 228–237.

Fu, M., Yu, X., Lu, J., and Zuo, Y. (2012). Repetitive motor learning induces coordinated formation of clustered dendritic spines in vivo. Nature *483*, 92–95.

Ganguly, K., and Carmena, J.M. (2009). Emergence of a stable cortical map for neuroprosthetic control. PLoS Biol. *7*, e1000153.

Ganguly, K., Dimitrov, D.F., Wallis, J.D., and Carmena, J.M. (2011). Reversible large-scale modification of cortical networks during neuroprosthetic control. Nat. Neurosci. *14*, 662–667.

Garcia-Junco-Clemente, P., Ikrar, T., Tring, E., Xu, X., Ringach, D.L., and Trachtenberg, J.T. (2017). An inhibitory pull–push circuit in frontal cortex. Nat. Neurosci. *20*, 389–392.

Gray, C.M., Maldonado, P.E., Wilson, M., and McNaughton, B. (1995). Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. J. Neurosci. Methods *63*, 43–54.

Greenberg, D.S., and Kerr, J.N.D.D. (2009). Automated correction of fast motion artifacts for two-photon imaging of awake animals. J. Neurosci. Methods *176*, 1–15.

Guizar-Sicairos, M., Thurman, S.T., and Fienup, J.R. (2008). Efficient subpixel image registration algorithms. Opt. Lett. *33*, 156–158.

Hangya, B., Pi, H.-J., Kvitsiani, D., Ranade, S.P., and Kepecs, A. (2014). From circuit motifs to computations: mapping the behavioral repertoire of cortical interneurons. Curr. Opin. Neurobiol. *26*, 117–124.

Harward, S.C., Hedrick, N.G., Hall, C.E., Parra-Bueno, P., Milner, T.A., Pan, E., Laviv, T., Hempstead, B.L., Yasuda, R., and McNamara, J.O. (2016). Autocrine BDNF–TrkB signalling within a single dendritic spine. Nature *538*, 99–103.

Hedrick, N.G., Harward, S.C., Hall, C.E., Murakoshi, H., McNamara, J.O., and Yasuda, R. (2016). Rho GTPase complementation underlies BDNF-dependent homo- and heterosynaptic plasticity. Nature *538*, 104–108.

Hessler, N.A., Shirke, A.M., and Malinow, R. (1993). The probability of transmitter release at a mammalian central synapse. Nature *366*, 569–572.

Hippenmeyer, S., Vrieseling, E., Sigrist, M., Portmann, T., Laengle, C., Ladle, D.R., and Arber, S. (2005). A developmental switch in the response of DRG neurons to ETS transcription factor signaling. PLoS Biol. *3*, e159.

Hira, R., Ohkubo, F., Masamizu, Y., Ohkura, M., Nakai, J., Okada, T., and Matsuzaki, M. (2014). Reward-timing-dependent bidirectional modulation of cortical microcircuits during optical single-neuron operant conditioning. Nat. Commun. *5*, 5551.

Hochberg, L.R., Bacher, D., Jarosiewicz, B., Masse, N.Y., Simeral, J.D., Vogel, J., Haddadin, S., Liu, J., Cash, S.S., van der Smagt, P., et al. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. Nature *485*, 372–375.

Hofer, S.B., Ko, H., Pichler, B., Vogelstein, J., Ros, H., Zeng, H., Lein, E., Lesica, N.A., and Mrsic-Flogel, T.D. (2011). Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. Nat. Neurosci. *14*, 1045–1052.

Hosp, J. a, Pekanovic, A., Rioult-Pedotti, M.S., and Luft, A.R. (2011). Dopaminergic projections from midbrain to primary motor cortex mediate motor skill learning. J. Neurosci. *31*, 2481–2487.

Huang, S., Huganir, R.L., and Kirkwood, A. (2013). Adrenergic Gating of Hebbian Spike-Timing-Dependent Plasticity in Cortical Interneurons. J. Neurosci. *33*, 13171–13178.

Hubel, D.H. (1957). Tungsten Microelectrode for Recording from Single Units. Science *125*, 549–550.

Hwang, E.J., Bailey, P.M., and Andersen, R.A. (2013). Volitional control of neural activity relies on the natural motor repertoire. Curr. Biol. *23*, 353–361.

Hwang, E.J., Dahlen, J.E., Mukundan, M., and Komiyama, T. (2017). History-based action selection bias in posterior parietal cortex. Nat. Commun. *8*, 1242.

Iacaruso, M.F., Gasler, I.T., and Hofer, S.B. (2017). Synaptic organization of visual space in primary visual cortex. Nature *547*, 449-452.

Jarosiewicz, B., Chase, S.M., Fraser, G.W., Velliste, M., Kass, R.E., and Schwartz, A.B.

(2008). Functional network reorganization during learning in a brain-computer interface paradigm. Proc. Natl. Acad. Sci. U. S. A. *105*, 19486–19491.

Jazayeri, M., and Afraz, A. (2017). Navigating the Neural Space in Search of the Neural Code. Neuron *93*, 1003-1014.

Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydın, Ç., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. Nature *551*, 232–236.

Karnani, M.M., Jackson, J., Ayzenshtat, I., Tucciarone, J., Manoocheri, K., Snider, W.G., and Yuste, R. (2016a). Cooperative Subnetworks of Molecularly Similar Interneurons in Mouse Neocortex. Neuron *90*, 86–100.

Karnani, M.M., Jackson, J., Ayzenshtat, I., Hamzehei Sichani, A., Manoocheri, K., Kim, S., and Yuste, R. (2016b). Opening Holes in the Blanket of Inhibition: Localized Lateral Disinhibition by VIP Interneurons. J. Neurosci. *36*, 3471–3480.

Kawaguchi, Y., and Kondo, S. (2002). Parvalbumin, somatostatin and cholecystokinin as chemical markers for specific GABAergic interneuron types in the rat frontal cortex. J. Neurocytol. *31*, 277–287.

Keck, T., Scheuss, V., Jacobsen, R.I., Wierenga, C.J., Eysel, U.T., Bonhoeffer, T., and Hübener, M. (2011). Loss of sensory input causes rapid structural changes of inhibitory neurons in adult mouse visual cortex. Neuron *71*, 869–882.

Khan, A.G., Poort, J., Chadwick, A., Blot, A., Sahani, M., Mrsic-Flogel, T.D., and Hofer, S.B. (2018). Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic interneuron classes in visual cortex. Nat. Neurosci. *21*, 851–859.

Koralek, A.C., Jin, X., Long II, J.D., Costa, R.M., and Carmena, J.M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. Nature *483*, 331–335.

Kuhlman, S.J., Olivas, N.D., Tring, E., Ikrar, T., Xu, X., and Trachtenberg, J.T. (2013). A disinhibitory microcircuit initiates critical-period plasticity in the visual cortex. Nature *501*, 543–546.

Kullmann, D.M., and Lamsa, K.P. (2011). LTP and LTD in cortical GABAergic interneurons: Emerging rules and roles. Neuropharmacology *60*, 712–719.

Kullmann, D.M., Moreau, A.W., Bakiri, Y., and Nicholson, E. (2012). Plasticity of Inhibition. Neuron *75*, 951–962.

Lamsa, K.P., Kullmann, D.M., and Woodin, M.A. (2010). Spike-timing dependent plasticity in inhibitory circuits. Front. Synaptic Neurosci. *2*, 8.

Li, W.L., Chu, M.W., Wu, A., Suzuki, Y., Imayoshi, I., and Komiyama, T. (2018). Adult-born neurons facilitate olfactory bulb pattern separation during task engagement. eLife *7*, e33006.

London, M., and Häusser, M. (2005). Dendritic Computation. Annu. Rev. Neurosci. *28*, 503–532.

Lu, J., Li, C., Zhao, J.-P., Poo, M., and Zhang, X. (2007). Spike-Timing-Dependent Plasticity of Neocortical Excitatory Synapses on Inhibitory Interneurons Depends on Target Cell Type. J. Neurosci. *27*, 9711–9720.

Madisen, L., Garner, A.R., Shimaoka, D., Chuong, A.S., Klapoetke, N.C., Li, L., van der Bourg, A., Niino, Y., Egolf, L., Monetti, C., et al. (2015). Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance. Neuron *85*, 942–958.

Maynard, E.M., Nordhausen, C.T., and Normann, R.A. (1997). The Utah Intracortical Electrode Array: A recording structure for potential brain-computer interfaces. Electroencephalogr. Clin. Neurophysiol. *102*, 228–239.

Mitani, A., Dong, M., and Komiyama, T. (2018). Brain-Computer Interface with Inhibitory Neurons Reveals Subtype-Specific Strategies. Curr. Biol. *28*, 77–83.e4.

Molina-Luna, K., Pekanovic, A., Röhrich, S., Hertler, B., Schubring-Giese, M., Rioult-Pedotti, M.-S., and Luft, A.R. (2009). Dopamine in motor cortex is necessary for skill learning and synaptic plasticity. PLoS One *4*, e7082.

Moritz, C.T., Perlmutter, S.I., and Fetz, E.E. (2008). Direct control of paralysed muscles by cortical neurons. Nature *456*, 639–642.

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.

Otsuka, T., and Kawaguchi, Y. (2009). Cortical Inhibitory Cell Types Differentially Form Intralaminar and Interlaminar Subnetworks with Excitatory Neurons. J. Neurosci. *29*, 10533–10540.

Packer, A.M., and Yuste, R. (2011). Dense, Unspecific Connectivity of Neocortical Parvalbumin-Positive Interneurons: A Canonical Microcircuit for Inhibition? J. Neurosci. *31*, 13260-13271.

Peters, A.J., Chen, S.X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. Nature *510*, 263–267.

Peters, A.J., Lee, J., Hedrick, N.G., O'Neil, K., and Komiyama, T. (2017). Reorganization of corticospinal output during motor learning. Nat. Neurosci. *20*, 1133–1141.

Pfeffer, C.K., Xue, M., He, M., Huang, Z.J., and Scanziani, M. (2013). Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. Nat. Neurosci. *16*, 1068–1076.

Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J.I., Huang, Z.J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. Nature *503*, 521–524.

Pnevmatikakis, E.A., and Giovannucci, A. (2017). NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. J. Neurosci. Methods *291*, 83–94.

Prsa, M., Galiñanes, G.L., and Huber, D. (2017). Rapid Integration of Artificial Sensory Feedback during Operant Conditioning of Motor Cortex Neurons. Neuron *93*, 929–939.e6.

Rudy, B., Fishell, G., Lee, S., and Hjerling-Leffler, J. (2011). Three groups of interneurons account for nearly 100% of neocortical GABAergic neurons. Dev. Neurobiol. *71*, 45–61.

Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista, A.P. (2014). Neural constraints on learning. Nature *512*, 423–426.

Sarihi, A., Jiang, B., Komaki, A., Sohya, K., Yanagawa, Y., and Tsumoto, T. (2008). Metabotropic Glutamate Receptor Type 5-Dependent Long-Term Potentiation of Excitatory Synapses on Fast-Spiking GABAergic Neurons in Mouse Visual Cortex. J. Neurosci. *28*, 1224–1235.

Shenoy, K. V., and Carmena, J.M. (2014). Combining decoder design and neural adaptation in brain-machine interfaces. Neuron *84*, 665–680.

Smith, S.L., Smith, I.T., Branco, T., and Häusser, M. (2013). Dendritic spikes enhance stimulus selectivity in cortical neurons in vivo. Nature *503*, 115–120.

Steinmetz, N.A., Koch, C., Harris, K.D., and Carandini, M. (2018). Challenges and opportunities for large-scale electrophysiology with Neuropixels probes. Curr. Opin. Neurobiol. *50*, 92–100.

Taniguchi, H., He, M., Wu, P., Kim, S., Paik, R., Sugino, K., Kvitsiani, D., Kvitsani, D., Fu, Y., Lu, J., et al. (2011). A resource of Cre driver lines for genetic targeting of GABAergic neurons in cerebral cortex. Neuron *71*, 995–1013.

Thévenaz, P., Ruttimann, U.E., and Unser, M. (1998). A pyramid approach to subpixel registration based on intensity. IEEE Trans. Image Process. *7*, 27–41.

Velliste, M., Perel, S., Spalding, M.C., Whitford, A.S., and Schwartz, A.B. (2008). Cortical control of a prosthetic arm for self-feeding. Nature *453*, 1098–1101.

Wilson, D.E., Whitney, D.E., Scholl, B., and Fitzpatrick, D. (2016). Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. Nat. Neurosci. *19*, 1003–1009.

Wilson, D.E., Smith, G.B., Jacob, A.L., Walker, T., Dimidschstein, J., Fishell, G., and Fitzpatrick, D. (2017). GABAergic Neurons in Ferret Visual Cortex Participate in Functionally Specific Networks. Neuron *93*, 1058-1065.e4.

Xu, T., Yu, X., Perlik, A.J., Tobin, W.F., Zweig, J.A., Tennant, K., Jones, T., and Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. Nature *462*, 915–919.

Yavorska, I., and Wehr, M. (2016). Somatostatin-Expressing Inhibitory Interneurons in Cortical Circuits. Front. Neural Circuits *10*, 76.

Yoshimura, Y., and Callaway, E.M. (2005). Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. Nat. Neurosci. *8*, 1552–1559.