Understanding the Role of M1 in Locomotion: Population Dynamics and Applications to Brain-Machine Interfaces

David Xing

B.S. Biomedical Engineering - 2014

Columbia University

This dissertation by David Y. Xing is accepted in its present form

by the School of Engineering as satisfying the

dissertation requirement for the degree of Doctor of Philosophy.

Date Jan 7, 2021 _____

_____

David Borton, Ph.D., Advisor

Recommended to the Graduate Council

Date Jan 8, 2021 _____

_____

Wilson Truccolo, Ph.D., Reader

Date Jan 8, 2021 _____

_____

Leigh Hochberg, Ph.D., Reader

Date Jan 8, 2021 _____

_____

David Sheinberg, Ph.D., Reader

Date Jan 8, 2021 _____

_____

Lee Miller, Ph.D., Reader

Approved by the Graduate Council

Date_____

_____

Andrew G. Campbell, Ph.D.

Dean of the Graduate School

David Xing

david_xing@brown.edu

160 Brown Street
Providence, RI 02906

Phone Number:
862-242-9216

EDUCATION

2010-2014.   B.S., Biomedical Engineering Major, Electrical Engineering Minor. Columbia University, Cum Laude.

2014-Present.   Pursuing a Ph. D.in Biomedical Engineering. Brown University.

WORK EXPERIENCE

Winter of 2012, 2013, & 2014.   Student aid at Columbia's Undergraduate Admissions office. Sorted student applications.

RESEARCH EXPERIENCE

Summer of 2010.   Columbia University. Research assistant for Professor Peter deMenocal at Lamont Campus, working on a project to track historic Atlantic ITCZ shifts. Sampled sea cores, isolated foraminifera, and ran analysis on the samples for data collection.

2012-2014.   Columbia University. Undergraduate research assistant in the Laboratory for Intelligent Imaging and Neural Computing headed by Professor Paul Sajda. Worked on projects modeling human visual perception. Participated in data collection, and wrote programs to perform data analysis and signal processing in MATLAB and FSL. Performed experiments in EEG and fMRI.

Summer of 2013.   Food and Drug Administration. Undergraduate research intern at the Neural Implants and Prostheses Laboratory headed by Dr. Cristin Welle. Devised, planned, and carried out a research project determining the neural-implant performance decay using a mouse model.

2013-2014.   Columbia University. Developed and prototyped an implant for autonomously detecting periprosthetic joint infections after total knee replacement surgery.

2014-Present.   Brown University. Developing a spinal-cortical prosthesis system to restore top-down control of locomotion after spinal cord injury. Investigating gait representation and encoding in motor cortex in the Brown Neuroengineering and Neuromotion Laboratory headed by Professor David Borton.

AWARDS AND SCHOLARSHIPS

2010.	Community Builders & Remodelers Association of New Jersey scholarship.
2010-2014.	Dean's List at Columbia University
2012.	Columbia WEP Grant to fund internship at Laboratory for Intelligent Imaging and Neural Computing. Received April, 2012. Beth Vanderputten (bv2145@columbia.edu, 212-854-9388).
2013.	ORISE Fellowship to fund research at the FDA. Received June, 2013. Melissa Goodman (Melissa.Goodman@orau.org, 865-241-1294).
2014.	ASAIO Student Design Competition, Finalist.
2014.	Collegiate Inventor's Competition, 3rd place, $7500.
2014.	Brown Institute for Brain Science NeuroPracticum award. Received September, 2014. Christopher Moore (christopher_moore@brown.edu, 401-863-1054).
2017.	Brain Science Graduate Research Award Fellowship. Received July, 2017. Ines Tomas Pereira (ines_tomas_pereira@brown.edu, 401-863-5926).
2020.	Society for the Neural Control of Movement Travel Scholarship. Received January 2020. Michelle Smith (michelle@podiumconferences.com).

PRESENTATIONS / POSTERS

Xing D, Civillico E, Huang S, Krauthamer V, Welle C. *Characterization of in Vivo Mice Electrode Implant Performance as Correlates of Frontal Limb Behavior*. Poster session presented at: 2nd annual Columbia undergraduate Research Symposium; 2013 Oct 3; New York, NY.

Huang S, Civillico E, **Xing D**, Park S, Knaack G, Krauthamer V, Welle C. *A Video-Based Behavioral Test Platform to Evaluate the Long-term Safety and Effectiveness of Neural Interface Technology*. Poster session presented at: IEEE EMBS Conference on Neural Engineering; 2013 Nov 6-8; San Diego, California.

Xing D, Aghagolzadeh M, Brandman D, Vargas-Irwin C, Truccolo W, Borton D. *Low Dimensional Dynamics of the Primary Motor Cortex during Natural Locomotion Captures Kinematic Information and Improves Decoding Performance for Brain Machine Interfaces*. Poster session presented at: 2015 Neuroscience. 45th Annual Meeting of the Society for Neuroscience; 2015 Oct 17-21; Chicago, IL.

Xing D. *Investigating the role of primary motor cortex in directed and un-directed locomotion in a rhesus macaque model*. Invited speaker at SNEAALAS Spring 2019 General Membership Meeting; 2019 March 27; Providence, RI.

Xing D, Truccolo W, Borton D. *Population dynamics in primary motor cortex during locomotion and obstacle avoidance*. Poster session presented at: 29th Annual Neural Control of Movement Meeting; 2019 April 24-27; Toyama, Japan.

Xing D. *Primary motor cortex differentially modulates movement during locomotion and voluntary actions*. Invited speaker at Neural Control of Movement Award Winners Symposium; 2020 June 17; Online.

PUBLICATIONS

Capogrosso M, Milekovic T, Borton DA, Wagner F, Moraud EM, Mignardot J, Buse N, Gandar J, Barraud Q, **Xing D**, Jianzhong Y, Ko D, Qin L, Detemple P, Denison T, Micera S, Bezard E, Bloch J, Courtine G. A brain spinal interface alleviating locomotor deficits after spinal cord injury. *Nature* 539(7628):284-288, 2016.

Xing D, Aghagolzadeh M, Truccolo W, Borton D. Low-dimensional motor cortex dynamics preserve kinematics information during unconstrained locomotion in nonhuman primates. *Front Neurosci.* 13:1046, 2019.

Xing D, Truccolo W, Borton D. The role of motor cortex in locomotion and obstacle-avoidance. *In preparation*.

SKILLS

Programming – proficient: Matlab, C#, some experience: C++, Python
Design software – Autocad, Fusion 360, Eagle, Adobe Illustrator
Animal models – Rhesus macaques (behavioral training, husbandry, surgery, electrophysiology)

Acknowledgements

This journey through graduate school has been long and arduous, but also highly rewarding. I'm grateful for the amount of growth I was able to experience as a scientist and academic, and for all the friends that have helped me along the way.

Firstly, I would like to thank my mentor and advisor, Dr. David Borton. He has always given me the freedom to explore and pursue my interests, but was always there to offer his help when I needed it. Because of his guidance, I will leave his lab a far more competent, wiser, and resilient researcher than when I first joined it. I will always appreciate his patience, candor, and genuine interest in my personal and professional development. I am also deeply indebted to Dr. Wilson Truccolo, who has given me a new appreciation for computational neuroscience. I will miss our long and thoughtful discussions and having a safety net to catch all my mathematical mistakes. I would also like to thank my other committee members for their valuable feedback, and all the other wonderful Brown faculty I have interacted with over the years. In particular, the animal care facility veterinarians, technicians, and staff have been essential in helping me carry out my animal experiments and surgeries. A special thanks to Veronica Bouvier for helping me prep for all those animal surgeries (and teaching me how to properly fold surgical drapes!), Roxanne Burrill for sterilizing so many equipment packs for me, Dr. Jessica Johnston, Dr. Lara Helwig, and Dr. Moira Sheehan for their expertise and guidance on general animal work, and all them for staying late into the night with me to help monitor animals after particularly long surgeries.

Of course, I wouldn't have anything to put in this thesis without the help of Aaron Gregoire whose assistance with animal training and surgery was indispensable and John Murphy who is the only reason I have an experimental rig to run. I would also like to thank all the members of the Borton lab for making my experience in graduate school a much more enjoyable one. In particular Elaina Atherton and Sophie Brown for assisting me with with histology and confocal imaging and Dmitrijs Celinskis for teaching me how to cryosection tissue. Finally, I would like to thank my parents who have been supportive throughout all these years and my sister whose artistic talents I can always rely on to provide me with quality figures.

Contents

List of Tables

List of Figures

Chapter 1

Background

At the risk of sounding melodramatic, one could posit that the ultimate purpose of the nervous system is to generate movement. Of course, perception, cognition, and other neural processes are all crucial components of the nervous system. Ultimately, however, these neural processes are all tools subserving the only instrument through which an individual has to exert their will on the world - their actions. One can endow a rock, rolling towards the pit of an active volcano, with consciousness, the capacity for fear, temperature receptors, and even omniscient knowledge of the state of the universe, but without the ability to alter its trajectory, it is forever doomed to its fate. It should come as no surprise, then, that the goal of this thesis is to investigate the neural control and representation of movement, and in particular, one specific type of motion.

Movements can be categorized into many different types of motor actions, reaching, kicking, nodding, etc., but there are few that are as important, and as ubiquitous across species, as locomotion. Generally speaking, locomotion is the movement that transports us from one location to another. This is usually accomplished with cyclic movements, be it the undulations of the *c. elegans* nematode or the intricate quadrupedal gait pattern of a horse canter. One could argue that locomotion precedes

the nervous system itself, as even individual cells, such as sperm cells, are capable of generating propulsion through cyclic movements. Of course, higher order organisms are capable of more than simply moving from place to place. They are able to integrate other complex movements, such as dexterous leg and foot manipulations, with locomotion to traverse uneven and erratic terrain. This particular facet of locomotion, the question of how the integration of locomotion with other movements is accomplished by the nervous system is an area of active research and will be one of the main components of this work. In addition to advancing our understanding of motor neuroscience, investigating how the central nervous system generates the necessary motor commands to actuate locomotor movements will be useful for the development of novel neurotechnologies aimed at recovering lost motor functionality. In this chapter, I will briefly review the state of our current understanding of the motor control of mammalian locomotion as well as the advancements that have been made in recent years in brain-machine interfaces for restoring locomotion.

1.1 Neural control of locomotion

1.1.1 Spinal circuitry

From the discovery of spinal central pattern generators (CPGs) by Graham Brown in the 1910's, our understanding of the neuroscience behind locomotion has been steadily progressing throughout the last century. Before Brown, it was unclear whether locomotion was the result of sequentially activated reflex pathways within the spinal cord or whether the cyclic activity was intrinsic to the spinal cord, without the need for afferent inputs. The reflex hypothesis states that when the limb is moved, cutaneous and proprioceptive signals activate interneurons in the spinal cord which in turn activate motor neurons through the standard reflex pathway. This produces a movement and new sensory information, starting the cycle anew(Sherrington, 1906).

Graham Brown utilized a spinal preparation, isolating the spinal cord from the rest of the CNS, and removed any afferent inputs by sectioning the dorsal roots. The absence of any inputs driving the reflexes should result in an absence of cyclic activity. However, he was still able to record alternating cyclic activity in the muscles, indicating that sensory information is not necessary for the generation of phasic activity, and that these oscillators are intrinsic, or central, to the spinal cord (G raham Brown, 1911). However, this is not to say that sensory inputs are not important or do not have an effect on driving the cyclic activity; in fact, they almost certainly do. Later studies have shown that a decorticated cat walking on the treadmill is able to change its gait to match the speed of the treadmill, indicating that the afferent inputs are able to directly drive the spinal CPGs (Lundberg and Phillips, 1973). Graham Brown simply demonstrated that these spinal CPGs, in the absence of other signals, are sufficient for generating cyclic activity. The intuitive next question is then, what is the mechanism that generates these alternating oscillations?
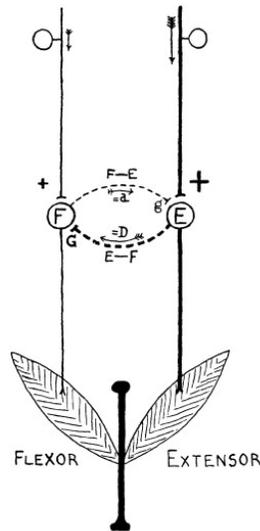
Figure 1.1: Graham Brown's half-center model for spinal CPGs. Adapted from (Stuart and Hultborn, 2008) with permission.

Grahm Brown proposed a half-center model as the neural circuit responsible for

generating the CPG rhythm. In this model, two neural populations responsible for the control of flexors or extensors reciprocally inhibit each other. This results in the domination of one muscle, e.g. extensors in the stance phase, until fatigue results in a reversal of the dominant population, e.g. flexors in the swing phase. This continues until the dominant muscle is reversed again, resetting the cycle. While the model may be lacking in complexity, the general principle of mutual inhibition is still one of the foundational components of spinal CPG models today. Over time, more complexity has been added to accommodate a wider range of phenomenon. Interneurons receiving input from 1a and 1b afferents have been added to incorporate the effect of sensory feedback on the CPGs (Perret and Cabelguen, 1980; G. Orlovsky, Deliagina, and Sten Grillner, 1999). The oscillator and motor neuron recruitment components were separated into two distinct levels to to allow for flexible pattern generation (Lafreniere-Roula and McCrea, 2005; McCrea and Rybak, 2008).

It should also be mentioned that there are likely separate CPGs for different limbs. For example, when spinalized cats were placed on a split belt treadmill, they were able to maintain their gait for the left and right side individually (Frigon et al., 2013). In a similar vein, spinalized cats are able to switch between different gaits (e.g. trop to gallop) when the treadmill speed was adjusted (FORSSBERG et al., 1980), providing evidence that the peripheral inputs could also change the inter-limb phase coupling of the CPGs. Intersegmental propriospinal neurons have been found to play a role in mediating inter-limb coordination (S. Miller and Van der Burg, 1973).

Spinal CPGs are also known to modulate spinal reflexes and descending input. For example, during the "stumbling reflex", stimulus to the dorsal aspect of the foot elicits an increase in flexion during the swing phase. One can imagine hitting an object or obstacle so this reflexive flexion would bring the limb closer to the body and

4

avoid a potential stumble. However, if a similar stimulus is applied during the stance phase, no flexion is evoked (D. M. Armstrong, 1986). Therefore, this reflexive response is gated by the phase of the gait cycle. Mechanistically, this phase-dependent response is necessary since lifting the leg while it is supporting the weight of the animal would likely cause a collapse in walking posture. Epidural spinal stimulation is also modulated by the gait phase. Gerasemiko et al. showed that the response amplitude to stimulation pulses varied depending on where in the gait cycle the limb was in when the stimulation occurred (Gerasimenko et al., 2006). Similarly, EMG responses to intracortical micro-stimulation in the motor cortex was phase-dependent (D. Armstrong and Drew, 1985). Taken together, these results indicate that spinal CPGs have a strong effect on the excitability of motor neurons to both natural and artificial inputs.

In addition to formulating more complex and realistic models for CPGs, recent work has made strides in identifying the cellular substrate underlying these models. Advances in genome sequencing techniques have begun to identify specific spinal interneuron populations that could be the embodiment of certain elements in these network models (Sathyamurthy et al., 2018). Genetic knockout techniques have also allowed for the mapping of specific receptor ligands, such as 8-OH-DPAT, buspirone, and quipazine, to CPG activation (Guertin, 2013).

Finally, while there has been many studies demonstrating the existence of spinal CPGs in felines and rodents, there are relatively few studies investigating the role of CPGs in primates. One study carried out in spinalized marmosets found that even in the absence of rhythmic afferents, the researchers were able to elicit rhythmic alteration between extensors and flexors after administering clonidine and NMDA. These results provide support for the existence of CPGs in non-human primates

(Fedirchuk et al., 1998). However, they found that other pharmacological agents, such as L-DOPA, and noradrenaline which normally induces rhythmic activity in spinalized cats were ineffective in marmosets. This inconsistency suggests that there still may be some differences in the CPG strength between primates and subprimates. Studies in humans are even more rare, which is understandable given the ethical challenges of obtaining a decerebrated, deafferented, (yet still functional) spinal cord. The strongest evidence for the existence of human CPGs so far comes from spinal stimulation. Researchers found that tonic epidural spinal stimulation in spinal cord injury patients induced alternating rhythmic movement in the legs (**Dimitrijevic1998EvidenceHumans**).

1.1.2 Involvement of motor cortex during basic locomotion

The findings obtained from spinal preparations has made it clear that descending input from supraspinal areas is not necessary to obtain cyclic firing of motor neurons. However, there are multiple regions involved in the control of locomotion in actual behaving animals. For example, the mesencephalic locomotor region (MLR) is involved in the initiation of locomotion. Shik and Orlevsky found that applying stimulation pulses (30-60Hz, 0.2-0.5ms pulse widths at 30μA) to this midbrain area induces locomotion in cats (Shik and G. N. Orlovsky, 1976). Given the importance of primary motor cortex in the generation of volitional movements, one would be interested in investigating what role it plays during locomotion.

Previous recordings from motor cortex have shown that the majority of the neural population fires cyclically during locomotion (**Drew1996RoleModifications**; T. Drew, 1988; Trevor Drew, Andujar, et al., 2008; Pearson and Gramlich, 2010; Krouchev and Drew, 2013; Trevor Drew, Jiang, and Widajewicz, 2002; Friel, Drew,
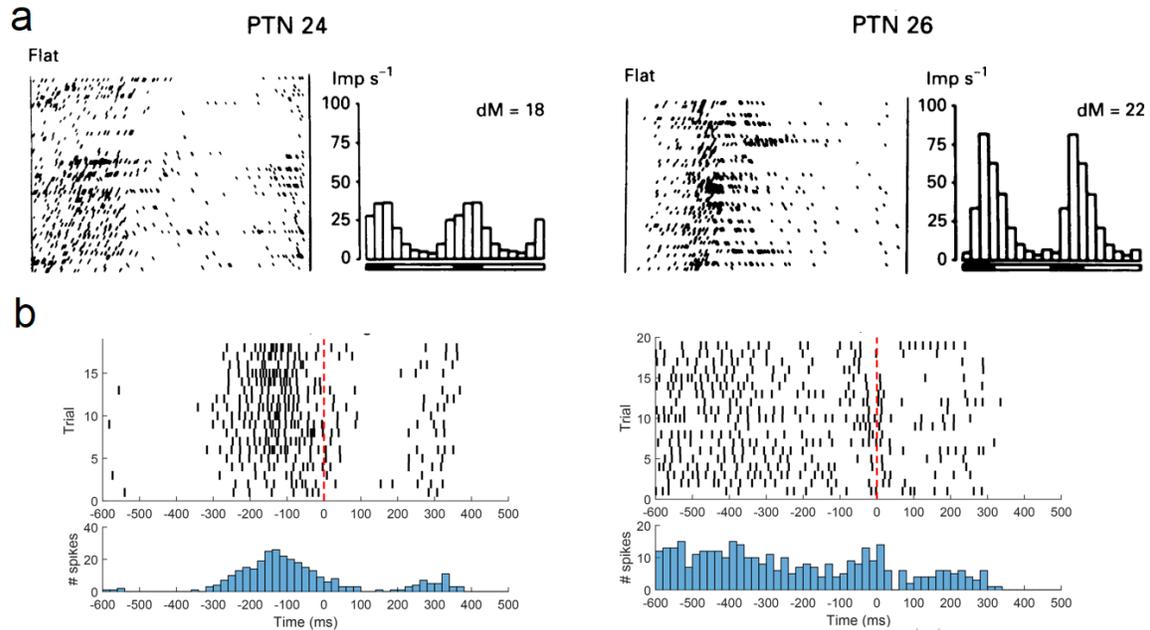
Figure 1.2: Activity of M1 neurons during locomotion. **a)** Recordings from felines during overground walking. Rasterplots represent the spike trains of one neuron across multiple trials. Histograms are binned spike counts across all trials. Here, only pyramidal tract, i.e. corticospinal, neurons (PTNs) are recorded. To identify PTNs, the medullary pyramids, which contains the descending axons of the corticospinal tract, were electrically stimulated. Stimulation of this area creates both antidromic and orthodromic action potentials along the nearby axons. Only the PTNs send axonal projections to the spinal cord so the presence of an antidromic spike immediately after the applied stimulation pulse would verify the identity of the recorded neuron. Adapted from (I N Beloozerova and M G Sirota, 1993) with permission. **b)** Recordings from rhesus macaques during treadmill walking (unpublished data). Dotted red line indicates the transition from stance to swing phase.

and Martin, 2007; I N Beloozerova and M G Sirota, 1993; Irina N. Beloozerova and Mikhail G. Sirota, 1998; Irina N Beloozerova and Mikhail G Sirota, 2003; Prilutsky et al., 2005) (Figure 1.2). That is, the neurons tend to always fire at specific phases within the gait cycle. However, it would be redundant for motor cortex to directly activate the downstream muscles cyclically. As we have shown in the previous section, central pattern generator circuits in the spinal cord are capable of generating their own phasic firing patterns without the need for descending control from supraspinal structures. The stimulation applied by (Shik and G. N. Orlovsky, 1976) to

the MRL was tonic, not phasic. Given that spinal centers are sufficient for carrying out locomotion, it is unclear what role the phasic firing pattern of cortical neurons are fulfilling. One possibility is that the phasic firing only occurs in interneurons and the pyramidal tract neurons which directly synapse to motor neurons are not actually modulated. However, this is not the case as both Trevor Drew's group and Irina Beloozerova's group specifically recorded from only PTNs (T. Drew, 1988; I N Beloozerova and M G Sirota, 1993). The more likely possibility is that the cyclic activity is the result of sensory feedback making its way to motor cortex through its recurrent connections to somatosensory cortex. Two types of sensory afferents could be candidates for inducing the phasic activity - proprioceptive and cutaneous. To decouple the two, researchers injected a local anesthetic into the distal aspect of the limb and compared the activity of the PTNs before and after the anesthetic. Little to no change was observed, indicating that it is not the rhythmic contact with the ground that is driving the cyclic activity in M1 (D M Armstrong and Drew, 1984b).

One other hypothesis is that rather than the afferents, it is the activity of the spinal CPGs themselves that is driving the activity in M1. The ventral spinocerebeller tract (VSCT) has been shown to carry activity of the spinal CPGs to the cerebellum (figure 1.3). Recordings from neurons projecting along the VSCT display phasic firing activity, even when the cord is deafferented. Therefore these cyclic modulations do not originate from sensory afferents, but carry the independent activity of the spinal interneurons (Arshavsky et al., 1972). Projections from the cerebellum can then relay the phasic signal to cortex through the ventral-lateral thalamus pathway (Middleton and Srrick, 1997). Beloozerova et al. carried out a lesion study where they destroyed the ventral-lateral thalamus bilaterally and recorded the activity of M1 neurons before and after the lesion. They found a significant decrease in the strength of modulation of the neural population during locomotion, indicating that

this pathway plays a significant role in driving the cyclic activity in M1(Irina N. Beloozerova and Mikhail G. Sirota, 1998).
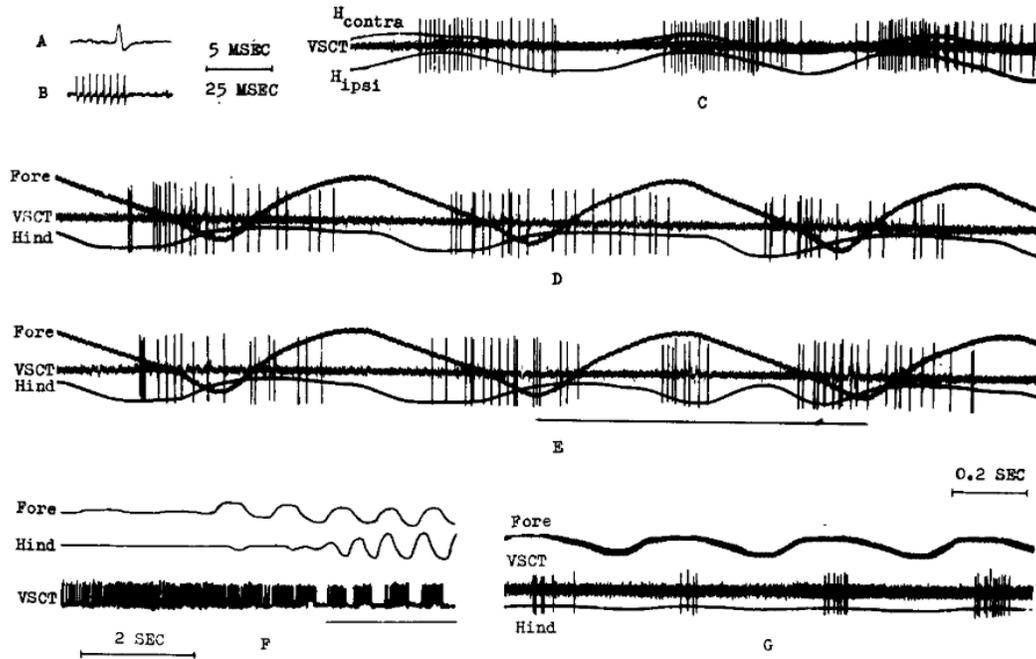
Figure 1.3: Activity of neurons from the ventral spinocerebellar tract in mesencephalic (transected at the midbrain) cat during locomotion. Top trace shows the activity of the neuron along with hip joint kinematics in a spinal cord with the dorsal roots intact. Remaining traces illustrate recordings from spinal cords where the dorsal root was severed. Adapted from (Arshavsky et al., 1972) with permission.

Recent recordings in non-human primates are in agreement with the previous findings in felines and rodents: many of the neurons in M1 are phasically tuned (figure 1.2b) (Foster et al., 2014; Fitzsimmons et al., 2009; Yin et al., 2014; Xing et al., 2019). Although there is evidence for the existence of CPGs in non-human primates (Fedirchuk et al., 1998), the importance of descending efferent signals from cortical areas could be greater than in felines and rodents. It is possible that the recorded neural activity may be actively contributing to the control of muscles rather than playing a perfunctory role. In one study, the corticospinal tract in rhesus macaques was interrupted by performing a lesion on the dorsalateral funiculus (figure 1.4a).

While they found that the animal was still able to perform locomotion, there was significant paw drag during the swing phase (figure 1.4b) (Courtine, 2005). However it appears that their partial hemisection was not very specific, and it is not clear whether the extent of lesion was solely confined to the coticospinal tract. Damaging other tracts or inter-segmental axons could be a confounding factor in evaluating the lesion results. A truly specific lesion to motor cortical neurons will be necessary to interrogate whether M1 is playing an active role in controlling the hind-limb muscles in the monkey. In chapter 5, I discuss the implementation of genetic tools to obtain a tract-specific lesion.
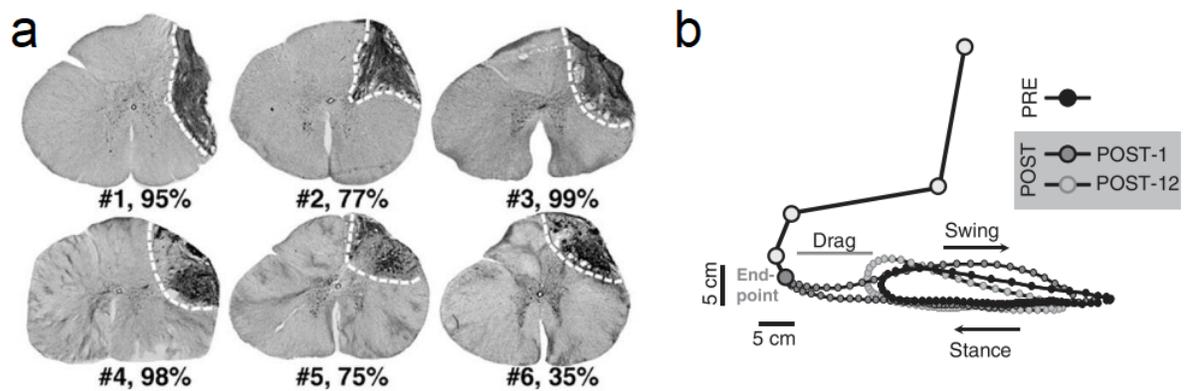
Figure 1.4: A unilateral partial hemisection was applied to rhesus macaques in order to interrogate the necessity of the corticospinal tract in locomotion. **a)** Histological assessment of the lesions. **b)** kinematics of the end-point (toe) as the animal walked on the treadmill before and after the lesion. Adapted from (Courtine, 2005) with permission.

1.1.3 Involvement of motor cortex during complex locomotion

While the causal role of motor cortex in locomotion is still uncertain, it's role in attentive, voluntary movements is much better established. For example, the same animals shown in figure 1.4 were completely unable to complete a foot-grabbing task

after the lesion, despite retaining some walking ability. Kuypers et al. performed lesions to the medullary pyramids in macaques, which, like in (Courtine, 2005), would disrupt the corticospinal tract and abolish any descending efferent signals from motor cortex. They also found that the animals retained the ability to walk overground and climb up cages, but were unable to reach and grab for food (Lemon et al., 2012). These results suggest that motor cortex may be deferentially engaged between locomotion and dexterous volitional movements.

To compare the behavior of motor cortex during these different behaviors, groups have recorded from M1 neurons during basic locomotion, and during obstructed locomotion. In the latter, an obstacle is presented during walking, and the subject must modify its gait in order avoid hitting the obstacle. This constitutes a voluntary movement and requires top-down input from higher cortical areas. Unsurprisingly, all studies consistently find that there is a change in the neural activity when the gait is modified to accommodate the obstacle (T. Drew, 1988; Yakovenko and Trevor Drew, 2015; Trevor Drew, Andujar, et al., 2008; I N Beloozerova and M G Sirota, 1993; Irina N. Beloozerova and Mikhail G. Sirota, 1998; Zelenin et al., 2011). Recent studies have focused on investigating whether the relationship between the neural activity in M1 and the movement is preserved between basic unobstructed walking and volitional movements.

Yakovenko et al. compares the neural activity during two different volitional movements. For the first task, they trained cats to walk on a treadmill and step over an obstacle; for the second task, trained the animals to place their forelimb over an obstacle while standing still. The movement and muscle activity were similar for both tasks. They found that the timing of the bursts of pyramidal tract neurons coincided with the activation of synergistic muscle groups for both obstacle and basic

11

locomotion (Yakovenko and Trevor Drew, 2015). These results have implications for the relevance of the cyclic activity observed in M1 during basic locomotion, as the only difference between the two tasks was that in the first task, the animals were continuously walking while the in the other, the animals were stationary. The consistency of the neural-muscle correlations may be further evidence that for felines, the activity of M1 during locomotion may not be directly responsible for controlling the muscles.
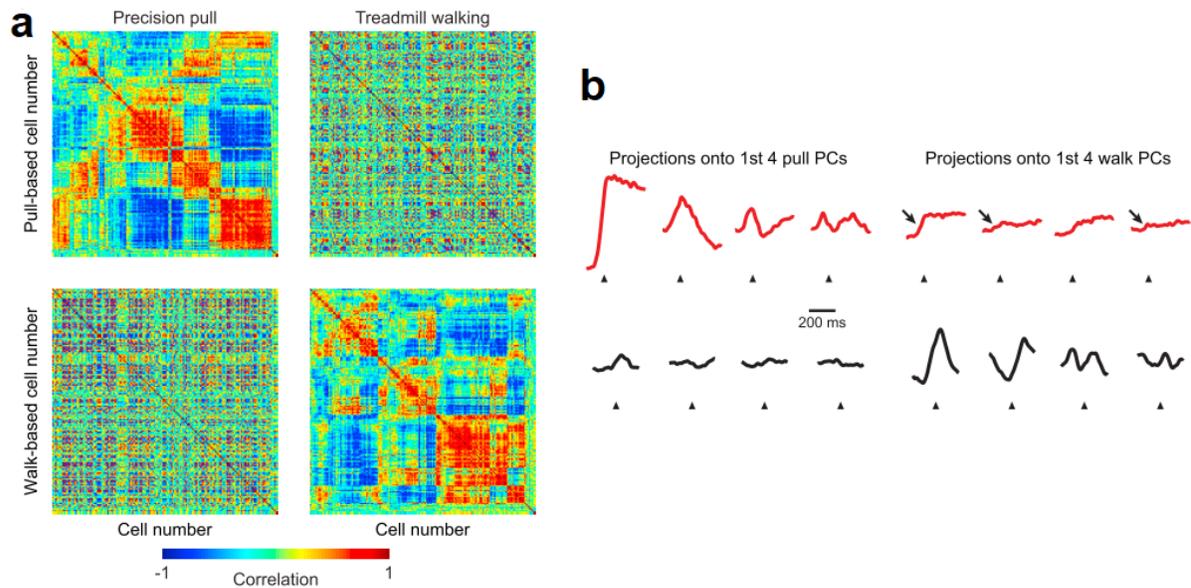
Figure 1.5: **a)** Neural correlation structure. Ordering of rows were preserved for the top two and bottom two plots. **b)** Left: projection of neural activity during lever pulling into either the PCA subspace of lever pulling (red) or the PCA subspace of walking (black). Right: projection of the neural activity during walking into either the PCA subspace of lever pulling (red) or walking (black). Adapted from (Miri et al., 2017) with permission.

A separates study investigated the activity of M1 neurons in mice. Rather than compare two volitional actions, here they compared the M1 activity during a volitional movement (stationary lever pulling) and during basic locomotion (Miri et al., 2017). They found that the correlation structure within the neural population were completely different between the different behaviors (figure 1.5). Additionally, they

looked at the combined activity of the population of neurons. The activity of the neural population can be represented in a low-dimensional subspace. While the dimensionality of these subspaces may be small, they can still account for most of the activity of the whole neural population. This is because there are correlations in the neural activity so it is redundant to represent this activity in the full high-dimensional space. Critically, if the activity of the neurons are correlated in the same way, then the subspace of that activity would be the same. However, if the correlation of the neurons suddenly change, then the subspace in which the new activity resides will be different from the previous subspace. Therefore if the correlation structure of the neurons are consistent between lever pulling and walking, then we would expect the subspace of their activity to be similar. The left plots in figure 1.5b correspond to the pulling subspace while the right plots correspond to the walking subspace. The same neural activity projected onto both of them results in drastically different signals, indicating that motor cortex engages separate subspaces when performing a volitional movement and basic locomotion. This is consistent with the behavioral findings of the lesion studies.

While this subspace analysis reveals the flexible nature of M1 network structure across different behaviors, it does not address how M1 would be able to integrate the two behaviors, such as, for example, carrying out a volitional movement while maintaining locomotion. In Chapter 4 I will investigate the low-dimensional structure of M1 network activity during such a behavior.

1.2 Brain-machine interfaces

A Brain machine interface (BMI) is a neurotechnology aiming to help patients who have experienced some sort of motor or sensory deficit to return to an operative lifestyle (Leigh R. Hochberg et al., 2012; Wodlinger et al., 2015; Bouton et al., 2016; Collinger et al., 2013; Capogrosso et al., 2016; Flesher et al., 2016). Motor BMIs read from the nervous system in order to generate an action, while sensory BMIs take in information about the world and write it in to the nervous system. In this work, I will be focusing on motor BMIs.

Most motor BMIs operate under the paradigm of 1) record signals from intact parts of the nervous system, 2) extract information regarding intended action from the recordings, and 3) use this data to either drive artificial actuators, such as robotic limbs or computer cursors, or, natural biological actuators, e.g. intact but enervated muscles. For example, neural firing rates from motor cortex may be recorded using a multielectrode array, and then forwarded to a processing computer. The firing rates are mapped through a mathematical model to a kinematic variable, such as the intended position of a joint in space. A robotic limb, for example, may then use this decoded positional information to bring itself to the intended location. A user can then voluntarily control the prosthetic limb through modulating the firing rates of the recorded neurons.

1.2.1 Input signals

There are many different parameters that can determine the efficacy of a BMI. On the recording side, the quality of the neural signal has an effect on how much information can be extracted from it. A highly noisy electrode, for example, may have

the relevant neural signal drowned out by the noise. There could be other trade-offs. EEG may be easy to use and requires no surgical intervention, but its spatial resolution is poor. And while intracortical electrodes provide excellent spatialtemporal resolution, it requires highly invasive surgery. Generally, progress on the recording front requires engineering advancements and development of new technologies.

1.2.2 Decoder implementations

The computational component of BMIs is another avenue for innovation. The decoder is the function that takes in the recorded signals and maps them to the output. One of the earliest implementations was the population vector algorithm which took the sum of the preferred vectors of each neuron, weighted by the firing rate (Georgopoulos and Grillner, 1989; Jarosiewicz et al., 2008; Ganguly, Dimitrov, et al., 2011). Wiener filter decoders, which took a linear combinations of the inputs with a fixed number of the inputs' past values, were also used (Ganguly and Carmena, 2009; Kim et al., 2008a). Perhaps the most widely used is the Kalman filter, which is based on following Gaussian linear dynamic state-space model:

$$
\begin{aligned}
\mathbf{X}_t &= \mathbf{A}\mathbf{X}_{t-1} + \boldsymbol{\epsilon}_t \\
\mathbf{Y}_t &= \mathbf{B}\mathbf{X}_t + \boldsymbol{\gamma}_t
\end{aligned}
\tag{1.1}
$$

where Y_t is the observation vector, X_t is the state being decoded, B is the observation matrix mapping the state to the observations, and A is the time evolution matrix updating the state using its previous values. The terms $\boldsymbol{\epsilon}_t$ and $\boldsymbol{\gamma}_t$ denote i.i.d. Gaussian random vectors. Under the above model and Gaussian initial conditions, the state posterior density $p(X_t|Y_{1:t} = y_{1:t})$, i.e. the probability density of the state at time t conditioned on given observations up to time t, remains Gaussian throughout. Thus,

one only needs to update at each time t the mean and the covariance of a Gaussian probability density. The Kalman filter solves recursively for the mean and the covariance (Bishop, 2006). These time-varying means and covariances can then be used for decoding and assessment of uncertainty.

An different approach would be to alter the inputs before they are given to the decoder, a process called feature selection. For example, one could apply PCA on the recorded signals to reduce the dimensionality of the feature space, which could alleviate problems with over-fitting. A recent study has found that extracting the low-dimensional dynamics from the input space before feeding it to a kalman decoder improved performance in a cursor control BMI (Kao et al., 2015). In chapter 2, I will discuss the viability of using a dynamical systems model in extracting underlying neural features for decoding kinematics during locomotion.

1.2.3   BMIs for restoring locomotion

Most of the early BMI systems were focused on restoring functionality usually associated with the forelimb, such as reaching and grasping, or controlling a computer cursor (Leigh R Hochberg et al., 2006; Leigh R. Hochberg et al., 2012; Serruya et al., 2002; Brandman, Cash, and Leigh R. Hochberg, 2017; Taylor, Tillery, and Schwartz, 2002; Kim et al., 2008b; Collinger et al., 2013). However, recent years have witnessed the emergence of hind-limb BMIs. For example, a system developed by Donati et al utilized EEG recordings to drive the stepping of a powered exoskeleton (Donati et al., 2016). I was fortunate enough to be involved in the development of a brain-spinal interface (BSI), a system aiming to restore walking ability after spinal cord injury (Capogrosso et al., 2016). Rather than using a mechanical actuator such as an exoskeleton, we leveraged the spinal cord itself to reactivate the subject's own muscles.

In this system, cortical recordings were obtained from microelectrode arrays implanted in M1. A classifier used the neural inputs to predict whether the subject was about to put the foot down on the ground, initiating the stance phase, or lift the foot off the ground, initiating the swing phase. In the case of the former, the extensors need to be activated to bear the weight of the body, while in the case of the latter, the flexors need to be activated to lift the foot up. We were able to selectively activate flexors or extensors by applying epidural spinal stimulation at specific spinal levels and voltage amplitudes. We were therefore able to activate the specific muscles at the specific times required to facilitate stepping. Figure 1.6 illustrates the patterned stimulation that was applied during treadmill walking. Application of our BSI system to a spinal cord injured animal saw the immediate recovery of stepping movements in the injured limb. While the system was highly effective, the decoding and processing of the inputs were rather basic. In Chapter 2 I investigate whether using dimensionality reduction as a pre-processing step could improve decoding accuracy for locomotion movements.
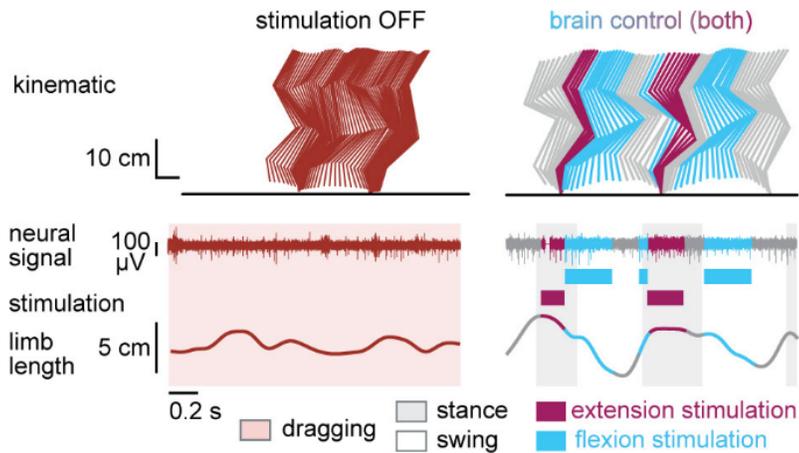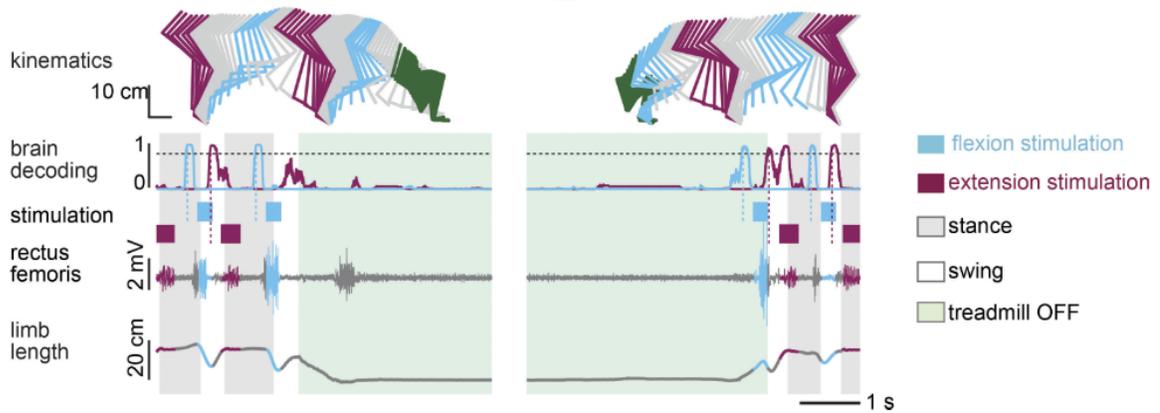
Figure 1.6: Top: Our classifier was able to predict the onset of flexion, extension, or no activation as the animal initiated and terminated locomotion on a treadmill. Bottom: After applying a spinal cord injury model, animals had difficulty using the ipsilateral limb. However, immediately after activating the BSI system, they recovered stepping ability in the limb. Adapted from (Capogrosso et al., 2016)

Chapter 2

Extracting locomotion kinematics from neural dynamics

Parts of this chapter have been previously published as an original article: (Xing et al., 2019)

2.1 Introduction

With the advent of large scale intracortical recordings allowing for the simultaneous interrogation of dozens to hundreds of neurons, the study of the role of motor cortex in the generation of movement has been steadily moving towards investigation of cortical dynamics in the brain on the population level. Ensembles of cortical neurons are organized as recurrently connected networks, introducing shared variability among the constituent cells, in turn constraining firing activity to a lower-dimensional space (Yu et al., 2009; Truccolo, Leigh R Hochberg, and Donoghue, 2010; Afshar et al., 2011; Mark M. Churchland et al., 2012; Sadtler et al., 2014; John P Cunningham and Yu, 2014; Pandarinath et al., 2018a; Gallego, Perich, L. E. Miller, et al., 2017). Dimensionality reduction techniques can be used to extract the coordinated neural activity of a population, and reveal structures that may be hidden at the isolated

single-neuron level. Previous investigations have shown that for reaching movements with the arm, semi-oscillatory dynamics, inferred using jPCA, are a common feature underlying a variety of different reach movements (Mark M. Churchland et al., 2012). Additionally, a common low-dimensional neural manifold underlying various wrist movements have be found using demixed principle component analysis, or dPCA (Gallego, Perich, Naufel, et al., 2018).

Besides dimensionality reduction, many of these techniques also employ dynamical systems models. These models address how the values at the current time step depend on the values at previous time steps, usually in the form of a temporal transition matrix. Poisson Linear-Dynamical-System (PLDS) is one such technique which employs both unsupervised dimensionality reduction as well as explicit temporal dynamics. PLDS maps low-dimensional latent states to the measured high-dimensional neural spiking activity through an observation point-process model and explicitly estimates the dynamics of these latent states as it evolves through time with a linear mapping (Macke et al., 2011; Truccolo, Eden, et al., 2005). Unlike PCA and dPCA, PLDS models the low-dimensional space as a state-space in a linear dynamical system, thereby explicitly accounting for the temporal relationships in the population (figure 2.1). Although these low-dimensional dynamics are inferred through an unsupervised process, they are able to retain relevant behavioral information. In reaching behaviors, explicit state-space models similar to PLDS have been shown to increase closed-loop BMI performance in cursor-control tasks (Kao et al., 2015) and decoding accuracy during forelimb reaching behaviors (Aghagolzadeh and Truccolo, 2014; Aghagolzadeh and Truccolo, 2016), demonstrating that only a small number of dimensions from the neural population space is needed to capture the movement kinematics during voluntary forelimb movements.
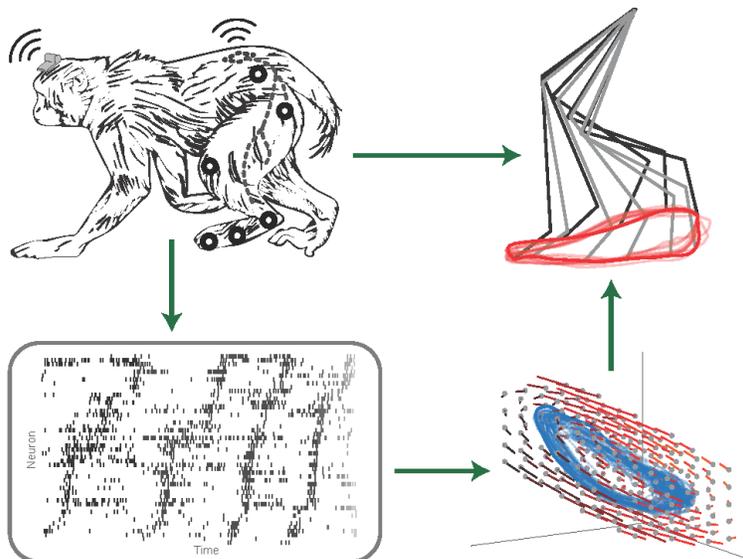
Figure 2.1: Diagram illustrating the the flow of information for neural decoding. Observable behavior such as leg joint positions can be simultaneously recorded with neural activity from the motor cortex during freely moving locomotion (top left). Leg kinematics follow a stereotyped cyclical pattern (top right). Ensemble activity in the leg area of the motor cortex is represented as a raster plot of sorted neurons (bottom left). Dimensionality reduction on population-level spike counts reveals cyclical trajectories constrained to a lower-dimensional subspace. PLDS explicitly incorporates dynamics into the dimensionality reduction model as a time-evolution matrix, represented by the slope field (bottom right). Three latent dimensions (LD) are visualized here. The information contained in these latent-state dynamics can then be used to decode the leg kinematics during locomotion. Illustration of the animal provided by Ellen Xing.

However, reaching actions with the arm are typically highly precise and have a strong voluntary control component, whereas locomotion movements have a higher degree of autonomy. For example, in felines, injection of the neurotoxins or ablation of the motor cortex does not affect the ability of the animal to walk along a flat surface, although their ability to step over a ladder and over obstacles is impaired (**Drew1996RoleModifications**; I. N. Beloozerova and M G Sirota, 1988). These results suggest that motor cortex may play a less active role in the control of limb movements during basic, unobstructed locomotion compared to movements that require top-down voluntary control. In non-human primates, inducing corticalspinal

tract lesions showed that although there is some locomotion deficits post-lesion, these recover quickly, while dexterous foot grasping remain severely impaired, even after 3 months post-lesion (Courtine, 2005).

These findings suggest that the contribution of M1 to the control of movements might be different during locomotion compared to during reaching, although how exactly the role of M1 is changing between these two behaviors is still not well understood. It has been well known for several decades that M1 is active during walking, and that cortical neurons are phasically tuned to the gait cycle (T. Drew, 1988; Trevor Drew, Andujar, et al., 2008; Irina N. Beloozerova and Mikhail G. Sirota, 1998). Yakovenko and Drew recorded from corticospinal neurons during reaching and during walking over an obstacle in cats and found that the firing onset phase of certain neurons correlated with the onset phase of muscle activation for both types of movements (Yakovenko and Trevor Drew, 2015), suggesting similar encoding of movements in M1 for both types behaviors. However, a recently published study in mice found that the population-level structure of M1 neurons is disparate during reaching and lever pulling compared to simple treadmill walking (Miri et al., 2017). These contrasting conclusions demonstrate that there is still a lack of consensus on the role of motor cortex during locomotion-related activities, and it is still unclear whether M1 is contributing to the control of the limbs in a similar manner during walking as during directed reaching movements. In particular, it has yet to be shown whether the latent state-space models that capture movement parameters during precise arm reaching would also be able to capture hind-limb movements during the potentially less engaging act of locomotion. While previous studies have found low dimensional representations of motor cortex activity in non-human primates during simple treadmill walking (Yin et al., 2014; Foster et al., 2014), they have not shown that movement kinematics are preserved in these dynamics (figure 2.1).

Here, we aim to determine whether the PLDS latent-state model is able to extract low-dimensional dynamics which are informative of the limb movements. We define informative as having the ability to decode hind-limb joint kinematics as well as gait phase, during various locomotion tasks such as treadmill, corridor, and horizontal ladder walking (figure 2.2a). To test this, we used either firing rates of the full recorded population of single neurons or the corresponding low dimensional dynamics as input features into a Wiener filter decoder which attempts to reconstruct the kinematics from these inputs. We demonstrate that for PLDS inputs, only a small number of dimensions are necessary to decode limb kinematics and gait phase as accurately as the full neural population. As far as we are aware, this is the first employment of explicit state-space models during both basic locomotion along a treadmill and corridor, as well as during directed locomotion along evenly and unevenly spaced ladders in nonhuman primates. Furthermore, it is also the first demonstration that cortical ensemble dynamics robustly captures behavioral information such as limb kinematics and gait phase during these different ambulatory behaviors.

2.2 Methods

2.2.1 Surgery

Five male rhesus macaques between 5 and 8 years of age were implanted with 96-channel microelectrode arrays (Blackrock Microsystems, UT) in one hemisphere of the leg area of primary motor cortex (M1), located medially along the precentral gyrus (He, Dum, and Strick, 1993). The details of the surgical implantation procedures have been described previously (Yin et al., 2014). Experiments complied with the European Union directive of September 22, 2010 (2010/63/EU) on the protection of animals used for scientific purposes in an AAALAC-accredited facility following

acceptance of study design by the Institute of Lab Animal Science (Chinese Academy of Science, Beijing, China). Experiments were approved by the Institutional Animal Care and Use Committee of Bordeaux (CE50) under the license number 50120102-A.

2.2.2 Kinematic Data Collection

To obtain kinematic data, white reflective markers were painted over the shaved skin overlaying the right iliac crest (crest), greater trochanter (hip), lateral condyle (knee), lateral malleolus (ankle), 5th metatarsophalangeal (metatarsal), and outside tip of the fifth digit (toe tip). Marker locations were identified by feeling for the bony projections of the above anatomic landmarks under the skin; see (Capogrosso et al., 2016) for a visualization of the marker locations. Videos were captured at 100Hz from four high-speed cameras placed around the enclosures. Motion tracking software (Simi motion systems, Germany) was used to determine the 3D spatial coordinates of each marker, after calibrating the cameras in 3D space using known calibration objects at the beginning and end of each session. The origin of the coordinate system was set to the iliac crest of the animal. The horizontal axis was set to the direction of walking, while the vertical axis was set to the height off the ground. Joint angle was set as the inner angle between the two limb segments. The polar nature of the joint angles can confound error calculations (for example, 0 degrees is closer to 350 degrees than 180 degrees, although 350-0 > 180-0), so instead of decoding the joint angles directly, we decoded the sine and cosine of each angle. Additionally, the total leg extension was calculated as the distance between the iliac crest and the metatarsal joint. Neural data, sampled at 30 kHz, was synchronized with the kinematic data using a camera-start trigger signal at the start of each trial of each recording session.

24

2.2.3 Tasks

Subjects carried out various behavioral motor tasks (figure 2.2a). Monkeys were trained to walk in an enclosed treadmill at speeds of 1.1, 1.6, 2.4, 3.2, 4.0, 4.8, and 6.4 km/h (TRM trials). The enclosure was approximately one and a half meters long by one meter high by half a meter wide. It was constructed from clear Plexiglas, which is both visually and RF transparent. Spontaneously, they would switch to a bipedal gait, or start walking backwards (BIP and BACK trials, respectively). In addition, monkeys walked along a straight 3m long corridor over a flat surface (CORR trials) and over a horizontal ladder with rungs either evenly spaced 35 cm apart or irregularly spaced (LAD trials). These tasks were self paced, with the average corridor walking speed at 3.24 km/h (0.79 km/h standard deviation) and the average ladder walking speed at 3.49 km/h (0.79 m/s standard deviation). Monkeys were trained for at least one month to walk to the end of the corridor or ladder to receive a food reward in response to an auditory beep and flash of light.

Recording was carried out over 1-2 days for each monkey. Each day consisted of recording different randomly interleaved sessions for each task. Each treadmill session consisted of approximately one minute of walking. These sessions were divided into epochs which contained only those time periods where the monkey was performing a single clear ambulatory behavior (e.g. epochs consisted of only bipedal walking, or only forward quadrupedal walking). Corridor and Ladder sessions consisted of walking back and forth along the corridor or ladder for a total of 20 to 30 rounds. These rounds contained only the steps performed at the middle of the corridor or ladder, to avoid the initiation and termination phase of locomotion.

2.2.4   Data Preprocessing

Neural data was obtained using the Blackrock Cerebus system, and saved onto a local computer through the Blackrock Central Suite program. The data was transmitted through a custom wireless headstage system (Yin et al., 2014), allowing for freely-moving recordings. The neural signal was band-passed at 500-5000Hz (second order Butterworth filter) and spikes were extracted by manual sorting using custom Matlab scripts (Laurens, Meng, and Angelaki, 2013), which performed PCA on spike amplitudes and peak velocities (although in practice, any commonly accepted spike sorter can be used to obtain the firing rates). Neurons (the total number ranging between 18-80 for each session) were extracted from the multi-unit activity for each recording day (Table 2.1) and the spike times (counts) were binned into 10ms intervals to match the 100Hz kinematic data.

Each epoch was divided into trials consisting of a single gait cycle by manually marking the time point of foot-strike and toe-off. The stance phase of a single gait cycle was defined as the time period between the foot-strike and toe-off and represents 0%-60% of the gait cycle while the swing phase was defined as the period between toe-off and the next foot-strike, and represents 60%-100% of the gait cycle. The time-varying gait phase percentage was linearly interpolated from the foot-strike and toe-off time points (defined as the 0%/100% and 60% mark, respectively).

Occasionally, neural data would become lost or corrupt for periods within a trial. Trials were manually inspected for data corruption, and if a trial was missing neural data, it was excluded from the analysis. All the trials shown in Table 2.1 had complete neural data. Additionally, for two of the recording sessions, we did not record kinematics, and for the other seven sessions where kinematics were recorded, there were occasional trials where we were unable to obtain the kinematics (for example, due to

video occlusion). Table 2.1 contains the number of trials that contain usable kinematics for each task during each recording day. However, despite having incomplete kinematics for some of the trials, we were able to obtain gait phase data for all of the trials shown in Table 2.1, including in the two sessions where we were unable to obtain any kinematics. Therefore, our gait phase decoding analysis has a sample size of 9 (figure 2.5c) while the kinematic decoding analysis has a sample size of 7 (figure 2.5d).

	Neurons	TRM Trials	BIP Trials	BACK Trials	LAD Trials	CORR Trials	Trials with Kinematics
Q1 Day1	56	96 (52)	43 (0)	4 (0)	35 (0)	18 (7)	59
Q1 Day2	80	91 (78)	61 (31)	4 (0)	14 (14)	5 (5)	128
Q2 Day1	63	129 (43)	12 (0)	0 (0)	6 (3)	6 (5)	51
Q2 Day2	51	217 (193)	22 (22)	0 (0)	22 (13)	4 (2)	230
Q3 Day1	46	77 (77)	6 (4)	0 (0)	10 (1)	5 (4)	86
Q3 Day2	41	79 (0)	0 (0)	0 (0)	16 (0)	12 (0)	0
Q4 Day1	51	224 (139)	0 (0)	0 (0)	0 (0)	0 (0)	139
Q4 Day2	18	51 (0)	0 (0)	0 (0)	3 (0)	20 (0)	0
Q5 Day1	39	37 (37)	0 (0)	0 (0)	26 (8)	9 (7)	52

Table 2.1: Number of trials and neurons recorded for each recording day of each animal. We were unable to obtain kinematics in some of the trials; the number of trials that did contain kinematics for each task are displayed in parenthesis, and the total number of trials that contains kinematics across all tasks is displayed in the last column. All the trials shown here contain complete neural data and gait cycle data.
.

2.2.5 Dimensionality reduction models

We used explicit state-space models to estimate the latent states of full population neural activity during locomotion. A LDS model assumes the neuronal ensemble activity as Gaussian linear observations, and uses expectation maximization (EM) learning to estimate the unknown model parameters and the latent states given only the observations (Macke et al., 2011; Aghagolzadeh and Truccolo, 2014). To account for the count process nature of ensemble spiking activity, we used a PLDS model, adding the assumption that the neural observations are conditionally Poisson given

latent states (Aghagolzadeh and Truccolo, 2016; Truccolo, Eden, et al., 2005). A Laplace approximation was used to compute the posterior density of the latent states given the neural observations. For decoding latent states from novel neural data, we used the mean of the state posterior density under the Laplace approximation as the estimate for the latent state (figure 2.2b). The algorithm details have been described in (Macke et al., 2011; Aghagolzadeh and Truccolo, 2016).

To compare PLDS with other dimensionality reduction techniques, we also tested decoders using low-dimensional inputs derived via principal component analysis (PCA), and also the activity of an optimal subsample of neurons from the full population, referred to as predictive subsampling (PSS). For PCA, the trials were concatenated across time and z-scored. The covariance matrix was computed, and the PCA projection matrix was constructed by eigenvalue decomposition - stacking the eigenvectors corresponding to the n largest eigenvalues, where n is chosen as the number of dimensions. We note that although PCA (and related methods) provide a low-dimensional representation of the ensemble activity, unlike PLDS, they do not explicitly account for any temporal dynamics in the latent states. PSS selects a subset of n neurons from the full population that optimize neural decoding through a greedy supervised learning algorithm, hence the name predictive. Details of this method can be found in a previous paper (Aghagolzadeh and Truccolo, 2016).

2.2.6 Comparison of low-dimensional neural trajectories

To generate the example neural trajectories shown in figure 2.3, we selected one recording session, Q1-day 1, which contained examples of all the tasks and treadmill speeds. In order to compute the average trajectories, the PLDS latent states were time-normalized to 0-100% of the gait cycle in steps of 1% using linear interpolation, with 0% representing the start of the trial (start of the stance phase) and 100%

representing the last time point of the trial (end of the swing phase), while 60% was set to the transition between stance and swing. This resulted in 100 time points for each trial, and the values were averaged at each time point. The distance metric was calculated for each trial as the Mahalanobis distance between the values of that trial and the distribution of a selected reference trial type in the full 12 dimensional space. For the comparison across tasks, reference task was treadmill walking and for the comparison across speeds, the reference was the 6.4 km/h trials. The Mahalanobis distance values were averaged across trials for each of the tasks/speeds. The standard error of the mean (SEM) was also computed and shown for each task/speed.

2.2.7 Frequency analysis

To compare the smoothness of the decoder outputs, the power spectral density (PSD) of the decoded kinematics was computed for each of the four decoders, along with the real kinematics. For each trial, we estimated the PSD using Thomson's multitaper method (*pmtm()* function in Matlab), and the values were averaged across all the trials of all the sessions. This was done for each of the decoded kinematic variables (figure 2.4).

2.2.8 Cross-Validation

We constructed four different decoders (which we will refer to as the Full Population, PSS, PCA and PLDS decoders) and determined the performance of each through a 10-fold cross-validation paradigm (figure 2.2b). The PSS, PCA and PLDS decoders employed a two-stage method where binned spike counts were projected onto a low-dimensional space, with dimensionality varying from 1 to 20. The coordinates of the neural data in this space were then utilized as input features for the Wiener filter described below. The Full Population decoder did not employ the dimensionality reduction stage and utilized the z-score normalized spike counts of all recorded

29
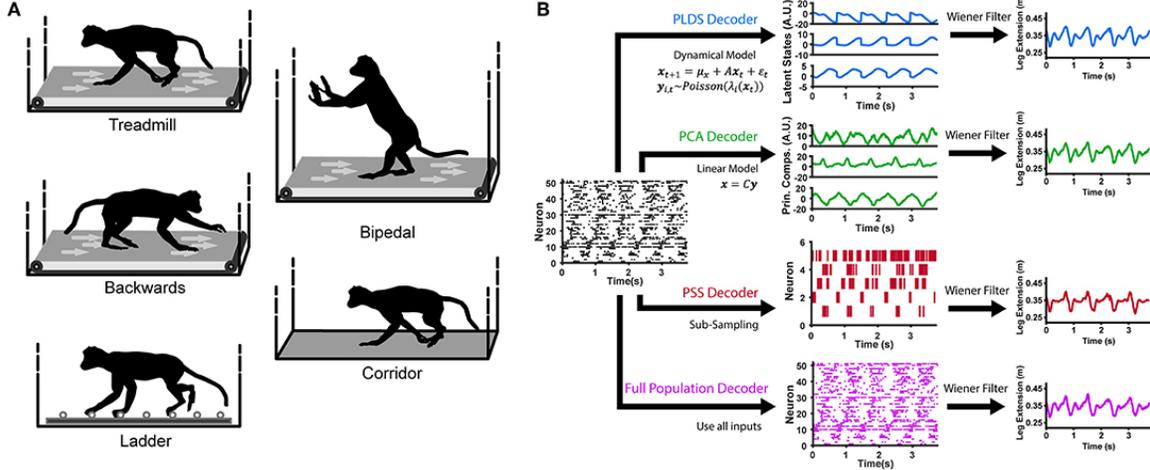
Figure 2.2: Locomotion behaviors and neural decoding with latent state models. **a:** Animals were trained to perform different locomotion tasks in a freely moving, untethered environment. The tasks include basic treadmill walking at various speeds, backwards treadmill walking, bipedal treadmill walking, corridor walking, and ladder walking. **b:** Construction of the neural decoders. For decoders implementing dimensionality reduction (PLDS, PCA, PSS), rasters of neural spike counts for each gait cycle (left plot) undergo either an orthogonal linear transformation to extract principle components (center, top), a count-process filter based on a state-space model to extract low-dimensional latent trajectories (center, middle) or a greedy search algorithm to obtain an optimal subsample of neurons (center, bottom). These low-dimensional features are then used as inputs to a Wiener filter for reconstructing the estimate of a desired locomotion variable, such as the leg extension distance (right). The Full Population decoder (bottom path) performs decoding directly on all recorded normalized neural spike counts, binned at 100ms. For decoding kinematics and gait, the trials of all tasks were combined and shuffled for each session.

neurons as input features (figure 2.2a). The PLDS inputs assumes spike counts, and utilized the 10ms bins, while the PCA, PSS, and full population utilized firing rates. The number of spikes is too sparse to estimate the firing rate using only 10ms counts, so to estimate the firing rates at each 10ms time step, the number of spikes in the current time bin as well as the previous nine time bins were summed and divided by 100ms to get the rate in spikes/s. For all decoders, the time steps of the inputs was 10ms.

Cross validation was performed on each recording session of each monkey. All

of the trials of all five tasks were combined and randomized. The trials of all the tasks were then divided into 10 cross-validation blocks. Each block consisted of 5-24 (for decoding kinematics) and 7-27 (for decoding gait) trials which were used as the testing set while the remaining nine blocks were used as the training set. Decoder coefficients were calculated from the training sets using the least-squares regression algorithm between the neural data (Full Population and PSS) or latent states (PCA and PLDS) and the measured output signal (kinematics, or gait phase). The decoder was then used to estimate the output signals of the testing set, and the decoding accuracy was calculated as the coefficient of determination (R^2) between the estimated and the true signal as defined by equation 2.1. n is the number of points of the kinematic or gait variable in the trial, y_i is the actual variable value at point i, $\hat{y}_i$ is the estimated value from the decoder at point i, and $\bar{y}$ is the average value of the kinematic or gait variable for the trial.

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \tag{2.1}$$

2.2.9 Decoder

A linear Wiener filter of order 10 (the filter order that gave the best decoding results under cross validation) was used to decode the kinematics and gait cycle phase from the neural data. The decoded signals included the horizontal and vertical positions of the hip, knee, ankle, metatarsal, and toe tip joint markers, the joint angles, the leg extension, and the percentage of the gait cycle phase. The decoder is described by the equation:

$$y_t = \sum_{n=0}^{9} A_n X_{t-n} \tag{2.2}$$

where y_t is the output signal being decoded at time t, and X_{t-n} is a vector of the decoder inputs at lagged times $t - n$, for $n = 0, \ldots, 9$. A_n is the corresponding vector

of regression coefficients computed from the training set.

In addition to continuous decoding, we used a kNN classifier as a discrete decoder to determine the gait state (swing or stance) at each time point. Here, the distance between the test and training samples was defined as the euclidean distance and k was set to 7. Decoding of the gait state was performed under 10-fold cross validation in the same manner as the continuous decoding.

Training Set Size Analysis

To investigate whether the PLDS decoder would maintain performance with restricted amounts of training data, we performed an analysis where the decoders were calibrated on incrementally smaller training sets. We again combined and randomized all the trials of all the tasks for each session, and split them into 10 blocks. For each cross-validation trial, a block was selected as a testing set, and the remaining nine blocks were combined into the training set. We then used 90%, 50%, 40%, 30%, 20%, 10% of the trials of the training set to calibrate the decoder before measuring decoding performance on the testing set. The trials used to train the decoder were randomly chosen.

Statistical tests comparing the R^2 values were carried out in Matlab (*signrank()* function).
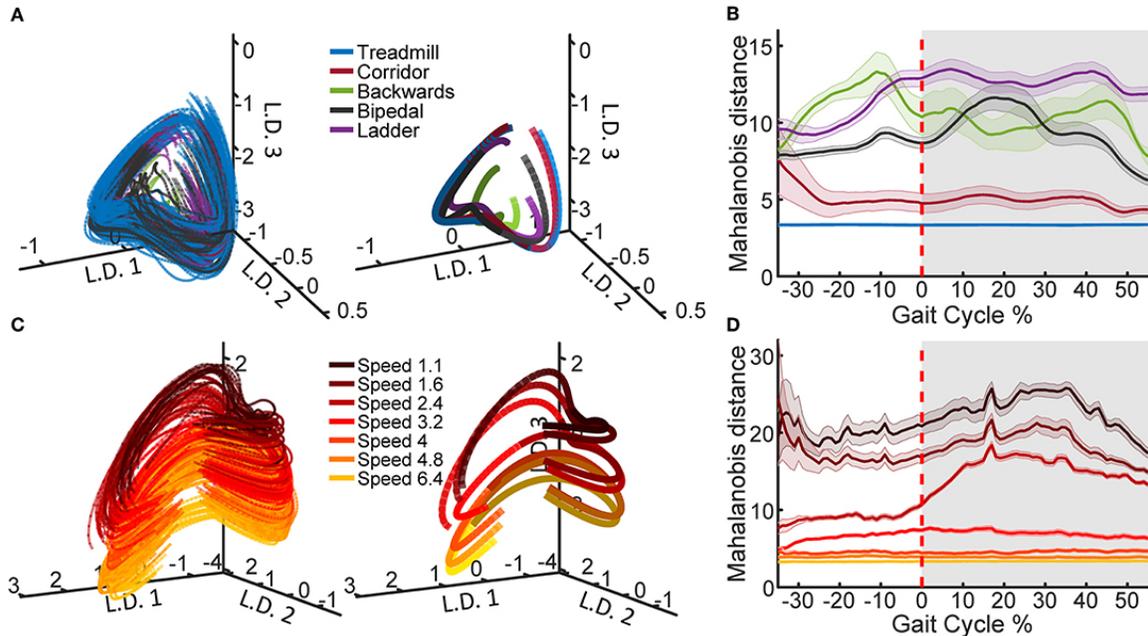
Figure 2.3: Latent state space of neural population activity during locomotion. **a:** Neural trajectories obtained from the PLDS model for the gait cycles of all tasks in animal Q1 on recording day 1. The trajectories in the first three latent dimensions (L.D.) for each individual gait cycle are shown (left), as well as task-averaged trajectories over the total gait cycle (right). **b:** The Mahalanobis distance between the trajectories of each of the tasks and the trajectories of the treadmill task, calculated on all 12 dimensions. Shaded region denotes stance phase (0-60% of the gait cycle), shaded bars denote 1 S.E.M. **c:** PLDS state-space trajectories in 3 dimensions for all treadmill walking trials at different locomotion speeds in animal Q2, session 2. Darker colors represent slower speeds; all speed values displayed in the legend are in km/h. As in (a), individual gait cycles (left) and gait cycles averaged for each speed (right) are shown. For the average trajectories in both (a) and (c), the darker shade represents the stance phase while the lighter shade represents the swing phase. **d:** The Mahalanobis distance between the trajectories of each of the speeds and the trajectory of the fastest speed (6.4 km/h).

2.3 Results

2.3.1 Structure of population dynamics varies across different locomotion tasks and speeds

Using the PLDS model, we extracted cyclic, low-dimensional neural trajectories during locomotion. Visualizing in three dimensions, the trajectories follow similar,

33

saddle-like rotations across different tasks, and across different walking speeds (figure 2.3a). The Mahalanobis distance between all of the task trajectories against just the treadmill walking trajectories demonstrates that the corridor walking trajectories are the most similar to treadmill walking (mean distance = 5.0392), followed by bipedal, backwards, and ladder walking in various orders depending on the phase of the gait (mean distance = 9.12, 10.66, and 11.97 respectively, figure 2.3 b). However, any interpretation of the neural trajectories for backwards walking should be treated with caution due to the low number of trials available (table 2.1). When training the PLDS model only on treadmill walking trials, there is a similar rotational structure across all walking speeds, however trajectories appear to separate along the third latent dimension (figure 2.3c). The change in state-space position along this dimension is reflected in the increase in Mahalanobis distance between trials of different speeds as the difference in speed increases (figure 2.3d). In general, the PLDS model was able to infer latent state trajectories that are closely phase-locked to the locomotor rhythm, and distinguishes relevant behavioral parameters such as task type and walking speed. We next describe how well kinematic and gait parameters could be decoded from these low-dimensional dynamics.

2.3.2 PLDS latent state trajectories capture limb kinematics and gait phase

We used several neural features (full population firing rates, PLDS or PCA latent variables, or a predictive sub-sample (PSS) of the population firing rates) as inputs to a Wiener filter decoder and measured how well each input feature could decode various kinematic variables under cross-validation. To determine the total number of dimensions to use in our decoders, we measured the PLDS decoder performance for all the kinematic variables while varying the input dimensionality from 1 to 20. We found that the performance plateaued at approximately 12 dimensions (figure 2.5b);

we used this number of dimensions for the remainder of the analyses, as well as for calculating the Mahalanobis distances in figure 2.3.
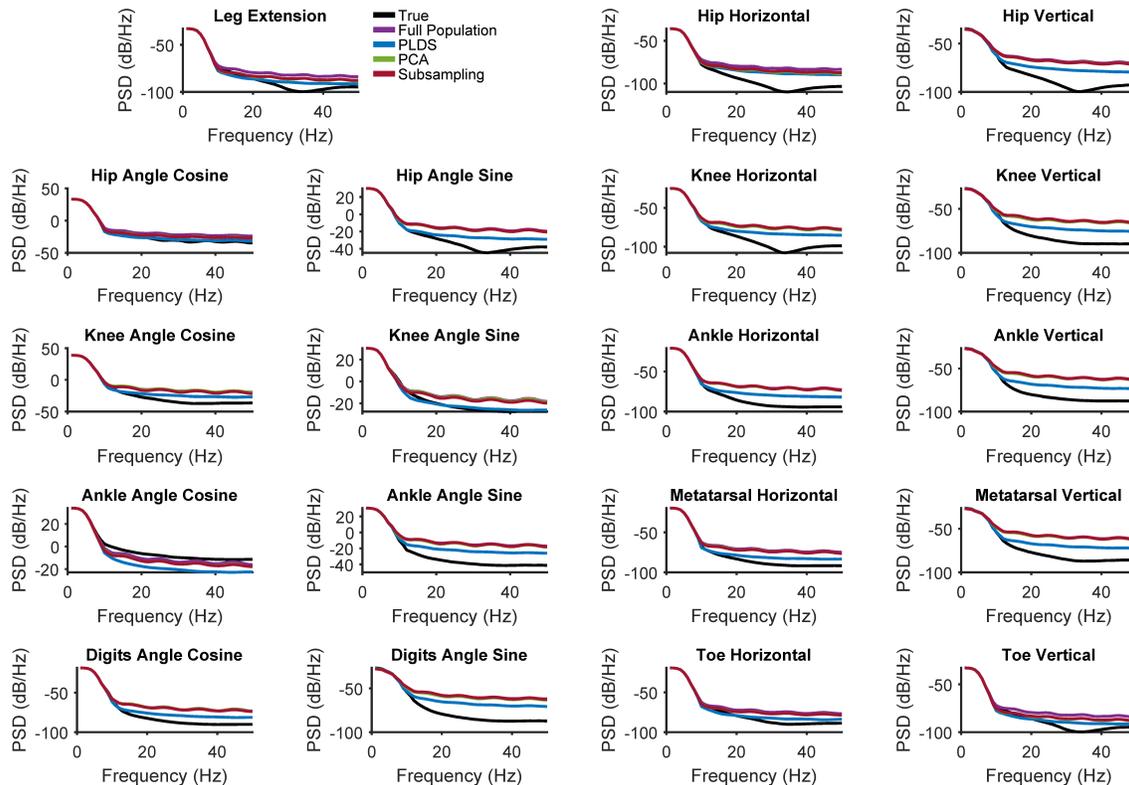
Figure 2.4: PLDS extract smooth latent trajectories. We assessed the "smoothness" of the decoded kinematics by computing the PSD function for each decoder. The power density is shown for the first 50Hz for each kinematic variable here.

First, we investigated whether using PLDS as a feature extractor would improve or deteriorate decoding of hind-limb kinematics. An example trace of the decoded leg extension for one trial is shown in figure 2.5a. For this example, the estimate from the PLDS decoder approximates the true signal more faithfully than the estimate from the Full Population decoder ($R^2 = 0.8294$ vs $R^2 = 0.5461$), as well as compared to the PCA and PSS decoders ($R^2 = 0.5355$ and $R^2 = 0.4980$, respectively). We then compared the cross-validation results for the various decoders across all monkeys and recording sessions that have kinematics (n=7) for each of the kinematic

variables, including the horizontal and vertical position of the joints, the joint angles (taking the sine and cosine to ensure valid R^2 values), and leg extension (two-tailed Wilcoxon signed-rank test, Benjamini-Hochberg false discovery rate correction for multiple comparisons, with FDR=10%). The PLDS decoder performance was not significantly different from the full-population decoder for any of the tested kinematic variables despite having only 12 input features, and indeed for all variables, the average PLDS decoder R^2 was higher than the full population, although none of them reached significance when corrected for multiple testing (figure 2.5d). However, the PCA decoder performed worse than the full population for many of the kinematic variables, as did the PSS decoder (see Table S1 for p-values), leading us to conclude that PCA and PSS features, at the same number of dimensions, do not preserve the kinematic information as completely as those obtained with the PLDS model. From the sample trial in figure 2.5a, the PLDS decoded signal appears much smoother than the other three estimations, which may account for the higher R^2. The power spectrum of the decoder outputs verifies that the PLDS decoder and the real kinematics have less power at high frequencies compared to the full population, PCA, and PSS decoders (Figure 2.4). Although the Wiener filter does provide some degree of smoothing by taking into account the previous 10 time bins in the decoding, the PLDS model smooths the intrinsic trajectories based on the activity of the whole ensemble, which may provide a more accurate estimation of the dynamics.

In addition to the above kinematic variables, we also investigated how well gait information, such as the phase of the gait cycle, could be estimated using each decoder. Trials consisted of a single gait cycle, and the phase of each cycle were standardized to a range of 0% to 100%. The phase was then estimated by each of the decoders, and the change in R^2 from the full population decoder is shown in figure 2.5c for each of the recording sessions (n=9). The decoding performance was similar to the decoding

of kinematics; there was no significant difference in R^2 between the PLDS decoder and full population decoder, and the PCA and PSS decoder performed worse than the full population decoder (see Table S1 for p-values).

In addition to regression, we wanted to establish whether dimensionality reduction with PLDS would improve classification of locomotion state. To test this, we divided each gait cycle into the swing or stance phase (Figure 2.6a). The PLDS classifier performed the best, correctly identifying the swing state 97.3% of the time and the stance state 98.3% of the time, yielding an accuracy of 97.95%, while the Full Population classifier only correctly identified the swing state 93.0% of the time, and the stance state 95.1% of the time, yielding an accuracy of 94.3% (Figure 2.6c). The PCA classifier performed similarly to the Full Population classifier, with an accuracy of 94.71%, while the PSS classifier had the lowest accuracy of 92.93%. These classification results are in agreement with the results from decoding gait phase with regression.

PLDS is advantageous with small training sets

One potential advantage of using PLDS to extract low-dimensional latent states is to mitigate the issue of overfitting. Higher dimensional feature spaces require greater amounts of training samples in order to build decoders that are generalizable to unseen data, and therefore it is advantageous to minimize the feature space dimensionality. To investigate whether PLDS would alleviate overfitting issues, we repeated the cross-validation analysis for decoding kinematics and gait phase using progressively smaller training datasets. Figure 2.7a and 2.7b shows the same decoded leg extension trial as figure 2.5a with varying training set sizes for the Full Population and the PLDS decoder respectively. Clearly, as fewer data points are used to train
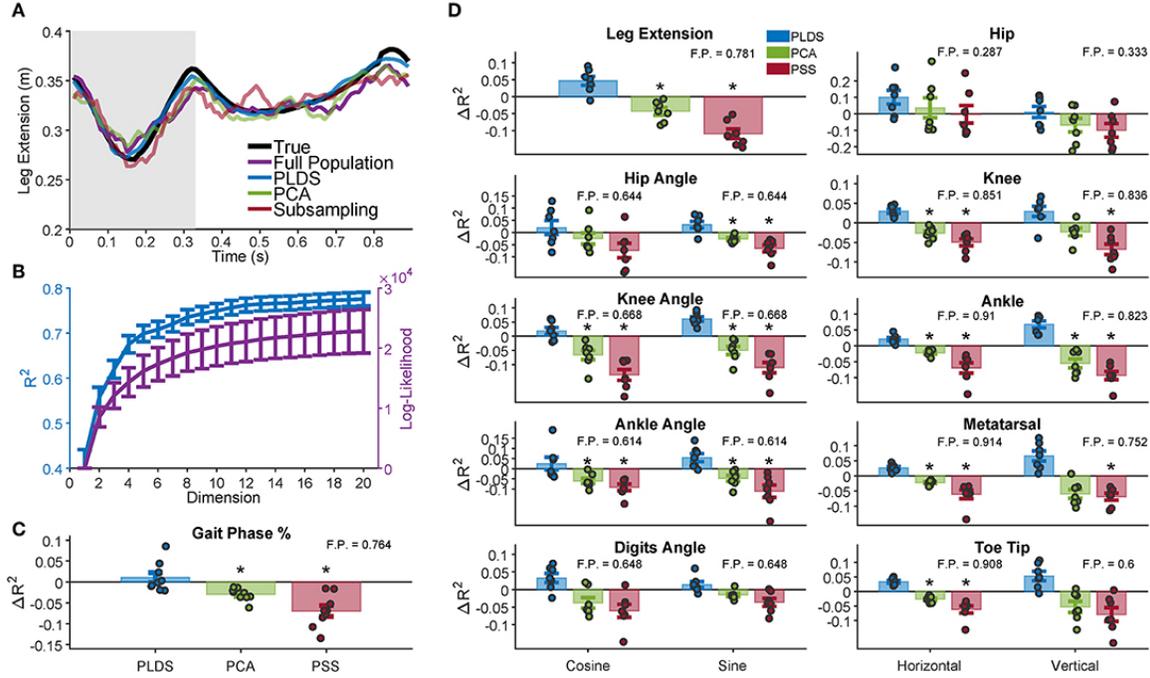
Figure 2.5: Decoding of leg kinematics and gait phase. **a:** A representative example of the limb extension kinematic variable during one gait cycle along with the decoded estimation for the four decoders. Grey background represents the swing period of the gait cycle, white background represents the stance period. In this example, the PLDS decoded signal was the most successful at reconstructing the true signal, with an R^2 of 0.8294, followed by decoding without any dimensionality reduction (Full Population: FP) with an R^2 of 0.5461. PCA and PSS decoding performed the worst, with R^2 of 0.5355 and 0.4980, respectively. The R^2 values for all of the animals and sessions are shown in **d**. **b:** Dimensionality analysis for PLDS reveals a plateau in decoding performance (blue) at approximately 12 dimensions for the latent states (plateau defined as when the increase in R^2 <0.001), error bars 1 SEM. Log-likelihood (purple) also plateaus at around 12 dimensions. **c:** Cross-validated decoding performance of dimensionality reduction techniques for decoding gait phase percentage. Bars represent the increase in R^2 over the Full Population decoder, with the inset number representing the average R^2 of the Full Population decoder. The full distribution of change in R^2s are displayed as dots over each bar. The PLDS decoder had the highest average R^2 for all the variables. Error bars: 1 S.E.M. Black stars denote significant difference from Full Population decoding R^2 (Wilcoxon Sign Rank Test; Benjamini-Hochberg false discovery rate correction for multiple comparisons, with FDR=10%). **d:** Same as in **c**, except for all analyzed kinematic variables rather than gait phase. Each dot represents a animal/recording session, and the bar plots show the average R^2 with error bars representing 1 S.E.M.

the Full Population decoder, the estimated signal quality significantly degrades, while for the PLDS decoder, the degradation is not as egregious, even when only using 10%
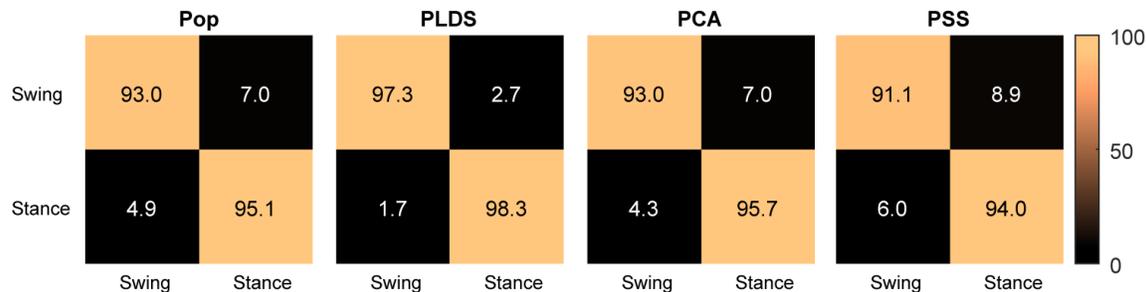
Figure 2.6: Classification of locomotion phases. Confusion matrices for swing vs stance classification. Rows are the actual states, columns are the kNN classified states. Numbers represent the percentage of samples classified in each class.

of the original training dataset. This benefit holds for decoding kinematics or gait at different dimensions. As shown in figure 2.7c, there is a clear drop off in the Full Population R^2s across all kinematic variables (from 0.832 at 90% to 0.574 at 10%), and the gait phase (from 0.533 at 90% to -1.203 at 10%). The PLDS decoder mean R^2s does decrease for all tested dimensions as the training set size is decremented, but the drop off is consistently smaller than that experienced by the Full Population decoder.

PCA and PSS also provided robustness against overfitting. At 12 dimensions, the mean R^2s for kinematics and gait phase of both PCA and PSS decoders remained consistently higher than in the case of the Full Population decoder as the training set size decreased (Figure 2.7d). In general, for kinematics and gait phase decoding, the PLDS decoder outperformed the PCA and PSS decoders, indicating that, although dimensionality reduction through feature extraction does address overfitting issues, explicitly modeling the intrinsic dynamics of neural ensembles with the PLDS model extracts features which are more informative of the hind-limb kinematics and gait with the same low number of dimensions.
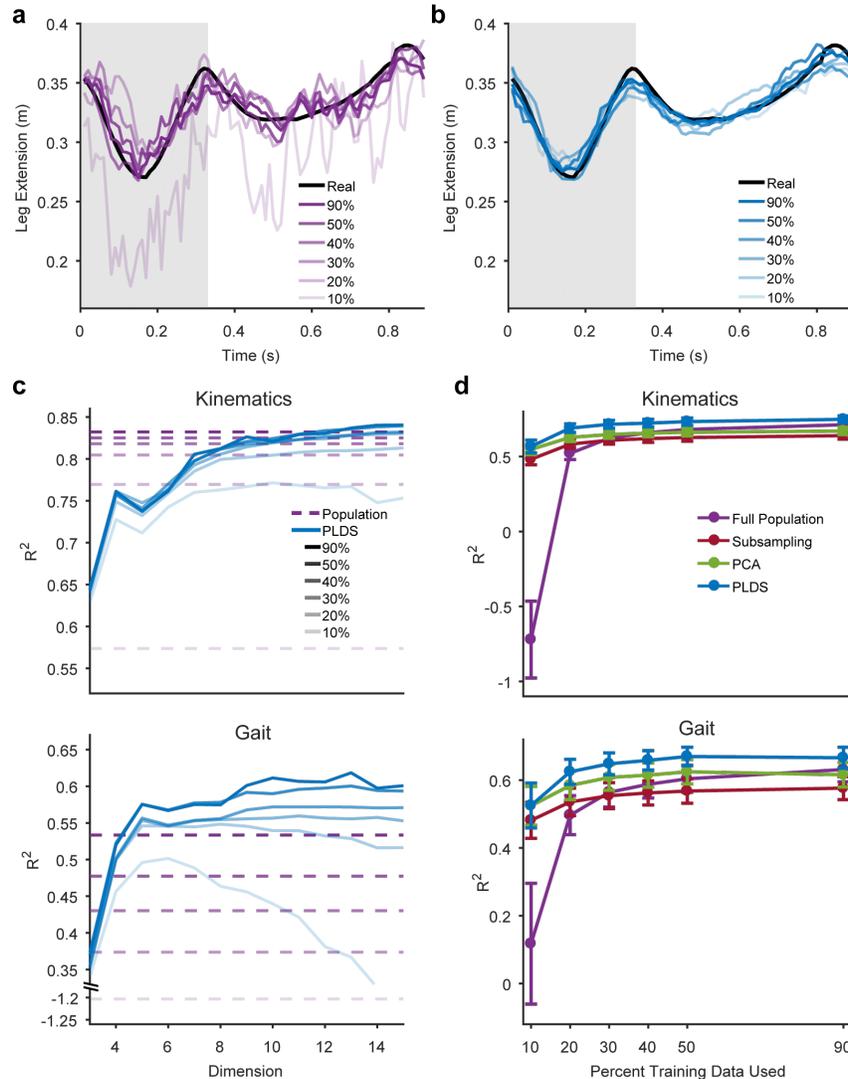
Figure 2.7: Decoding performance results for small training-set analysis. **a:** The same example trace of limb extension for the trial displayed in Figure 2.5a. The decoded traces of the Full Population decoder at various training set sizes are shown. Lighter traces indicates less training data used. Grey background denotes swing phase. **b:** Same as in **a** except the traces are from the PLDS decoder. **c:** Plot of the R^2 values of the session shown in **a** and **b** for different numbers of dimensions used (x-axis), and different training set sizes (opacity). **d:** Decoded R^2 as a function of the percentage of training data used for all three dimensionality reduction and Full Population decoders. 12 dimensions were utilized for the dimensionality reduction techniques. In general, smaller training set sizes resulted in performance drop offs, although this was far more pronounced with Full Population decoding. Error bars: 1 S.E.M. The coefficient of determination could be negative because they were computed on test data.

2.4 Discussion

While the PLDS and PCA models are both dimensionality reduction techniques, a key difference of PLDS is it explicitly models the temporal dynamics using a linear dynamical system. Other techniques have been developed which also models these dynamics such as jPCA (Mark M. Churchland et al., 2012), or Hypothesis-guided dimensionality reduction (HDR) (Lara, J. P. Cunningham, and M. M. Churchland, 2018). These other techniques introduces some constraints to the form of the dynamical system in order to test specific hypotheses about the structure of cortical dynamics. For example, jPCA limits the temporal transition matrix to the set of skew-symmetric matrices in order to extract rotational dynamics, while in (Lara, J. P. Cunningham, and M. M. Churchland, 2018), HDR was used to divide the projected dimensions into linear dynamical systems states and dimensions orthogonal to those states that are condition-invariant (in order to find similarities between two different brain regions). We did not wish to impose any additional constraints about the form the dynamical systems so we used the general form PLDS model for this analysis.

Both the PCA and the PLDS model extracted neural trajectories that were oscillatory in low-dimensional space. PLDS explicitly models temporal dynamics of the latent states with a state transition matrix, resulting in smoother single-trial cyclical trajectories (figure 2.2b, figure 2.4). The structure of the trajectories for treadmill walking at different speeds are in agreement with previous studies, where the neural latent-states separated along one dimension as walking speeds increased while still conserving the rotational structure (Foster et al., 2014). We were also able to obtain qualitatively distinct neural trajectories during different tasks (Figure 2.3a), although again, rotational structure is preserved across all tasks.

The dimensionality of neural data required to represent hind-limb kinematics was

approximately 12 dimensions, when empirically determined as the plateau in kinematic decoding performance (figure 2.5b). Previous studies have estimated the intrinsic dimensionality in forelimb motor cortex during center out tasks to be around 10-20 dimensions (Sadtler et al., 2014; Vargas-Irwin et al., 2015; Yu et al., 2009). Despite the more constrained movements of the hind-limb during locomotion, the dimensionality between leg and arm area of M1 and between reaching and walking are surprisingly similar. Such similarity could suggest comparable levels of cortical involvement during these different behaviors (although the actual structure of the cortical activity could be different). In 1989, Georgopoulos and Grillner proposed the hypothesis that reaching movements in primates may have evolved out of precise gait adjustments during locomotion (Georgopoulos and Grillner, 1989), and others have suggested that similarities in neural and muscle activation onsets between those two types of movements support this hypothesis (Yakovenko and Trevor Drew, 2015). The similarities in neural dimensionality between voluntary reaching and walking would also be consistent with this view. However, due to the low sample size and the use of only one dimensionality reduction technique in this study, further experiments exploring additional models may be necessary before any definitive conclusions can be drawn on the dimensionality of leg-M1 population activity during locomotion.

It is well understood that neurons in motor cortex are correlated to muscle activity and kinematics during locomotion in cats (Trevor Drew, Jiang, and Widajewicz, 2002; Trevor Drew and Marigold, 2015; Irina N. Beloozerova and Mikhail G. Sirota, 1998), rodents (DiGiovanna et al., 2016; Rigosa et al., 2015; Miri et al., 2017; Song et al., 2009), and non-human primates (Foster et al., 2014; Yin et al., 2014; Fitzsimmons et al., 2009). Here, we showed whether the extracted low dimensional dynamics from PLDS preserve the kinematic information that is present in the neural activity. In terms of the decoding performance, PLDS with only 12 dimensions was able to

reconstruct all the kinematic variables in addition to the gait phase just as well as the full population decoder that contains 18-80 dimensions. There is some variance in the distribution of PLDS decoder improvements across animals and sessions (figure 2.5). This could be due to the different populations of neurons that are recorded from in different subjects. These populations are not homogeneous and may represent the true underlying cortical dynamics by varying degrees. Other dimensionality reduction techniques such as PCA or PSS were unable to achieve the same decoding performance at the same number of dimensions, indicating that an explicit model of temporal dynamics of the low-dimensional states, such as PLDS, is crucial for decoding the kinematics accurately. These results reflect the improved decoding performance using linear dynamical systems models during cursor control from arm area (Kao et al., 2015; Yu et al., 2009) and during reaching and grasping behaviors (Aghagolzadeh and Truccolo, 2016), suggesting that low-dimensional dynamics play an important role in both types of movements. Our findings differ in that we did not see a statistically significant increase in performance when using latent state input features, whereas the arm decoding studies did see an improvement. However, given that the R^2 of every single kinematic variable was on average higher with the PLDS decoder compared to the full population decoder, this may be due to our low sample size and lack of statistical power. Additionally, the power is lowered by the use of parametric tests and the large number of kinematic variables tested.

We should also mention that our decoders utilized firing rate inputs, and thus carries the assumption of rate coding rather than temporal coding. Recent findings have suggested that temporal coding may play a larger role in motor control than previously thought (Srivastava et al., 2017). One future extension of this study would be to include models that utilize precise spike timings. Additionally, the Wiener filter decoder we employed is a linear decoder, and although it was able to reconstruct the kinematics fairly accurately, other non-linear models could be used to further im-

prove decoding performance. Finally, we should mention that this study was limited to higher-level control areas such as the motor cortex, though lower-level structures such as brain-stem or spinal cord have been shown to also exhibit intrinsic population dynamics (Bruno, Frost, and Humphries, 2017).

In conclusion, our study investigates whether unsupervised dimensionality reduction can infer latent neural states reflecting ensemble dynamics, while preserving information about the kinematics and gait phase of the hind-limb during various locomotion tasks. We show that dynamical systems models, which have been shown to decode forelimb reaching kinematics, were able to extract robust, stereotyped low-dimensional state-space trajectories, and that these trajectories capture hind-limb movements during directed locomotion (e.g. ladder walking), as well as autonomous locomotion (e.g. basic treadmill and corridor walking). As far as we are aware, this is the first demonstration of explicit state-space models of neural dynamics robustly decoding kinematic and gait information during primate locomotion. These results also points to the potential of using PLDS in hind-limb BMIs, although direct testing in a closed-loop system would be required before any determination of the usefulness of PLDS as a feature extraction step can be made. Additionally, although the PLDS model is highly non-linear in the observations, the state dynamics remains linear. Recently, newer techniques have been developed to extract neural dynamics using recurrent neural networks (Pandarinath et al., 2018a) which enable extraction of non-linear state dynamics and have been employed in arm reaching tasks. Such models may extract the underlying neural dynamics more accurately and may out-perform PLDS in terms of decoding of kinematics. One potential extension to this study in the future would be to apply these non-linear models to hind-limb locomotion behaviors as well.

Chapter 3

Development of a fully-implanted hind-limb brain-machine interface

3.1 Introduction

This chapter details the development of a fully-implanted brain-spinal interface (BSI) system which was the result of a collaboration between Medtronic and our laboratory. The fully-implanted BSI aims to provide lower-limb rehabilitation using newly developed Medtronic technology, and implement the system in a non-human primate (NHP) model for collecting pre-clinical data. An previous study has demonstrated effective rehabilitation after an induced spinal cord injury (T7/T8 hemisection) in rhesus macaques using a closed-loop system driving epidural spinal stimulation with the Medtronic Activa RC device (Capogrosso et al., 2016). However that system relies on additional hardware such as a transcutaneous micro-electrode array pedestal to provide cortical recordings, which may pose challenges for clinical translation. The goal of this project is to improve upon the previous system by leveraging the on-board sensing and processing capabilities as well as the upgraded telemetry module of the new Medtronic RC+S device to create a fully embedded system. This work is the

result of the collaborative efforts of multiple members of the Brown neuromotion lab, including myself, Radu Darie, and Marc Powell.
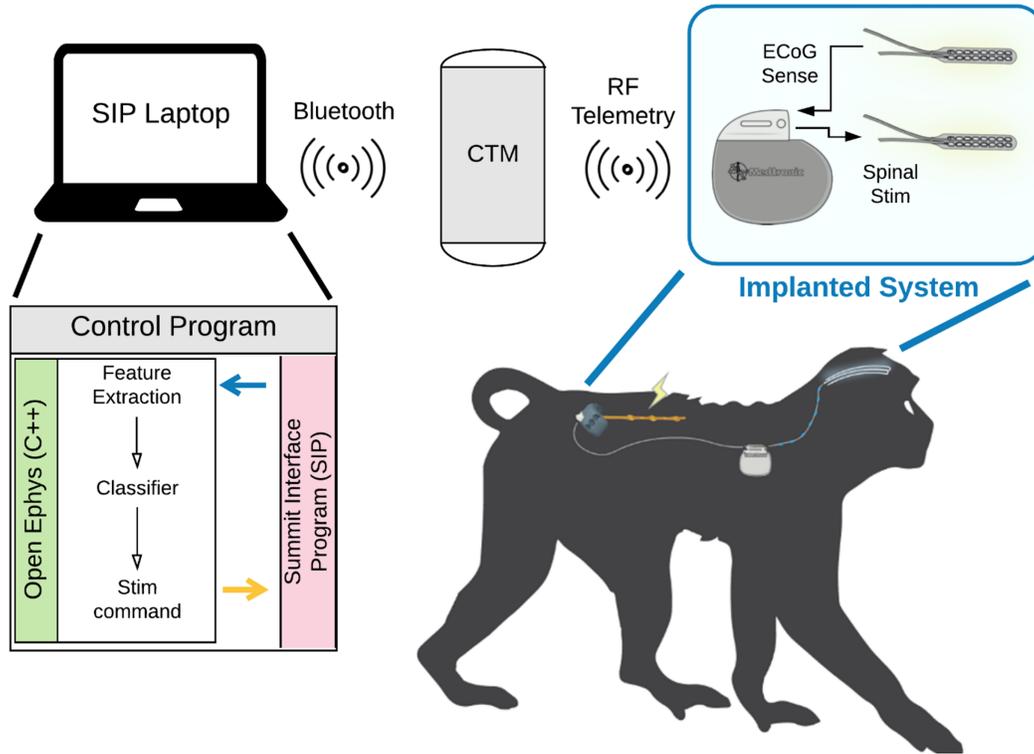
Figure 3.1:   **Full-implanted BSI approach**. Overview of the fully-embedded BSI system.

Figure 3.1 illustrates the design of the fully-embedded BSI. This system utilizes electrocorticography (ECoG) electrodes implanted over the leg area of primary motor cortex (M1). The ECoG signals will serve as inputs to the BMI and are forwarded to a classifier which decides whether the user is about to transition into the swing phase or the stance phase of the gait cycle during walking, much like in the previous system in (Capogrosso et al., 2016). This determination then triggers the start of a stimulation program across selected pairs of electrodes in an epidural spinal electrode array implanted over the lumbar spinal cord. The stimulation programs are chosen and optimized to provide selective activation of hind-limb extensors (to facilitate the foot-strike motion for transitioning into the stance phase), or flexors (to facilitate

foot-lifting motion for transitioning into the swing phase).

We will implement the classifier in two different ways. First, we will leverage the embedded linear discriminant analysis (LDA) processer in the RC+S device to perform decoding of the ECoG signals. However, the amount of experimenter control of the onboard LDA implemenation is limited, and it is unclear at this time whether the limited LDA algorithm is sufficient for providing accurate classification of gait events. We will therefore also implement a computer-in-the loop variation where the ECoG data is streamed from the implanted RC+S to a dedicated decoding computer running potentially more powerful (but more computationally intensive) algorithms. The stimulation triggers are then streamed back to the INS which will then execute the stimulation program.

Development of the system involves:
- Testing and verifying the RC+S sensing and stimulation capabilities.
- Designing and manufacturing ECoG and epidural spinal electrode arrays, implement connectorization schema with RC+S device
- Developing software for interfacing with Medtronic's Summit API
- Developing software for performing neural decoding on ECoG inputs
- Validate closed-loop system in benchtop tests
- Train animal subjects to walk in a treadmill and calmly interact with any necessary elements (e.g. CTM, recharging strap, ect.)
- Implant RC+S, ECoG, and spinal electrode in animal subjects
- Validate closed-loop system in animal subjects

In this chapter, I will focus on my contributions to the development of the system, which mainly comprises of hardware validation and software development for the implementation of the closed-loop system.

3.2 Hardware Validation
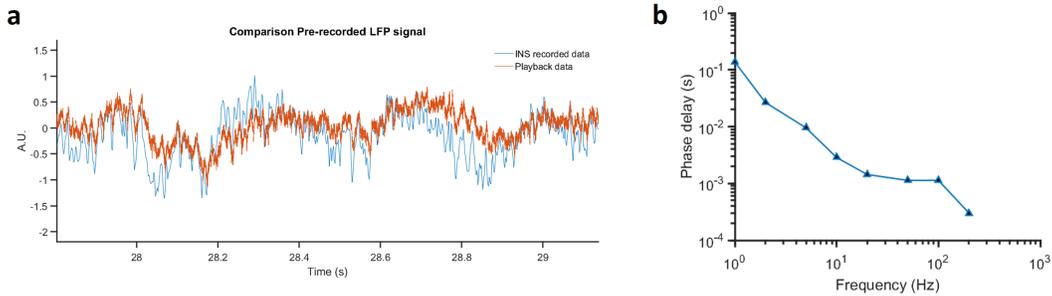
Figure 3.2: **ECoG sense channel recording**. **a)** Playing back previously recorded LFP data into the INS through the sense leads. **b)** Overall group delay from the filters applied on the recorded data.

To test the sensing capabilities of the system, we connected one 1x8 test lead (Medtronic lead kit, 3778-60) to one of the bores of the implanted neural simulator (INS). We fed in a 10mVpp 10Hz, sine wave across two electrodes using a function generator, and we were able to reproduce a sine wave in the saved sense data from the INS. From our testing, we determined we were able to obtain 500Hz sampling rate using 4 sense channels, or 1000 Hz sampling rate using 3 sense channels. Next, we tested how well the INS was able to capture neural signals. We had previously acquired local field potential (LFP) signals from motor cortex using micro-electrode arrays in one of our non-human primate subjects. We played back this LFP signal through the sense lead, and compared the INS recorded signal and the actual LFP signal.

The INS sense channel was able to capture the recorded signal fairly well (figure 3.2a). However, there is a clear phase offset between the actual signal and recorded signal, but only at low frequencies. There is clearly a non-linear phase delay introduced through the filters, and to characterize this phase delay, we recorded sine waves at 1, 2, 5, 10, 20, 50, 100, and 200 Hz (we were unable to get an accurate sine wave

at 500Hz), and measured the time offset. Figure 3.2b shows the group delay for the whole recording process.
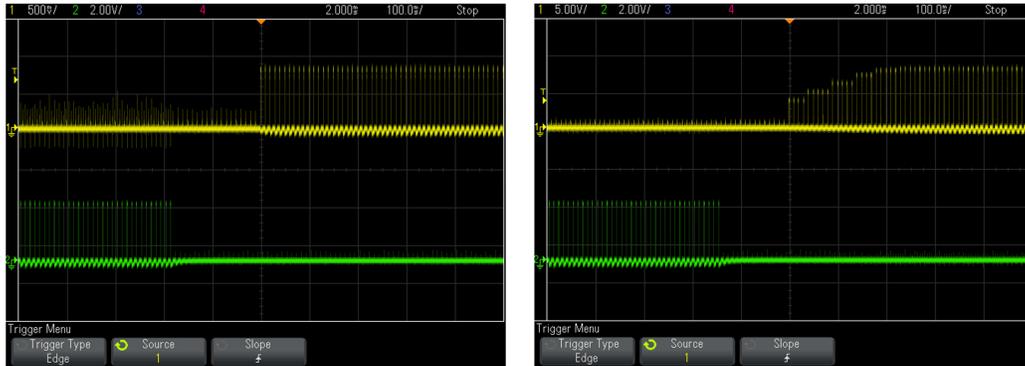
Figure 3.3:   **Spinal stimulation pulses**. Turning on stimulation (yellow trace) or off stimulation (green trace) using either a 1.5kOhm resistor (left) or a 15kOhm resistor (right).

Next, we wanted to test the stimulation capabilities of the RC+S system. Using an oscilloscope, we verified that the pulse width, amplitude, and frequency were indeed what we defined them to be in the stimulation program. We also tested how amplitude change commands are carried out by the INS. Using a 1.5kOhm resistor across the stimulation cathode and anode, we found that the INS was able to adjust the stimulation amplitude within one stim pulse. However, when using a 15kOhm resistor, it takes the INS approximately 200 ms to ramp up the current to the correct level (figure 3.3). These results indicate that we will need to ensure the impedance of our spinal stimulation electrodes are below 15kOhms to maintain high-performance switching of stimulation programs.

Finally, we tested the reliability of the timing of the INS clock relative to other data acquisition devices we are using. Namely, we use video cameras to capture kinematic data of animals walking, and the timings of the foot-strike and foot-lifting events are obtained through these videos. Therefore, we have to synchronize the

49

camera frames with the INS streamed data. Due to imprecision between the clocks of the two devices, there will be some drift which will cause time misalignment the longer the session goes on. We performed a drift test by sending in pulses to both the INS and the DAQ which controls the cameras over two hours, and measured how much the pulses become misaligned. We determined the drift rate to be approximately 43.2ms/hour. The results of these tests indicate that to keep our alignment resolution to be within 10 ms, we should resynchronize the INS with the DAQ approximately every 14 minutes.

3.3 Computer-in-the-loop interfacing and implementation of the BMI control loop

Medtronic provides an Application Programming Interface (API) for streaming data from the implanted device, as well as controlling stimulation from a host PC. First, we wanted to make sure we were able to reliably pull the streamed data from the INS. Using the API's listener design pattern, we implemented a callback function which would write the data received from the INS to a text file. We were able to reconstruct sine waves and LFP signals we sent in though the sense leads (figure 3.2). To determine the reliability of the telemetry system, we characterized how many packets were dropped during streaming by looking at the time-domain packet number associated with each packet received from the INS. If the packet number did not increment by one (or wrap around to 0), this means that packets were dropped. We found that we get a packet drop rate of about 0.162%. We also characterized how consistent the packet size was. We found that at 500Hz sampling rate (which is what we will be using for the BSI system), most packets contain about 25 samples. Occasionally, the number of samples per packet jumps up to 50, and a few times it

went up to 75, as shown in figure 3.4a.
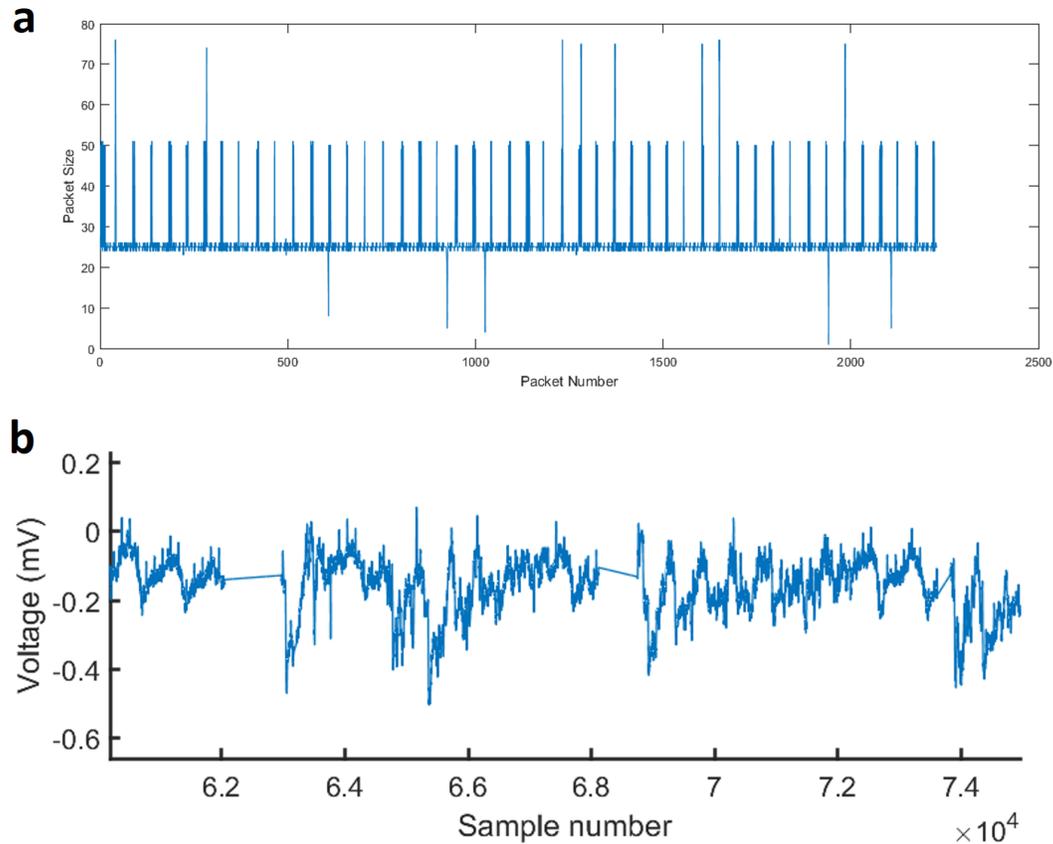
Figure 3.4:   **Data packets from real-time streaming of sense channels**. **a)** Number of samples contained in each time-domain sense packet from the INS. **b)** Recorded data from the INS with dropped packets. Software detected when data was dropped, and automatically filled in missing points with linear interpolation.

Because there are occasional packet drops, we had to implement a method to handle missing data. We decided to use linear interpolation to fill in missing data points. Because the number of samples per packet is not constant, we had to estimate how many points to fill in for the dropped packets. We used information in the packet headers of the current packet and the last packet to determine the approximate number of samples that was lost, n_p. We then inserted n_p points after the last packet using the linear equation:

51

$$d_i = \frac{d_l + d_n}{n_p} * i + d_l, \text{for} i = 1...n_p \qquad (3.1)$$

Where d_l is the value of the last data point in the last packet, and d_n is the value of the first data point of the current packet. We were able to fill in missing data in real time by adding this code to the callback function. We induced packet drops by moving the INS away from the wireless receiver and verified that the linear interpolation was working, see figure 3.4b.

Finally, we realized that the callback functions for handling streamed data packets was asynchronous. Each time a new packet came in, a new thread was spawned. When saving the data to disk, this originally caused issues as new threads would spawn before the writing of the last data packet would finish, causing data-race conditions, and jumbled values. We created thread-safe buffer classes, using the ReaderWriterLockSlim class in the .NET framework to make all data handling thread-safe.

To carry out the decoding of the sense data to detect foot-strike or foot-lift gait events, we decided to use an open-source signal handling and processing platform that was partially developed at Brown, Open-Ephys (http://www.open-ephys.org/gui). There are two main advantages with using Open-Ephys. First, the platform was designed to acquire data from a "source", process the data along a "filter" chain, and finally send out the processed data to "sinks". These "sources", "filters", and "sinks" are implemented as plugins that can be dragged and dropped to create a signal-chain for processing data streams. This architecture is exactly what we need to implement the decoding. Second, developers have already created many plugins which we can leverage for handling the data from the INS. For example, there exists a visualizer plugin, which will graph the data being acquired, in real time (figure 3.7). What we have done is created three plugins for the Open-Ephys system: a "source" plu-

52
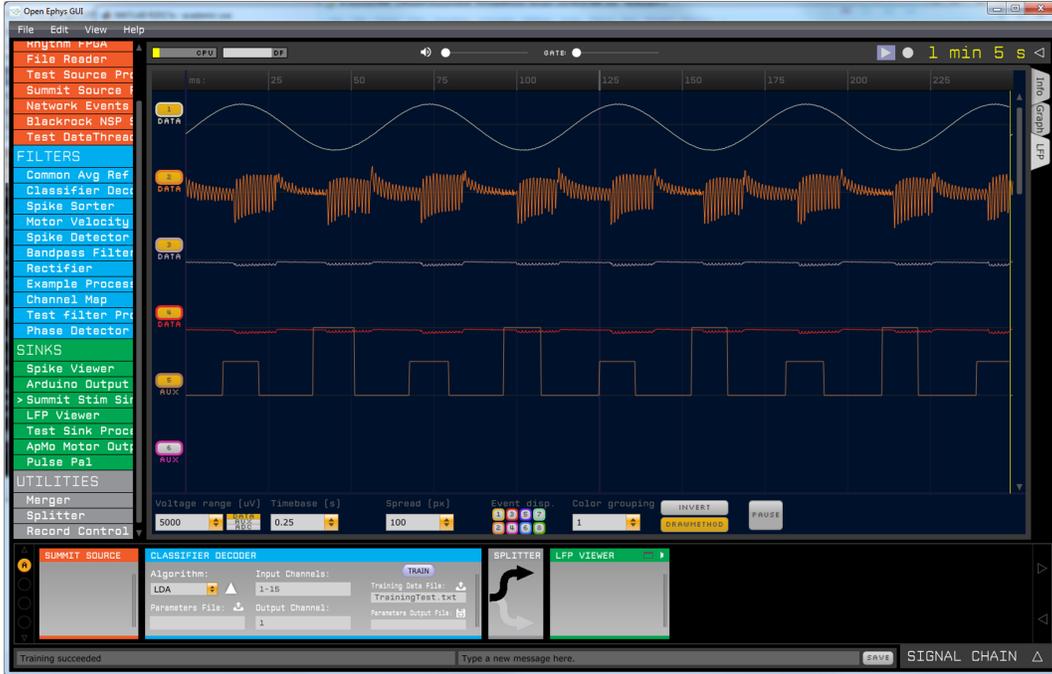
Figure 3.5: **Open-Ephys GUI interfacing with the INS**. Viewer plugin of the Open-ephys program. Top four plots is displaying the data that is being acquired by 4 sense channels of the INS. The first sense channel is connected to a function generator outputting a sine wave, while the remaining three channels are left floating. The last plot is the output of the classifier built to detect peaks and troughs of the first channel.

gin which accepts data from the Summit API, a "filter" plugin which will perform classification of the streamed data to into foot-strike, foot-off, or neither classes, and a "sink" plugin which will stream the classification results back to the Summit API to trigger the appropriate stimulation program. Unfortunately Open-Ephys is developed in C++, while the Summit API uses C, so we can not directly integrate the Summit API functions into an Open-Ephys plugin. We implemented a solution which runs a Summit program in the background, acquiring data from the INS using the listener pattern; this program then sends out the data saved in its pbuffers through a network socket using ZeroMQ (ZMQ, http://zeromq.org/). We serialize the data using our own serialization scheme. We then wrote a Open-Ephys Source plugin which connects to the same ZMQ socket, and deserialize the data, which can then be passed on to the classifier. Similarly, after we finish processing/decoding the data, a

Open-Ephys Sink plugin sends the decoded class through another ZMQ socket back to the Summit API program which will call the necessary functions to carry out stimulation. The Summit API program also handles saving the data from the INS to disk.

To ensure that the Open-Ephys interfacing, INS interfacing, and data-saving components do not block or interfere with each other, the Summit API program spawns a thread that handles each of these operations. The Summit manager and thread-safe data buffers are shared between these threads. Figure 17 illustrates the overall architecture of the Summit-OpenEphys system, and defines the threads that are used and their corresponding loop times.
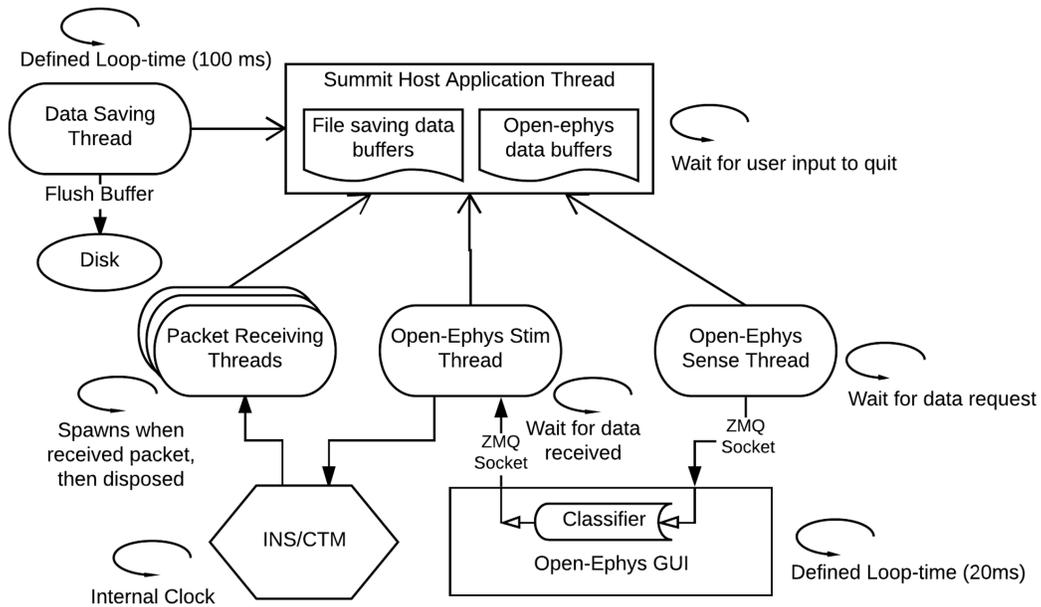
Figure 3.6: **Software threading overview**. Overview of the system for the computer-in-the loop classifier program. The Summit Host Application Thread is the main program and it interfaces with the INS/CTM by sending out stim commands. It also spawns packet handling threads whenever a new data packet is received from the CTM. It spawns three other threads: one dedicated to saving data to disk, one dedicated to streaming data to the Open-Ephys program via ZMQ, and one dedicated to receiving data from Open-Ephys via ZMQ. The Open-Ephys GUI is its own program running the classifier on the streamed INS data.

3.4   Closed-loop bench-top test

Finally, before implanting into an animal, we wanted to perform a system test on the benchtop to validate the functionality of our compete BSI implementation. We need to verify that we were able to stream in a signal from the INS through the sense leads, pass that data to Open-Ephys as inputs to the classifier, and perform the correct stimulation based on the decoded class. The decoder we utilized was a LDA classifier. The optimal LDA weighting coefficients, w, are obtained by calibrating the classifier on a training set using eigenvalue decomposition. The computation of class distances was implemented in an Open-Ephys plugin.

For the benchtop test, we used a well defined input signal, a 20mVpp, 0.2Hz sinusoid. This signal was connected across one sense channel, using electrode 0 of the test lead as the anode and electrode 1 as the cathode. The test lead was connected to the top bore of the INS. We used an additional test lead connected to the bottom bore of the INS as the stimulation leads. We used the Medtronic deep-brain stimulation test leads rather than our ECoG and Spinal stimulation leads so that the condition of the pads could be kept as pristine as possible for implantation, as we had only obtained one pair of leads for each set. The classifier was trained to detect the sinusoid peaks (¿ 8mV), which will serve as a stand-in for foot-strike events, and troughs (¡ -8mV), which will serve as a stand-in for foot-lift events. If a peak was detected, stimulation at 100Hz and 1mA, with 300 us pulse width was initiated across channels 8 and 10 in the stimulation lead, and if a trough was detected, the same stimulation parameters was initiated across channels 13 and 15. These channels was connected to an oscilloscope to visualize the stimulation activations. The input sinusoid was also connected to the scope to see the timing of the input and the resultant stimulation output.

Figure 18 displays the channels on the oscilloscope during the closed-loop test.

The classifier is able to reliably detect the peaks and troughs, as evidenced by the pulses on both stimulation channels. However, there appears to be around a 1.5s delay between peak or trough onset, and stimulation output. The origin of this delay is currently being investigated.

We have additionally implanted the complete BSI system in a rhesus macaque, but unfortunately, before we were able to perform any data collection, an infection developed around the implantation site and the INS had to be explanted. The project is still ongoing, and we will continue to develop and validate the fully-implantable BSI system.
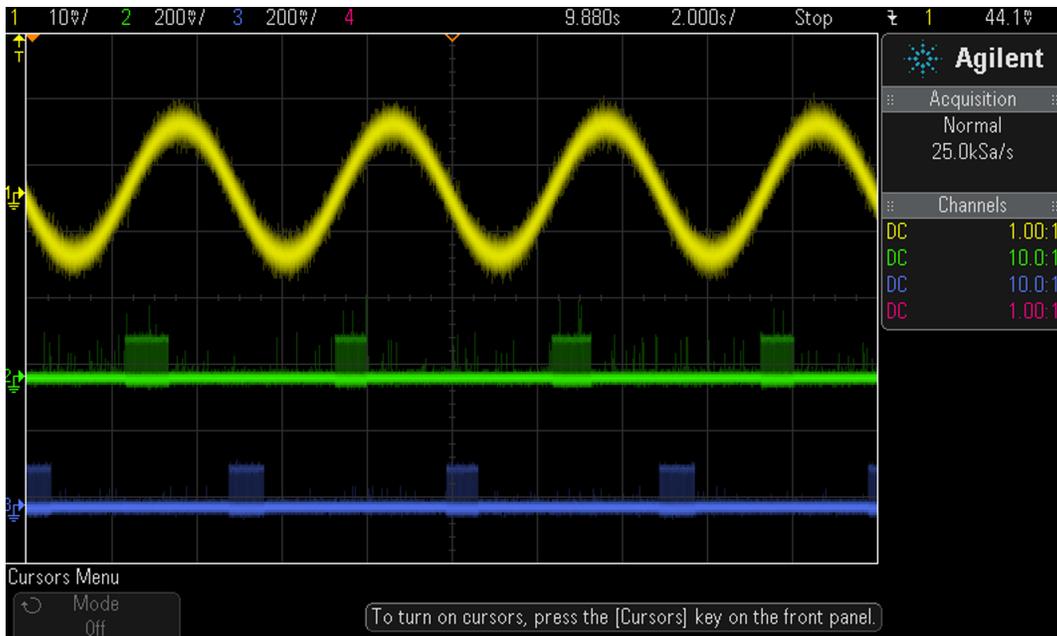
Figure 3.7: **Closed-loop bench-top test**. Closed-loop benchtop test. Top yellow trace displays the sinusoid input to the sense leads of the BSI system. Middle green trace shows voltage across stimulation electrodes activated when troughs are detected, and bottom blue trace shows voltage across stimulation electrodes activated when peaks are detected.

Chapter 4

M1 neural dynamics during obstacle avoidance

Parts of this chapter are being prepared for publication as an original article

4.1 Introduction

Locomotion is an essential component of our motor repertoire, enabling us to gather food and avoid predators for survival. Many animals, including primates, also have the refined ability to skillfully position their limbs to carry out precise movements such as reaching, grasping, and manipulating objects in their environment. Such movements integrate visual cues from the environment(Irina N Beloozerova and Mikhail G Sirota, 2003; Armer et al., 2013; Rivers et al., 2014), continuous sensory feedback from the limbs(Yakovenko and Trevor Drew, 2009), and decision making processes (Raposo, Kaufman, and A. K. Churchland, 2014) from cognitive areas to formulate and execute a motor plan. While the goal of basic locomotion is to simply create forward propulsion, it is complex locomotion, which combines basic locomotion with skillful volitional movements, that allows us to react to irregular or unexpected elements in the environment. During complex locomotion, such as obstacle avoidance,

animals must carry out all the above computations while additionally maintaining the underlying locomotion cycle (T. Drew, 1988; Yakovenko and Trevor Drew, 2015; Trevor Drew, Andujar, et al., 2008; I N Beloozerova and M G Sirota, 1993; Irina N. Beloozerova and Mikhail G. Sirota, 1998; Zelenin et al., 2011).

Figure 4.13a illustrates a few of the neural processes that are involved in successfully carrying out complex locomotion. Motor cortex is essential for carrying out precise adjustments to the movements of the limb while walking (Courtine, 2005; Lemon et al., 2012; Irina N. Beloozerova and Mikhail G. Sirota, 1998) and must integrate top down control of the muscles with incoming afferent information from downstream areas (Irina N. Beloozerova and Mikhail G. Sirota, 1998; Arshavsky et al., 1972). While the question of how primary motor cortex (M1) generates volitional reaching and grasping movements from a stationary position has been the focus of much research (Georgopoulos and Grillner, 1989; Mark M. Churchland et al., 2012; Pandarinath et al., 2018b; Lara, J. P. Cunningham, and M. M. Churchland, 2018), how M1 integrates the necessary computations to produce the correct volitional movement while maintaining locomotion is still poorly understood, especially in primates.

Previous studies in felines have demonstrated that subsets of M1 neurons are phasically modulated with the gait cycle during locomotion(D M Armstrong and Drew, 1984a; T. Drew, 1988; Pearson and Gramlich, 2010; Krouchev and Drew, 2013; Trevor Drew, Jiang, and Widajewicz, 2002; I N Beloozerova and M G Sirota, 1993; Irina N. Beloozerova and Mikhail G. Sirota, 1998; Irina N Beloozerova and Mikhail G Sirota, 2003; Prilutsky et al., 2005), and that many experience a shift in activity when the animal must step over an obstacle(Trevor Drew, Andujar, et al., 2008; Friel, Drew, and Martin, 2007). A recent study identified neurons whose activity consistently coincides with the onset of hind-limb muscle activation during both basic locomotion
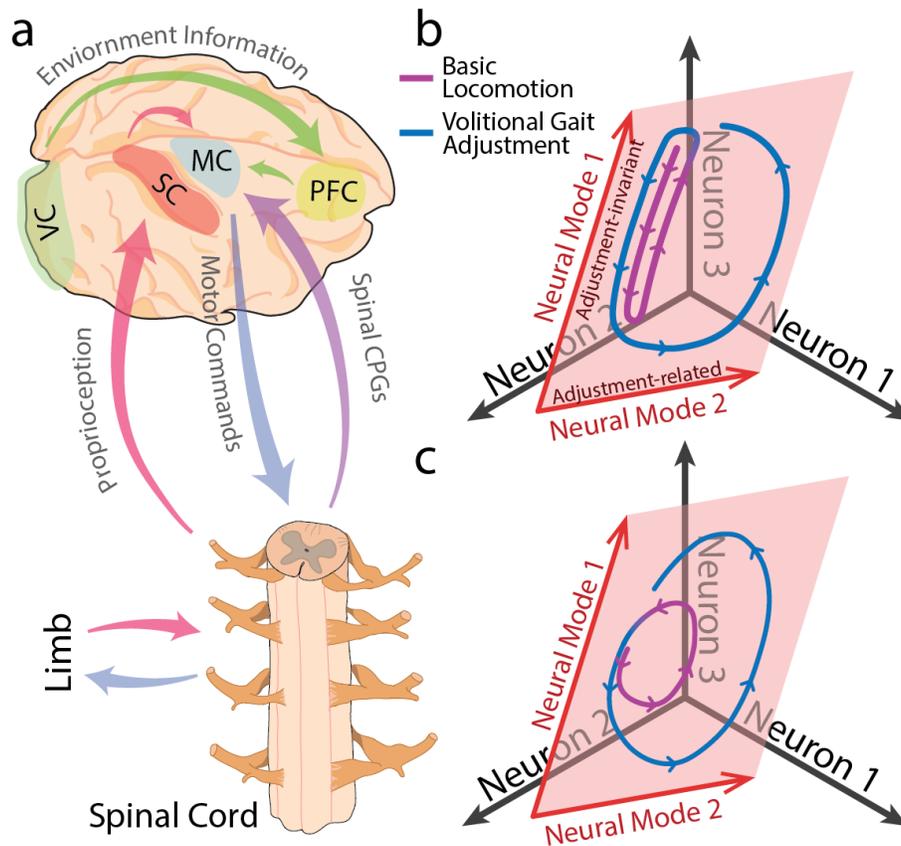
Figure 4.13: **Neural modes related to volitional adjustments during locomotion a)** Example pathways of some of the converging inputs onto motor cortex during volitional gait adjustments. Proprioception and spinal CPGs provide information about the limb position in the gait cycle while visual information from the environment both drives decision making to react to complex environments as well as facilitates the calculation of the precise movement needed to successfully navigate the terrain. Motor cortex (or perhaps upstream areas such as parietal cortex) must integrate all of these components to generate a motor command, which is relayed through the spinal cord to the limb. MC - motor cortex, SC - somatosensory cortex, VC - visual cortex, PFC - prefrontal cortex. **b,c)** Two possible strategies that motor cortex could employ to carry out volitional gait adjustments. The full three dimensional space represents all possible combinations of firing rates of three toy neurons. The neural activity is confined to a two-dimensional subspace (red plane) spanned by two neural modes (red arrows). The curve represents the time-varying neural activity during one stride of basic locomotion (purple) followed by a stride with a volitional gait adjustment (blue). In **b**, the neural activity during basic locomotion is mostly confined to the first neural mode, while the second neural mode encodes the movement modifications. In **c**, both neural modes are utilized during basic locomotion and both are modified by motor cortex during the volitional movement.

and obstacle walking(Yakovenko and Trevor Drew, 2015). Unfortunately, most studies to date are limited to characterizing the properties of each neuron individually. In recent years, the dynamical systems view of movement generation in motor cortical areas has emerged as an effective way to explain coordinated activity of populations of neurons recorded simultaneously(Gallego, Perich, L. E. Miller, et al., 2017; Sadtler et al., 2014; Gamaleldin F Elsayed and John P Cunningham, 2017; Williams et al., 2018; Athalye et al., 2017; Stavisky et al., 2017; Russo et al., 2018). In this view, the role of neurons is not solely to encode any particular movement covariate, such as kinematics, kinetics, or muscle activity, but to act as elements of a coordinated network carrying out the necessary computations for successful goal-oriented movement, such as sensorimotor transformations or gait integration.

Prior research has shown that much of the coordinated activity of M1 neurons live in a subspace of much lower dimension than the space of all possible neural states (Mark M. Churchland et al., 2012; Kao et al., 2015; Gallego, Perich, Naufel, et al., 2018; Xing et al., 2019). They have also shown that the bases that span this subspace, which we refer to as neural modes, may be responsible for different computations (e.g. figure 4.13b). For example, in delayed reaching tasks, cortical activity during movement preparation and movement execution appear to lie in different subspaces, and therefore utilize different sets of neural modes(Gamaleldin F. Elsayed et al., 2016). A separate study has found that the generation of a wide range of different movements result from a superposition of a set neural modes common to all movements (task-invariant modes) and a set of neural modes which captures the variability between movements (task-dependent modes)(Gallego, Perich, Naufel, et al., 2018). Along the vein of locomotion, Miri and colleagues have found that in mice the M1 neurons occupy different subspaces during basic locomotion compared to during a volitional movement (lever pressing from a stationary position)(Miri et al., 2017). However,

they did not investigate how the underlying neural modes of M1 integrate the precise volitional movements with ongoing locomotion. Here, we wanted to ask whether the same neural modes underlying basic, unobstructed locomotion are being utilized by motor cortex to carry out a volitional gait adjustment, or whether cortex engages separate neural modes to carry out the adjustment. In the case of the former, we would expect to see modulation of all of the neural modes that span the low-dimensional subspace in response to a volitional change (e.g. figure 4.13c), but in the case of the latter, we would expect the change in neural activity be contained to a specific set of neural modes, while a different set maintains the same activity as during basic locomotion (e.g. figure 4.13b).

To investigate this question, We recorded neural activity from leg area of M1 in behaving animals while they performed basic locomotion on a treadmill as well as stepping over an incoming obstacle. We found a shift in neural activity directly preceding the increase in step height during the stride placing the limb over the obstacle, suggesting the presence of an efferent signal in M1 inducing the volitional gait adjustment. Using dimensionality reduction, we found that most of the variance related to this efferent signal appear to be contained to a distinct set of neural modes, while other neural modes are unaffected by the gait modifying movements. Therefore, M1 appears to engage two distinct subspaces, one for maintaining the locomotor rhythm, and one for encoding the momentary increase in M1 engagement during the gait modification.
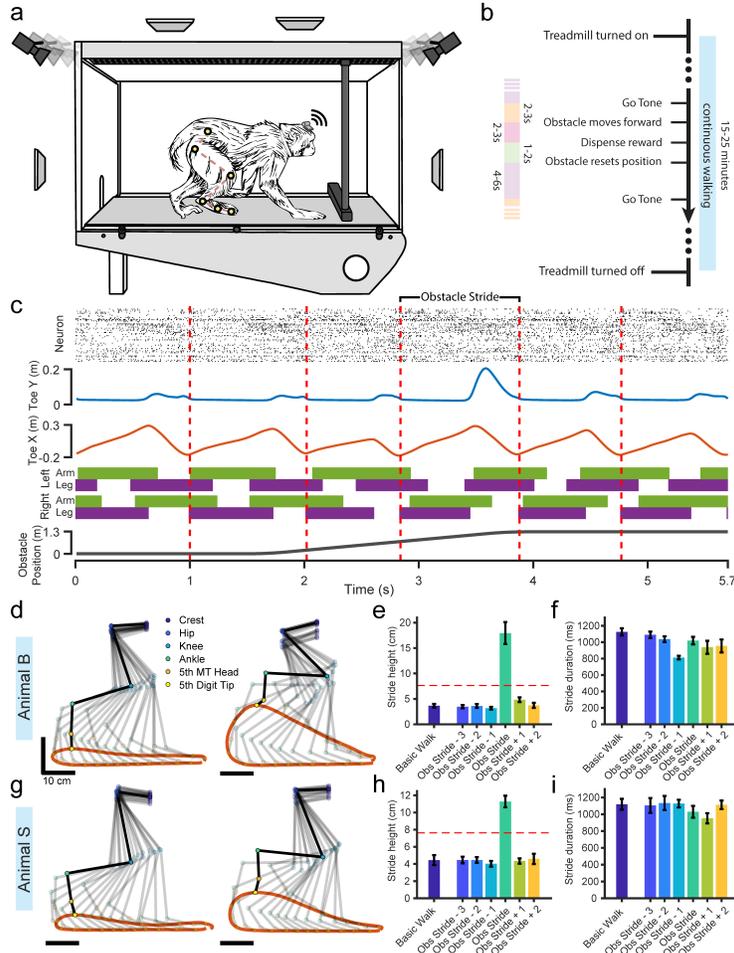
Figure 4.14: **Experimental paradigm and behavior. a)** Treadmill enclosure and behavioral apparatus where animals performed locomotion and obstacle stepping. Antennas surrounding the enclosure wirelessly collected neural data while video cameras recorded the positions of six joint markers painted on the hind limb. **b)** Behavioral paradigm. Each trial consisted of a single obstacle run. **c)** Example data from a single trial of Animal B, showing the stride bringing the leg over the obstacle along with three strides before the obstacle stride and two strides after the obstacle stride. Top: raster plot of single units from the implanted Leg-M1 MEA. Blue trace: height of the right toe tip. Orange trace: horizontal position of the right toe tip, normalized to the right illiac crest. Second to bottom: gait pattern across all four limbs. Solid bars represents contact with the ground (stance phase); purple - hind-limbs, green - forelimbs. Bottom: Obstacle horizontal position. **d,g)** Stick diagram of the right hind limb during one stride of basic, unobstructed walking (left) or during the stride stepping over the obstacle (right). Orange trace represents the trajectory of the toe tip. Stick figures are spaced 60ms apart, with the dark stick figure highlighting the limb during the maximum height of the toe tip. **e,h)** Maximum height reached by the toe tip for either continuous unobstructed (basic) walking or the stride over the obstacle and the surrounding strides. Error bars represent standard deviation. Dotted red line indicates the height of the obstacle. **f,i)** Duration of each stride, error bars are standard deviation. **d-f**: Animal B, **g-i**: Animal S.

4.2 Results

4.2.1 Obstacle avoidance paradigm

In order to interrogate the cortical response to volitional gait adjustments during locomotion, we trained two rhesus macaques to carry out an obstacle avoidance task. Animals were placed inside a plexiglass enclosure on top of a motorized treadmill and continuously walked at 2.2km/h. Within the enclosure was an actuated Styrofoam bar that we were able to move towards the animal at 2.2 km/h, giving the perception of an obstacle approaching along the treadmill belt (Figure 4.14a). Animals were trained to raise their limbs to step over the obstacle without touching it while maintaining their ongoing locomotion. After the obstacle reached the end of the treadmill, a motor would rotate the bar out of the path of the animal and it would be brought back to its starting position to initiate the next trial (Figure 4.14b). A "go" tone would alert the animal of the oncoming obstacle, and the operator would manually time the start of the obstacle movement to a particular gait event (e.g. the moment the right hand made contact with the floor) in order to maintain consistency in the obstacle avoidance movement across trials. Although there was some jitter in the timing of the obstacle initiation relative to the gait cycle, the variance was rather small (standard deviation: 2.7% of the gait cycle for animal B, and 2.8% for animal S). Each trial consisted of one obstacle run which included the stride where the leg was lifted over the obstacle, three strides before the obstacle stride, and two strides after the obstacle stride (figure 4.14c). We will denote the stride where the limb is lifted over the obstacle as Stride 0, and the surrounding strides by their relation to the obstacle stride (e.g. stride -3 refers to three strides before the obstacle stride). Our dataset contains a total of 38 walking obstacle avoidance trials from animal B and 43 trials from animal S. Additionally, we collected data from the animals as they continuously walked on the treadmill without any obstacle movement, which we will

refer to as basic walking trials. In this case, each trial consisted of a single stride, and we obtained a total of 49 basic walking trials from animal B and 66 trials from animal S.

We implanted 96-channel micro-electrode arrays (Blackrock Microsystems, Utah) in the leg area of left M1 (figure 4.19) and were able to isolate 50 neurons from animal B and 42 neurons from animal S. We also recorded the kinematics of six joints (crest, hip, knee, ankle, 5th knuckle and 5th toe tip) of the contralateral (right) hindlimb using painted markers and video capture (figure 4.14a). Example neural activity and kinematics of an obstacle stepping trial, along with the gait pattern and obstacle position are shown in figure 4.14c.

4.2.2 The movement pattern of stepping movements is preserved during obstacle avoidance while the amplitude is increased

The limb movement throughout all the strides, including the obstacle stride, follows the same stepping pattern. Horizontal and vertical positions of the joints increased and decreased at similar phases of the gait cycle. The main difference between the obstacle stride and the other strides is the magnitude of the movements. Both animals had to raise their hindlimbs considerably higher than during their normal walking gait in order to clear the obstacle (figure 4.14d-e, g-h). During basic walking, the average stride height was 3.66cm for animal B and 4.45cm for animal S, and increased to 17.95cm and 11.28cm respectively. However, the stride durations were similar for all the strides, with the exception of the stride immediately before the obstacle stride in animal B, which was shorter in duration than usual (figure 4.14f, i). In this animal, the gait cycle lasted on average 1.12s during basic walking
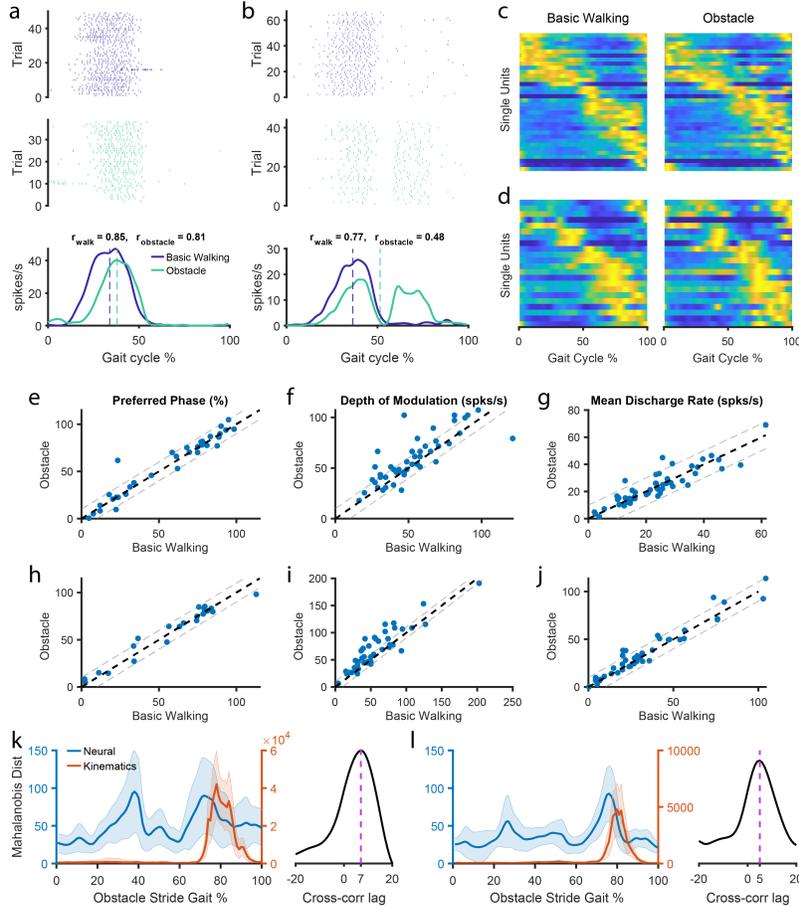
Figure 4.15: **Most M1 neurons are phasically tuned to the gait cycle and increase their depth of modulation during obstacle stepping. a)** Example raster plots and PETHs from a M1-leg neuron from Animal B. Spike times were normalized to the gait cycle duration for each trial (see methods). Top raster plot is for basic walking trials while bottom raster plot is for the obstacle stepping trials. r values represent the dispersion of the neural activity around the gait cycle, and the dotted lines represent the preferred phases of the neuron. **b)** Same as in **a** but for an example neuron from Animal S. **c)** Activity of all phasically modulated neurons for Animal B during both basic walking (left) and the stride stepping over the obstacle (right). Neurons are sorted by ascending preferred phase of the basic walking trials, for both basic walking and obstacle stepping plots. The activity of each neuron is normalized to its maximum firing rate. **d)** Same as in **c**, but for Animal S. **e,f,g)** Changes in preferred phase, depth of modulation, and average firing rate respectively between the basic walking stride and the obstacle stride. Each point represents an individual neuron. Thick dotted line indicates no change, thin dotted lines delineate a change in 10% for preferred phase and 10 spks/s for depth of modulation and average firing rate. **h-j)** Same as in **e-g** but for Animal S. **k)** Mahalanobis distance between the stride over the obstacle and the stride before any obstacle movement (three strides before the obstacle stride) for the population of neural firing rates (blue) or the kinematic variables (orange). Right plot shows the cross-correlation between the neural and kinematic distances across multiple gait percentage lags with the dotted line indicating the peak lag. Positive indicates neural lagging kinematics. **l)** Same as in **k** but for Animal S. Using average stride durations, 7% of the gait cycle for Animal B corresponds to ms, while 5% of the gait cycle for Animal S corresponds to ms.

110

and decreased to 0.822s in the stride before the obstacle, although the obstacle stride increased back to 1.03s. The stride duty factor was similar for both animals, 67% for animal B, and 69% for animal S. In both animals, the right limb was the trailing limb over the obstacle.

The obstacle did not start moving until two strides before the obstacle stride in animal b and one stride before the obstacle stride in animal S, meaning that from the perspective of the animal, the first gait cycle in the obstacle avoidance trials should be essentially the same as during basic walking. Indeed, the stride height and duration are virtually identical between these two strides (figure 4.14e-f, g-h). Between the next two strides, the animals could see and were aware of the obstacle moving towards them, although the stride height remained unchanged for these strides. After clearing the obstacle, the stride height returns to normal pre-obstacle ranges. In summary, despite the large change in kinematics in the obstacle stride, the subjects returned to normal walking quickly after avoiding the obstacle, and with the exception of animal B taking a smaller and quicker stride right before the obstacle stride, they did not drastically alter their gait leading up to the obstacle stride.

4.2.3 No gross reorganization of M1 neurons during walking obstacle avoidance.

We then looked at the response properties of individual neurons in leg area of M1. Figure 4.15a, b illustrates the activity of two example neurons during both basic walking (purple) and the the stride over the obstacle (teal). In order to compare the activity across trials, which had some slight variations in duration, the spike trains were time-normalized to 0-100% of the gait cycle (see methods, figure 4.20a,b). Consistent with previous M1 recordings during locomotion, the majority of neurons

tended to increase their firing rate during specific phases of the gait cycle. Although many of these neurons fired at around a single phase of the gait cycle (e.g. figure 4.15a, figure 4.20c), we also found neurons that had multimodal peri-event time histograms (PETHs). We utilized circular statistics on the unimodal PETHs to determine the preferred phase and dispersion (r) of the neurons around the gait cycle (figure 4.20d). We applied the Rayleigh test to identify neurons which were not significantly modulated to the gait cycle and further classified unimodal modulated neurons as weakly modulated if they had a r value less than 0.15 during basic walking. Figure 4.20e shows the distribution of dispersion values and demonstrate that most of our recorded neurons are strongly modulated. For animal B, 7/50 neurons were multimodal, 3/50 neurons were not phasically modulated, and 11/50 neurons were weakly modulated, for animal S, 6/42 neurons were multimodal, 3/42 were not phasically modulated, and 17/42 were weakly modulated. Since preferred phase is only well defined for neurons that are both significantly modulated to the gait cycle and have a unimodal PETH, we excluded neurons that are multimodal or weakly modulated from preferred phase calculations (but not depth of modulation or mean firing rate calculations).

We found that for both animals, the preferred phases of the population of recorded neurons spanned the whole gait cycle (figure 4.15c, d), consistent with previous studies in felines and macaques (I. N. Beloozerova and M G Sirota, 1988; Trevor Drew, Jiang, and Widajewicz, 2002; Fitzsimmons et al., 2009; Xing et al., 2019). Additionally, neurons tend to fire at the same phase during the obstacle avoidance step as during basic walking (Figure 4.15c-e,h). Only 4/34 (animal B) and 2/22 (animal S) of the neurons had shifts in preferred phase greater than 10% of the gait cycle. However, most neurons tended to increase their depth of modulation during the obstacle step (average increase of 6.37 spks/s for Animal B, 11.62 spks/s for Animal S), consistent with the larger amplitude of the movement (Figure 4.15f,i). The mean firing rates
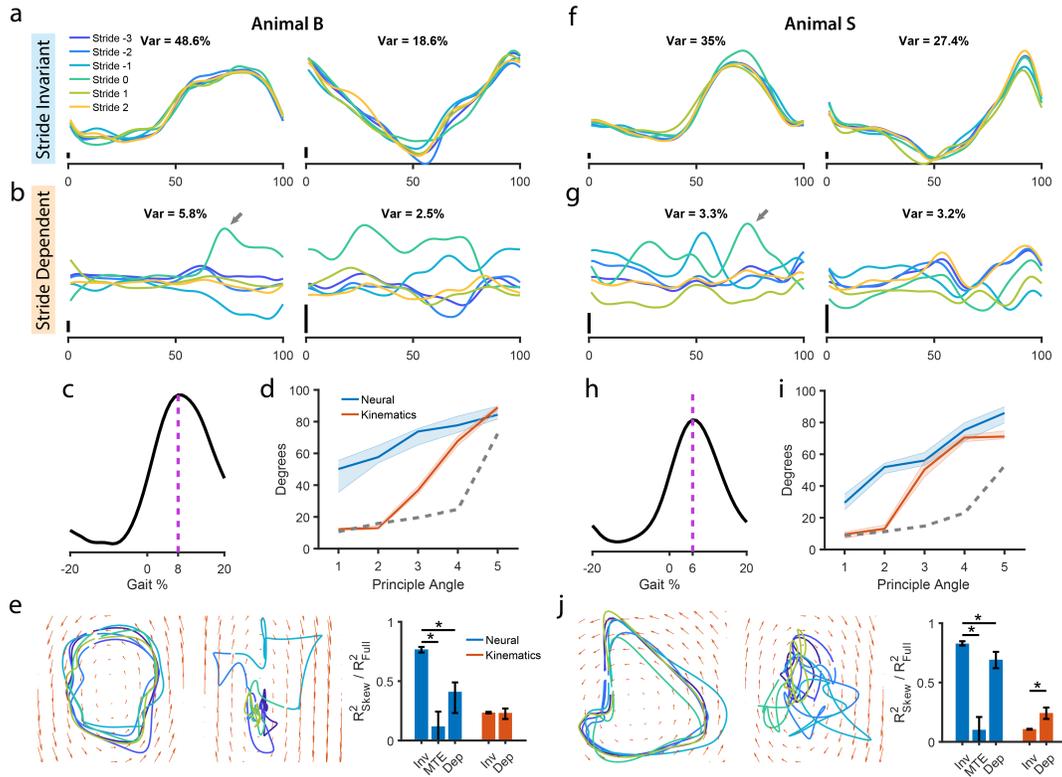
Figure 4.16: **dPCA separates neural activity into obstacle stride invariant and stride dependent subspaces a-b, f-g)** Top two stride-invariant (**a, f**) and stride-dependent (**b, g**) demixed components for Animal B. Each of the strides surrounding the obstacle stride (Stride 0) are plotted individually. Numbers at the top represent percentages of the total variance for each component. Arrows indicate the increase in the stride-dependent neural modes during the obstacle stride that account for the peaks in **c** and **h**. Vertical bars indicate 0.1. **c, h)** cross-correlation between the change in kinematics and the change in activity within the stride-dependent subspace across multiple gait lags, similar to figure 4.15**k, l**. **d, i)** Principle angles between the stride-invariant subspace and stride-dependent subspace for both neural dPCA and kinematics dPCA. Error bars represent 95% confidence intervals for 500 bootstrap resamples. Dotted grey line represents the 97.5th percentile of the null distribution of completely overlapping subspaces. **e, j)** First two stride-invariant components (right plot) or stride-dependent components (center plot) are plotted against each other to visualize rotational structure (or lack thereof). The slope fields for the rotations inferred with jPCA are shown as orange arrows. Left plot: ratio of the jPCA model R^2 to an unconstrained LDS model R^2 to quantify the strength of rotations. For dPCA used on neural activity, maximum tensor entory (MTE) was used to generate surrogate datasets and the R^2 ratios was also calculated for the stride invariant subspace of the surrogates. Error bars are 95% confidence intervals for 500 bootstrap reshuffles. Stars indicate statistically significant difference (500 label shuffles, $\alpha = 0.05$).

also saw an overall increase, although the change is not as large (average increase of 0.52 spks/s for Animal B, 2.25 spks/s for Animal S, Figure 4.15g,j).

The increase in depth of modulation of the neurons is not surprising given the drastic increase in the step height during the obstacle stride. Lifting the leg over the obstacle involves a conscious decision to move the limbs in a specific way, which requires descending input from motor cortex. We therefore hypothesized that there would be a change in neural activity preceding the modifying movement to the gait. We determined when the changes in the obstacle avoidance stride occurred by calculating the Mahalanobis distance between the obstacle stride and the stride before any obstacle movement (three strides before the obstacle stride) at each percentage of the gait phase. Unsurprisingly, we saw a large increase in the Mahalanobis distance of the kinematics during the the swing phase (figure 4.15k,l orange plot). We also observed a bump in the Mahalanobis distance of the neural population immediately before the kinematic increase (figure 4.15k,l blue plot), confirming the presence of an efferent signal in M1 related to the lifting of the leg over the obstacle. Cross-correlation analysis determined that the neural modulation preceded the kinematics change by 7% of the gait cycle in animal B and 5% in animal S, which, using the average stride durations, corresponds to 72.1ms and 52.3ms respectively.

Although the summary statistics of the neural activity suggests that there is not any large change in the activity of the neurons besides an overall increase in the firing rates corresponding to an increase in step height, the response profiles for individual neurons can be quite complex. For example, the neuron in Figure 4.21f decreased its firing rate during its activation in the stance phase, but a second peak emerged during the swing phase when stepping over the obstacle. Figures 4.21b and 4.21d also demonstrates neurons with complex changes in firing activity. How can we interpret

these complicated response profiles? One strategy is to not consider each neuron individually, but as a part of latent neural modes underlying the activity of the whole population.

4.2.4 dPCA reveal division of neural modes into obstacle-related and obstacle-invariant subspaces

We have previously shown that the population activity can be accounted for by a low number of neural modes during locomotion. We wanted to investigate whether M1 modifies these neural modes during the obstacle step or whether it employs a separate set of modes to account for the increase in volition. We therefore wanted to see if we could find a subspace which is invariant to the need for descending input and preserves it's time-varying activity throughout all the strides during obstacle stepping.

One dimensionality reduction technique, demixed principle component analysis (dPCA), decomposes neural activity into subspaces associated with specific task parameters (Kobak et al., 2016). We used dPCA to extract subspaces that were unrelated to the gait adjustment (stride-invariant) and subspaces which captured the change in population activity during the gait adjustment (stride-dependent). Despite these additional constraints, the 10 dPCA components were able to explain 87.2% and 88.4% of the variance for animal B and animal S respectively.

The stride-invariant subspace accounted for most of the neural variance (77.1% for animal B, 80.7% for animal S), and indeed showed very little variation across strides (figure 4.16a, e), while the stride-dependent activity only account for 11.9% of the variance in animal B and 11.0% for animal S. We observed large deviations in the stride-dependent modes during the stride immediately before and the stride

over the obstacle (figure 4.16b,f), suggesting that these modes capture an increase in M1 engagement during those strides. Additionally, when we compared the timing of these shifts in relation to the changes in the kinematics during swing phase, and found that, like with the population firing rates, the activity of the stride-dependent neural modes preceded the gait changes, (figure 4.16c,h).

The nature of the dPCA algorithm means that although it is able to find subspaces that are invariant and subspaces that are dependent on the specified parameters (gait modification in our case), it does not guarantee that the stride-invariant and stride-dependent subspaces are orthogonal. For example, we were also able to find stride-invariant and stride-dependent subspaces for the kinematics, although these subspaces could be overlapping. Principle angles have been used in previous studies to determine the alignment between two subspaces. We calculated the principle angles between the stride-invariant and stride-dependent neural subspaces, and found that they were significantly greater than the bootstrap null distribution (figure 4.16d,i). The principle angles between the kinematics subspaces were much smaller, indicating greater overlap between the stride-invariant and stride-dependent components.

Finally, we wanted to determine if there were any rotational dynamics present in these neural modes. We hypothesized that since the locomotor rhythm was a common factor across all strides, the stride-invariant subspace would contain rotational dynamics, while the stride-dependent subspace, which represents the precisely timed motor adjustments on the gait, would not. We used jPCA to fit a rotational dynamical system to both sets of neural modes, and quantified the rotational strength as the ratio of the $R2$ of the jPCA model to the $R2$ of an unconstrained dynamical system model (Mark M. Churchland et al., 2012). Figure 4.16e,j illustrates the fitted rotational dynamics of the first two neural modes in both subspaces. In agreement
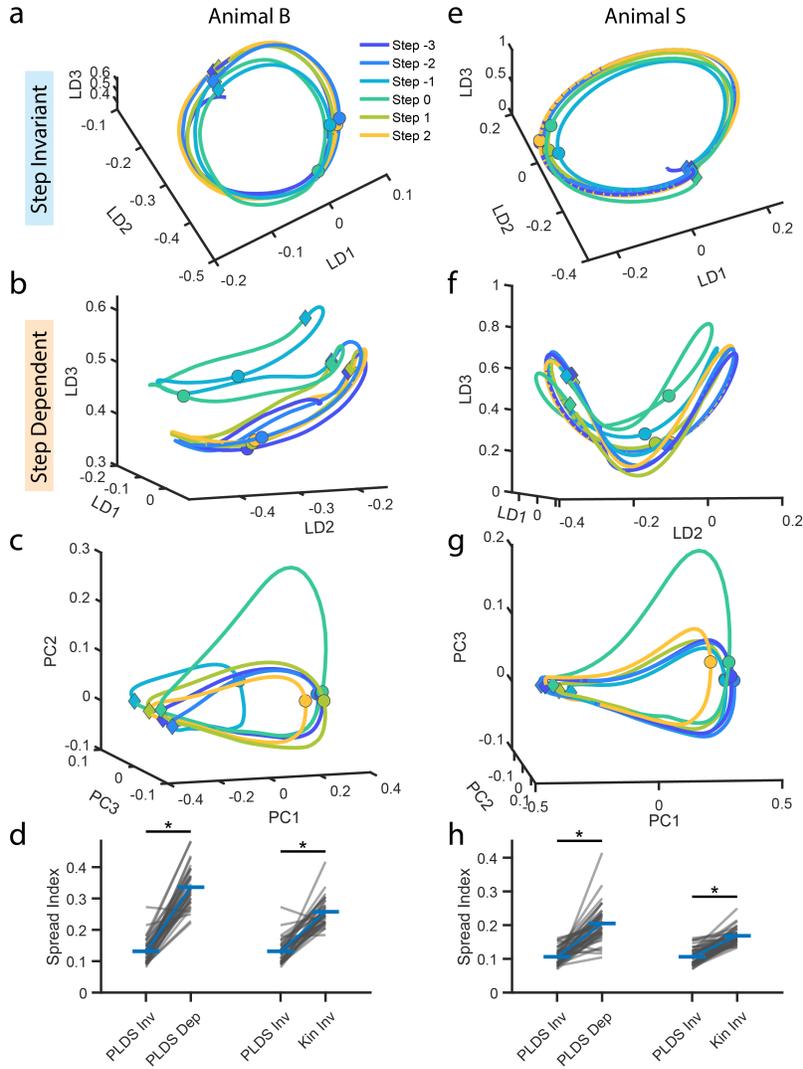
Figure 4.17: **PLDS also extracts neural modes invariant to obstacle stepping. a,e)** Neural trajectories from three neural modes, or latent dimensions (LD), inferred from the PLDS model. Each of the strides surrounding the stride over the obstacle (Stride 0) are plotted individually. The view angle of the trajectories is chosen to maximize the overlap between all of the strides according to a spread index (see methods). The projection into this view represents the stride invariant subspace. Circular dots indicate the transition from the stance phase to the swing phase in each stride, while diamond dots indicate the transition from the swing phase to the stance phase of the next stride. **b,f)** Same plot as in **a** and **e** but with a view angle which highlights the differences across strides, representing a stride dependent subspace. **c,g)** Top three principle components of the kinematics for the same strides as the PLDS plots. **d,h)** Spread index quantifying how much the trajectories overlap or diverge across strides (see methods). We compared the amount of divergence between the trajectories in the stride invariant neural subspaces with the stride dependent neural subspaces (left) as well as the stride invariant neural subspaces with the stride invariant kinematic subspaces (right). Stars indicate statistical significance (Wilcoxon signed rank test, $\alpha = 0.05$). **a-d**: Animal B, **e-h**: Animal S

with our hypothesis, we found that there was significantly more rotational structure in the stride-invariant than the stride-dependent subspace. As an additional control, we generated surrogate datasets using maximum tensor entropy (MTE) and found little rotational structure in these surrogates (figure 4.16e, j bar graphs). Finally, although we chose to use a dimensionality of 10 in our dPCA model (5 dimension for stride-invariant, 5 for stride-dependent subspaces) based on previous dimensionality reduction studies, we obtained similar results across a range of dimensions (supp figure 4b-h).

4.2.5 Unsupervised dimensionality reduction show similar separation of neural subspaces

dPCA is a powerful technique which uses labeled data to extract task-related subspaces. To ensure that our results are not simply the result of the supervised nature of dPCA, we also employed an unsupervised dimensionality reduction model agnostic to the stride type. One such model, Poisson linear dynamical systems (PLDS), explicitly infers the dynamics through a time-evolution matrix, resulting in smooth single trial neural trajectories without the need for trial averaging (figure 4.17a-b,e-f) (Aghagolzadeh and Truccolo, 2014; Xing et al., 2019; Macke et al., 2011). Using the Akaike information criterion, we estimated the dimensionality of the M1 population to be approximately nine (figure 4.23a). These nine neural modes were able to reconstruct the spiking activity of the recorded neurons fairly accurately (figure 4.23b,c), although they accounted for less variance of the neural population than the top nine dimensions of PCA (figure 4.23d,e). However, we have previously shown that the latent neural modes inferred from PLDS were able to predict the movement kinematics just as well as the full population firing rates and better than PCA components during locomotion and indeed this was still the case for our obstacle avoidance data (figure

4.23f,g), indicating that PLDS is an appropriate model for extracting behaviorally relevant latent states from our recorded neural population.

The activity in the first three latent dimensions for each of the strides are shown in in figure 4.17a-b for animal B and e-f for animal S. The same three modes are shown in both the stride-invariant and stride-dependent plots, just from different angles (which is in itself just a projection, or linear combination, of the 3 neural modes into a 2-dimensional subspace). We defined the stride-invariant view as the projection which results in the greatest amount of overlap across all of the strides. We quantified the amount of overlap using a spread metric (see methods), and optimized for the projection which minimized this spread metric. Figure 4.17a and e illustrate these optimal overlap trajectories. There are two notable features in these projections that are consistent in both animals. First is the clear rotational structure in this subspace, corroborating our results from the dPCA analysis. Second is the high amount of invariance between the trajectories across the different strides, despite the large change in movement during the obstacle stride. We emphasize that unlike dPCA, the PLDS model is not designed to specifically identify neural modes invariant to any particular variable, so it is entirely possible that the modulation in response the the voluntary gait adjustment pervades all the PLDS neural modes (e.g. as in figure 4.13c). Indeed, the top three principle components of the kinematics demonstrate this case, as there are no projections resulting in the same amount of overlap between the obstacle stride and the unobstructed strides (figure 4.17c, g). When comparing smallest achievable spread index for the kinematics and the PLDS trajectories, the kinematics was larger by about 2.13 times for animal B and 1.58 times for animal S (figure 4.17d, h). Displaying the PLDS trajectories along a different angle reveals the modulation of the latent activity in response to the voluntary intervention (figure 4.17b, f). Like in dPCA, the trajectories appears to diverge away from the basic
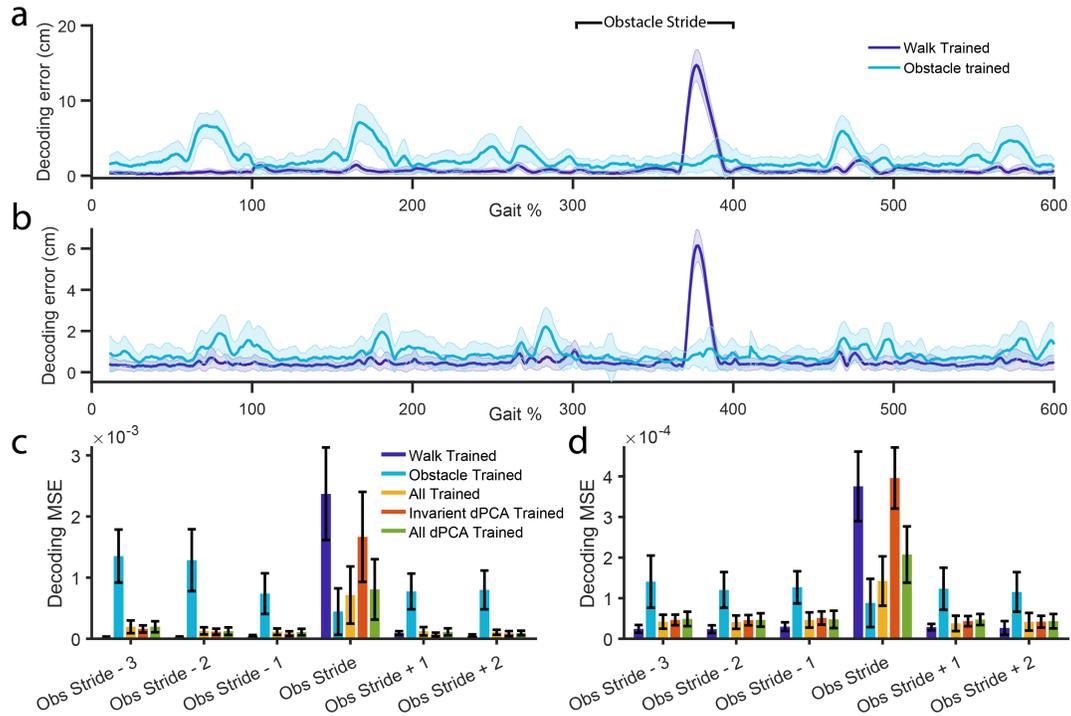
Figure 4.18: **Decoding does not generalize across basic walking and obstacle stepping. a,b)** Error in decoding of the toe tip height for decoders trained only on the data from the stride before any obstacle movement (three strides before the obstacle stride, shown in purple) or decoders trained only on data from the stride stepping over the obstacle (shown in blue). 300-400% represents the stride over the obstacle, error bars indicate standard deviation. **c,d)** Mean squared error for decoding toe height in each of the strides surrounding the obstacle stride, trained on different sets of training data. *Walk Trained* and *Obstacle Trained* are what is shown in **a** and **b**. *All Trained* indicates a decoder trained on data from all of the strides (tested under cross validation). *Invariant dPCA Trained* indicates a decoder trained on data from all of the strides, but only using the stride invariant dPCA components. *All dPCA Trained* indicates a decoder trained on data from all of the strides, and using both the stride invariant and stride dependent dPCA components. **a** and **c**: Animal B, **b** and **d**: Animal S.

walking activity (purple and dark blue traces) during the obstacle step and the preceding step (light blue and teal traces), before returning to the pre-obstacle region (green and orange traces). These results indicate that M1 maintains a consistent set of rotating neural modes throughout obstacle avoidance stepping while employing a separate set of modes to encode the actual gait adjustment movements.

4.2.6 Decoding of movement kinematics does not generalize between basic locomotion and obstacle stepping

The existence of neural modes which seem to represent the efferent intervention of M1 onto the locomotor movements suggests that any decoders trained on neural data during basic locomotion (and therefore not requiring any intervention) would not capture the information represented in these modes. Therefore, we hypothesized that decoders would not be able to generalize across strides during obstacle stepping. To test this, we employed a Wiener filter to try to decode toe height, since that was the most pertinent kinematic variable to clearing the obstacle. We trained the decoders with neural and kinematic data from only the stride before any obstacle movement (walk trained) or the stride over the obstacle (obstacle trained) and measured the decoding performance throughout all the strides. When using a decoder trained on one stride to decode the kinematics within that stride, we employed leave-one-out cross validation. We found that while the walk-trained decoder was able to reconstruct the kinematics fairly accurately for the strides before and after the obstacle stride, there was a large amount of error when decoding the swing phase of the obstacle stride (figure 4.18a, b, blue). Similarly, the obstacle trained decoder was able to accurately decode the toe height during the obstacle stride, but was unable to do so with the surrounding strides (purple).

Finally, to directly test the necessity of the stride-dependent neural modes for capturing the neural-kinematic relationship during the obstacle stride, we trained a decoder on just neural modes from the stride-invariant subspace of dPCA, as well as a decoder trained on all neural modes, including those from the stride-dependent subspace. We trained the decoder using the activity from all the strides and performed leave-one-out cross validation. We found that while the decoder trained on just the stride-invariant neural modes performed well on the strides before and after the obstacle stride, it performed poorly on the obstacle stride itself, despite having trained on obstacle stride data. When we included the stride-dependent modes, there was no effect on the surrounding strides, but increased the decoding performance during the obstacle stride (figure 4.18c, d), suggesting that much of the movement information involved in the gait adjustment is contained in these neural modes.

4.3   Discussion

Using our obstacle avoidance paradigm, we were able train non-human primates to consistently perform a volitional gait adjustment during continuous locomotion. In general, animals maintained their normal walking movements until the obstacle stride, during which the only deviation was an increase in the amount of flexion around the joints to lift the leg higher. Our recordings from leg-M1 neurons during basic locomotion exhibit properties that are consistent with earlier studies, such as phasically tuned neurons, and the activation times of the recorded population spanning the whole range of the gait cycle (I. N. Beloozerova and M G Sirota, 1988; Trevor Drew, Jiang, and Widajewicz, 2002; Fitzsimmons et al., 2009; Xing et al., 2019). Additionally, the overall increase in discharge frequency in our recorded neurons during the stride over the obstacle is in agreement with previous recordings of

felines performing obstacle avoidance (T. Drew, 1988; Yakovenko and Trevor Drew, 2015). We also identified a peak in the neural changes immediately prior to the changes in the kinematics, indicating the existence of an efferent motor component involved in raising the legs over the obstacle.

In subprimate mammals, lesion studies have demonstrated that motor cortex is not necessary for the generation of walking movements (G raham Brown, 1911; Sten Grillner, Georgopoulos, and Jordan, 1997), and the control of basic unobstructed locomotion is thought to be managed by subcortical and spinal circuits (McCrea and Rybak, 2008; Gerasimenko et al., 2006). However, conscious volitional movements, such as reaching, do require cortical input, so during complex locomotion, the motor cortex is faced with a unique challenge of integrating movements generated intrinsically with movements generated in extrinsic areas.

While it is not as clearly established in non-human primates as in subprimates that motor cortex merely plays a peripheral role during basic locomotion, there have been some studies suggesting that volitional movements and basic locomotion requires different amounts of cortical engagement. Macaques were still able to walk within days after a lesion to the corticospinal tract, albeit with some dragging of the paw, but their ability to carry out fine foot movements were almost completely abolished (Courtine, 2005). Similarly, Kuypers et al showed that after a pyramidotomy, macaques could still walk and climb up cages, but were unable to reach for food (Lemon et al., 2012). These findings suggest that even in non-human primates, regions outside of motor cortex are responsible for, at least partially, generating the movements necessary for locomotion. Therefore, the motor cortex of our animals carrying out obstacle avoidance must be able to generate the necessary adaptive movements, while taking into account the movements being generated by spinal circuits.

One strategy that could be employed is to send an efference copy of the motor commands generated by the spinal circuits up to cortex (e.g. figure 1,a purple arrow), allowing it to generate the volitional movements within the proper context of the ongoing locomotor movements. Indeed, it has been proposed that the cyclic activity observed in motor cortex during basic locomotion is a copy of the spinal CPG activity, conveyed through the spinocerebellar tract to the dentate nucleus and routed through the ventrolateral thalamus before arriving at motor cortex (Arshavsky et al., 1972; Irina N. Beloozerova and Mikhail G. Sirota, 1998). Miri et al have found that this rhythmic activity in cortex during locomotion lies in the null space of volitional movements, such as lever pressing, and therefore may serve an ancillary rather than direct role to movement generation (Miri et al., 2017). This ancillary role could potentially be to inform cortex of the state of the limb within the gait cycle, should the need for an adaptive movement arise.

The existence of stride-invariant neural modes are in agreement with this hypothesis. The activity within these neural modes remain consistent throughout all gait cycles, even when the kinematics become drastically different during the obstacle stride. These neural modes also exhibit strong rotational dynamics, consistent with the idea that they are copies of spinal CPGs. Additionally, these neural modes were sufficient for decoding the toe kinematics during basic locomotion, but fails during the volitional gait adjustment, suggesting that they do not carry information involving the efferent signal for the adjusting movement. However, further experiments will be needed to truly determine the origin of these signals in non-human primates.

Meanwhile, the variance related to gait adapting movement appear to be contained to a separate set of neural modes. The integration of the visual information along

with the calculation of the precise movement required to clear the obstacle could be subserved by these neural modes. We emphasize that this division of neural processes would have been difficult to recognize when analyzing the single neuron properties of M1. It is the dynamical systems framework that allows us identify the structure underlying the activity of a cortical population (Gallego, Perich, L. E. Miller, et al., 2017; Pandarinath et al., 2018a; Mark M. Churchland et al., 2012).

Finally, our findings have implications for the development of brain-machine interfaces aiming to restore hind-limb functionality for patient with motor deficits. In recent years, there has been the development of closed-loop systems for restoring walking ability after spinal cord injury(**weg**; Capogrosso et al., 2016; Donati et al., 2016). These systems utilize electrophysiology recordings from cortex to drive either spinal stimulation or movement of an exoskeleton, and have demonstrated remarkable success in allowing subjects to walk again. These systems currently only aim to restore basic locomotion, and not precise directed leg movements. Our decoding results indicate that future BMIs aiming to restore a wide range of hind-limb movements should include both basic locomotion as well as volitional movements during decoder calibration in order to achieve optimal performance.
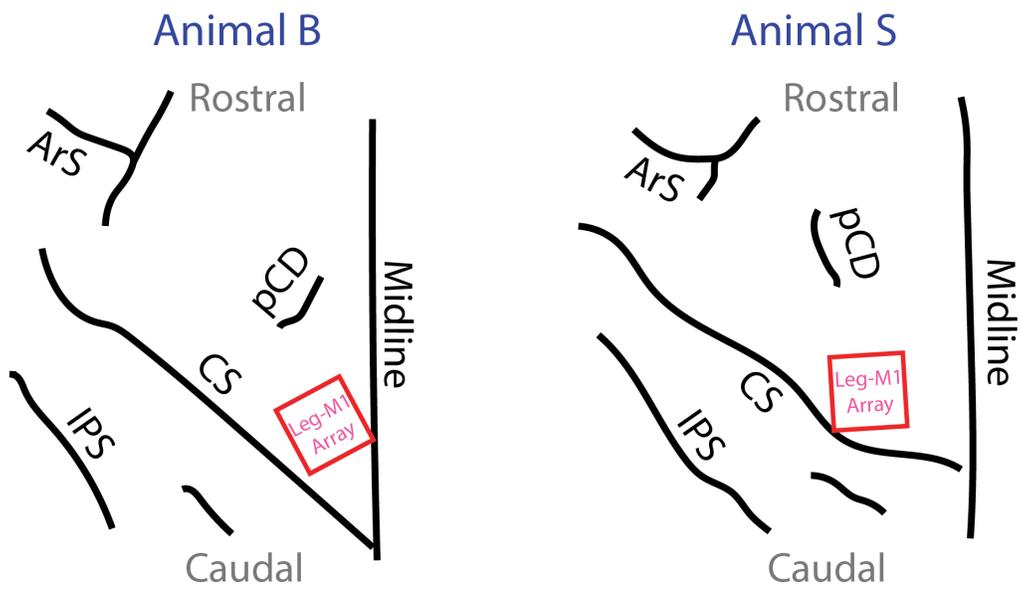
Figure 4.19:   **Microelectrode array implant locations.** Red rectangles indicate where the Utah arrays were implanted in leg-area of M1 in the left hemisphere. *ArS*: Arcuate Sulcus, *pCD*: pre-Central Dimple, *CS*: Central Sulcus, *IPS*: Intraparietal Sulcus.

4.4 Methods

4.4.1 Animal husbandry

All experimental and surgical procedures were approved by the Animal Care and Use committee at Brown university. Two male rhesus macaques, aged 7 and 9 years old were housed in individual cages and trained to perform the obstacle avoidance walking task. Positive reinforcement was used and food and water were provided ad lib throughout training.

4.4.2 Experimental setup

All behavioral tasks were carrying out inside a treadmill enclosure. The treadmill was purchased commercially (JogADog, MI), and a custom plexiglass box (177.8cm long x 47.6cm wide x 91.4cm tall) was constructed above it (figure 4.14a). Animals were able to move freely inside the enclosure and were not tethered in any way. To encourage consistency of movements across trials and also to protect the obstacle components of the apparatus, plexiglass walls were placed below the ceiling and in front of the animal, removing 43.8cm of the top of the enclosure and 72.4m of the front of the enclosure from the available space the animals were able to move in (not shown in figure 4.14, shown in supp movie 1).

A 5.08cm high by 4.45cm wide by 42.86cm long rectangular Styrofoam bar served as the obstacle. The bar was attached to a stepper motor which rotated the obstacle into and out of the path of the animal. The motor was then attached to a belt linear actuator (Igus, RI) which moved the obstacle back and forth along the length of the treadmill. The whole obstacle apparatus was attached to the ceiling of the treadmill enclosure such that the top of the obstacle bar was 7.62cm above the treadmill floor.

Additionally, a speaker was placed in the ceiling to play audio tones, and a small slot in the front of the enclosure allowed food rewards to be placed on the treadmill belt and carried to the animal. Treadmill speed, obstacle speed, and timing of the audio tones were measured using a hall effect angular position sensor, stepper quadrature encoder, and microphone, respectively.

Eight cameras were positioned around the enclosure and captured video of the animals performing the tasks at 100Hz (SIMI Reality Motion Systems GmbH, Germany). Camera calibration was performed after each recording session to determine the positions and angles of each camera relative to each other, allowing for 3D triangulation of any markers that appear in at least two cameras. UV floodlights were placed around the recording room to enhance visibility of our UV reactive joint markers (see kinematics). The cameras were synchronized to each other and to the neural data with TTL sync pulses. Additionally, 16 radiofrequency antennas were placed above and around the treadmill enclosure to receive the neural data transmitted through our wireless headstage (figure 4.14a).

The video capture was controlled through the Simi Motion software and the neural data capture was controlled through Blackrock Microsystem's Central software. The treadmill, obstacle, and audio tone playback were controlled through a custom C++ program.

4.4.3 Behavioral Tasks

Animals performed either basic unobstructed locomotion by walking on the treadmill without any other interactions, or obstacle avoidance by stepping over an incoming obstacle during walking. Each of these tasks were carried out in blocks. Before
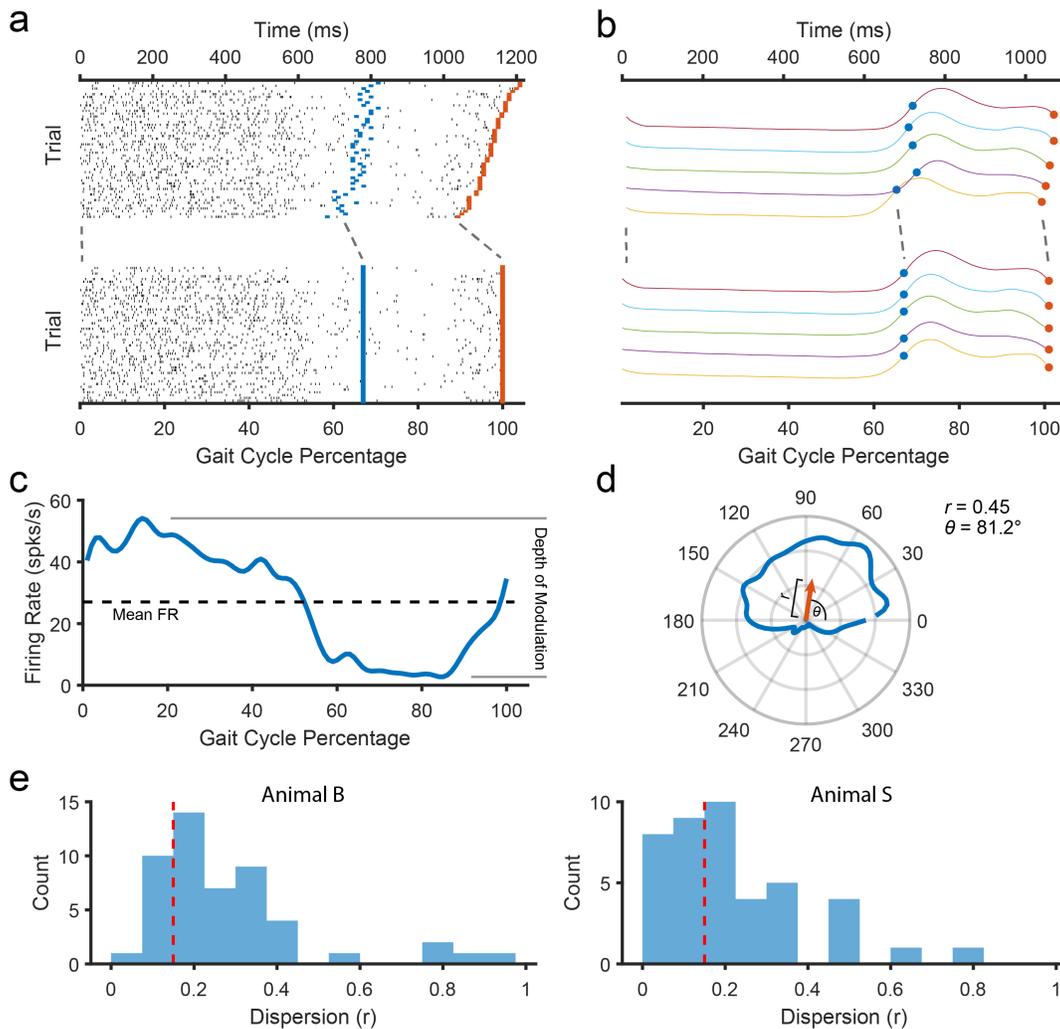
Figure 4.20: **Data normalization and circular statistics. a)** Example raster of a neuron from Animal B across multiple trials normalized to the gait cycle. Blue lines indicate foot-off time points (stance to swing transitions) and orange lines indicate foot-contact time points (swing to stance transitions). **b)** Same as in **a** but for kinematics (toe height shown here). **c)** PETH of the neuron shown in **a** after normalization. Dotted black line indicates the average firing rate across the gait cycle, and depth of modulation is calculated as the amplitude of the firing rate across the gait cycle. **d)** The PETH in **c**, but shown as a polar plot. Circular statistics were used to calculate the average directional vector (orange arrow). The magnitude of the vector is the dispersion, r, of the neuron and the angle, θ, is the preferred phase of the neuron. **e)** Distribution of dispersion values for the whole recorded population of neurons for Animal B (left) and Animal S (right). Dotted red line indicates the cutoff value of 0.15 for classifying a "weakly modulated" neuron.

entering the treadmill enclosure, animals were trained to enter a primate chair which allowed us to attach the wireless recording headstages (Cereplex W, Blackrock Microsystems, UT). The fur on the hind limb was shaved and the joint markers were also painted on at this time. Animals were then allowed to enter the treadmill enclosure where they were able to move freely.

During obstacle avoidance blocks, the treadmill was first turned on at 2.2 km/h. The obstacle would be in position in the front of the treadmill, but unmoving for the first minute to allow the animals to setting into a natural walking rhythm. At the start of each trial, a "go" tone was played to indicate that the obstacle was about to move. The operator would wait for a specific point in the gait cycle before starting the obstacle movement. The obstacle would move forward at 2.2 km/h until it was past the animal, and then rotate up out of the way of the animal. The obstacle was moved back to the front of the treadmill and rotated down into the path of the animal to assume the starting position for the next trial. After stepping over the obstacle and while it was moving back into position, a "success" tone was played and a food reward was placed in the front of the treadmill which would be carried by the moving belt to the animal. After a few seconds to allow the animal to eat the reward and resume normal walking, the "go" tone would play again to initiate the next trial (figure 4.14b). At the end of the obstacle avoidance block, the treadmill would be turned off.

During basic walking blocks, the treadmill would be turned on and the animal would walk continuously without any obstacle or food interaction for 2-5 minutes. We would also include some of the strides at the beginning of the obstacle avoidance block before the obstacle was moved for the first time as basic walking trials. We excluded the first two strides after the treadmill was turned on and the last two strides

before the treadmill was turned off to avoid any transition effects. All animals were trained to proficiently step over the obstacle without hitting it before experimental recordings were initiated.

4.4.4 Surgery

All surgical procedures were performed under general anesthesia induced through isoflurane. A craniotomy was performed and 96-channel multielectrode arrays (Blackrock Microsystems, UT) were inserted into leg area of primary-motor cortex (leg-M1) which was identified via cortical landmarks (figure 4.19). Electrodes were platinum, 1.5mm in length, and attached to a percutaneous pedestal that was fixed to the skull. Animals were given at least one week to recover after the implantation surgery before resuming the behavioral tasks. Animal B experienced some motor deficits in the contralateral limbs initially after the surgery, although it was unclear whether the cause was from neurological damage during the implantation or poor positioning on the surgery table. However, he was able to fully recover walking ability within a week of walking on the treadmill post-surgery.

4.4.5 Kinematics

We used UV reactive colored body paint to mark the positions of 6 joints of the right hindlimb. We identified the joints through bony landmarks and painted circular markers over the iliac crest (crest), greater trochanter (hip), femur lateral epicondyle (knee), lateral malleolus (ankle), 5th metatarsal head (knuckle), and 5th distal phalanx (toe). Multi-camera video tracking was used to determine the 3D position of each of the joints (SIMI Reality Motion Systems GmbH, Germany). The direction of the treadmill movement was determined through markers placed on the side of the
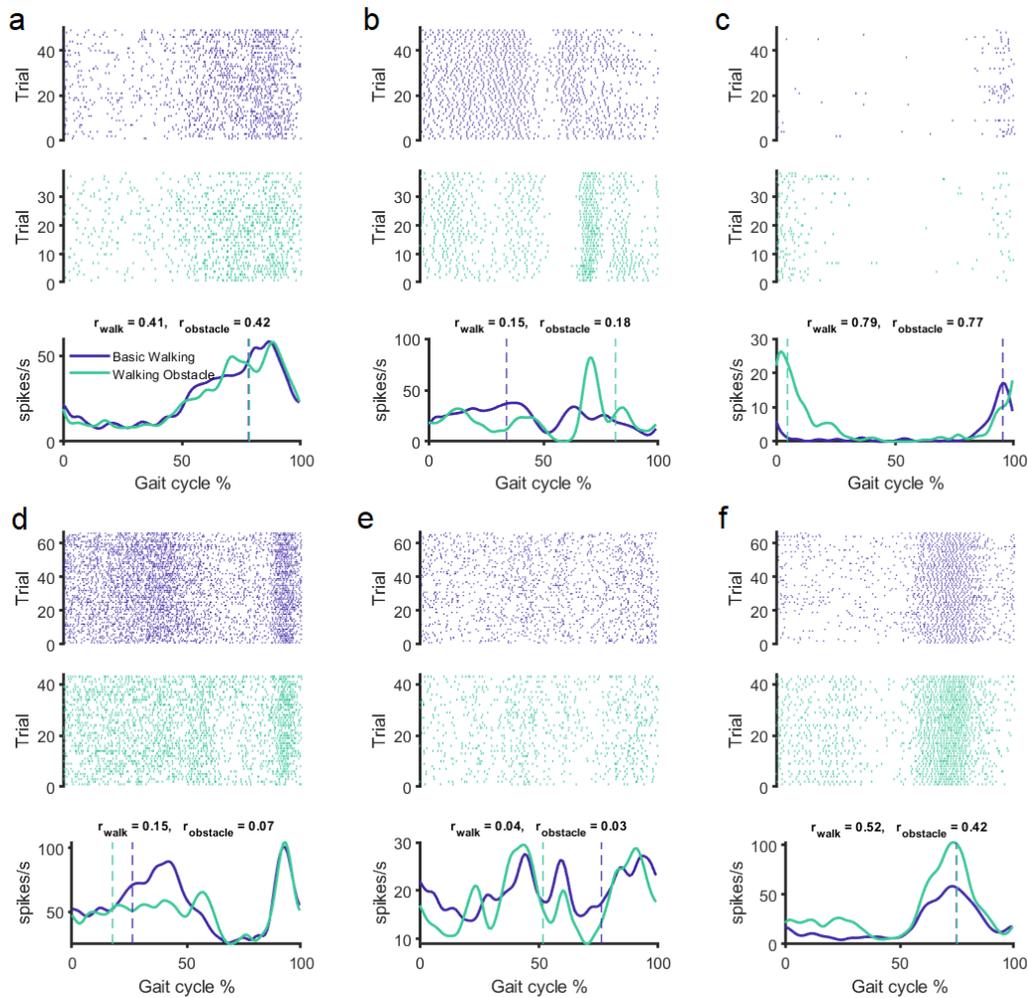
Figure 4.21: **Example M1 neurons during basic walking and obstacle stepping** Same plots as shown in figure 4.15. **a-c)** Neurons from Animal B, **d-f)** Neurons from Animal S.

treadmill and the kinematics axes were rotated so that the x-axis corresponded to the direction of walking, the y-axis corresponded to the height, and the z-axis corresponded to medial-lateral movement. Because the animals were able to freely move back and forth along the length of the treadmill, we normalized the x-position of each of the joints to the x-position of the iliac crest.

Additionally, the timing of gait events (e.g. right hand off, left foot strike, ect.) were obtained manually by inspecting the video and marking the frame when the event occurred. One stride was defined as the time from one heel-strike of the limb to the next heel strike; the stance phase was defined as the period from the first heel-strike to the next toe-off, while the swing phase was defined as the toe-off to the next heel-strike.

4.4.6   Neural data processing

Intracortical recordings were obtained at 30kHz, band-pass filtered (300-3000Hz), and thresholded at 4x the standard deviation for spike events. Spikes greater than 1.15V in amplitude were rejected as noise. For each channel, we used Wave Clus superparamagnetic clustering to semi-manually extract waveform templates, and used subtractive waveform decomposition for automated template matching of the thresholded spikes. Spike counts were obtained by binning the number of spike events into 10ms bins corresponding to each frame of the video data. Smoothed firing rate values were computed by convolving with a Gaussian kernel (s.d. 40ms).

Because each stride could vary in length, to compare across strides, we normalized the neural and kinematic data to the gait cycle. The start of the stride (heel-strike) was defined to be 0%, the toe-off moment of each stride was defined to be the duty

factor (67% for animal B and 69% for animal S), and the end of the stride (next heel-strike) was defined to be 100%. We used piecewise-linear time warping to interpolate the data at each percentage point from 0-100% (Matlab *interp1()* function). Peri-event time histograms (PETHs) were computed by averaging the normalized spike counts across trials and then smoothed with a gaussian kernel (s.d. 20ms). Because of the unconstrained movement of the animals and wireless transmission of neural data, there would be periods where none of the antennas were able to receive the signal from the headstage. We excluded trials from analysis if more than 5% of the timepoints contained dropped signal.

4.4.7   Neural response characterization

Because the strides are cyclic, we utilized circular statistics to characterize the response profiles of each neuron (Trevor Drew and Doucet, 1991). The circular mean vector of the PETH was calculated by representing each value of the PETH as a polar vector and averaging across all percentge points (figure 4.20d). The preferred phase of the neuron was defined as the angle of the mean vector while the dispersion was defined as the magnitude. The depth of modulation was calculated as the difference between the maximum and minimum firing rate of the PETH, and the mean firing rate was calculated as the average firing rate across all the gait percentages. To identify neurons whose activity was uniform across the gait cycle, we applied the Rayleigh test at $\alpha = 0.05$ with Bonferroni correction for multiple testing. We classified neurons as multi-modal if there were more than one peak that was greater than 50% of the depth of modulation for more than 10% of the gait cycle. Finally, we classified neurons as strongly modulated if they had a dispersion value greater than 0.15. All processing was done in Matlab with the Circ-stat toolbox.

To quantify the magnitude of change in the kinematics and neural activity during the obstacle stride, we calculated the Mahalanobis distance between the values during the obstacle stride and the values during the stride before any obstacle movement (three strides before the obstacle stride), which served as the baseline reference. At each gait percentage of the reference stride, we obtained the inter-trial distribution of the neural or kinematic activity in 15 dimensional space. For kinematics, this was the 3 spatial positions of the hip, knee, ankle, knuckle, and toe. For neural data, using the full 50 or 42 dimensional neural space resulted in singular matrix issues so we instead used the top 15 dimensions after performing PCA on the neural firing rates. After obtaining the mean and covariance matrix of this reference distribution, the Mahalanobis distance was computed for each trial of the obstacle stride at that gait percentage. The distance was calculated for all percentages of the gait cycle. The average and standard deviation of the distances across trials is what is shown in figure 4.15k-l. We also computed the cross-correlation between the average Mahalanobis distance of the neural activity and the average distance of the kinematics at various time lags ranging from -20 to 20 gait percentages. We used the Matlab functions *mahal()* and *crosscorr()* to implement the analyses.

4.4.8   Demixed principle component analysis

We used dPCA to find task-specific subspaces within the population activity (Kobak et al., 2016). dPCA is similar to PCA in that they are both linear dimensionality reduction methods that projects high dimensional time series data into a lower dimensional space via a decoder matrix:

$$\mathbf{X} = \mathbf{BY}, \tag{4.1}$$
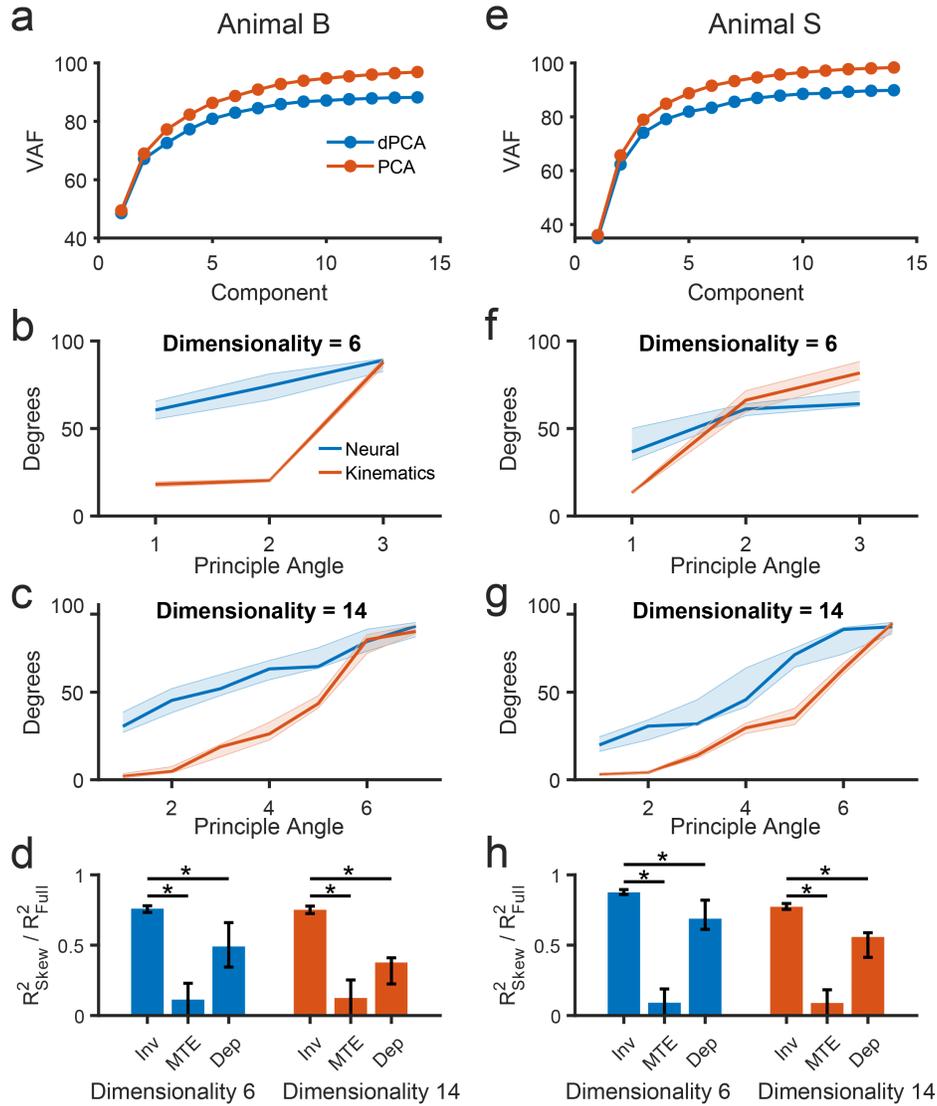
Figure 4.22: **dPCA results are robust to model dimensionality. a,e)** Cumulative of variance accounted for of the dPCA components (blue), or PCA components (orange). **b,f)** Principle angle between the stride invariant and the stride dependent subspaces for a dPCA model with dimensionality of 6 (3 for invariant subspace, 3 for dependent subspace). Error bars indicate 95% confidence intervals for 500 bootstrap reshuffles. **c,g)** Same as in **b** and **f**, but for a model with dimensionality or 14 (7 for invariant subspace, 7 for dependent subspace). **d,h)** Ratio of jPCA model R^2 to an unconstrained LDS model R^2 for the stride invariant (Inv) and stride dependent (Dep) subspaces. Also shown is the ratio for the stride invariant subspace from maximum tensor energy surrogate data (MTE). Error bars are 95% confidence intervals for 500 bootstrap reshuffles. Stars indicate statistically significant difference (500 label shuffles, $\alpha = 0.05$). **a-d** are for Animal B, **e-h** are for Animal S.

where $\mathbf{Y}$ is a n x p matrix representing the mean-subtracted PETHs of n neurons at p samples/datapoints;

$\mathbf{X}$ is the m x p matrix of the activity of m latent states, with m $<$n;

$\mathbf{B}$ is the m x n decoder matrix that projects each of the n dimensions to the lower m-dimensional space.

We refer to each of the m variables as a *latent dimension* or *neural mode*. However, unlike PCA which finds the projection that maximizes the variance accounted for by each of the m dimensions regardless of any task-related parameters, dPCA attempts to find subspaces that are related to these parameters. In dPCA, the rows of the decoder matrix can be split into groups, each corresponding to a particular task variable. For example, if there are j task variables, then the decoder matrix can be divided into

$$\mathbf{B} = \begin{bmatrix} \mathbf{D_1} \\ \mathbf{D_2} \\ \vdots \\ \mathbf{D_j} \end{bmatrix} \tag{4.2}$$

where $\mathbf{D_j}$ is the decoder matrix corresponding to the *jth* parameter. Note that the number of rows in each of the $\mathbf{D}$'s can be different, so long as they sum to m. Therefore, to obtain the $\mathbf{B}$ matrix, we will need to find each of the $\mathbf{D}$ matrices and stack them on top of each other.

In order to obtain the decoder matrix related each task parameter, the $\mathbf{Y}$ matrix is decomposed into marginalizations with respect to all possible combinations of the task parameters. If we let $\mathbf{F}$ be the set of all j parameters, then we can decompose

$$\mathbf{Y} = \sum_{f \in P(F)} \mathbf{Y}_f + \mathbf{Y}_{noise}, \tag{4.3}$$

92

where P($\mathbf{F}$) is the power set of $\mathbf{F}$. The marginalization is obtained by subtracting out the data averaged across all the parameter subsets except the parameters of the marginalization of interest, similar to ANOVA. Sets containing a single parameter correspond to marginalizations that vary with just that parameter, while sets containing multiple parameters correspond to the interaction terms. dPCA then finds a separate projection $\mathbf{D}_f$ into a low dimensional subspace corresponding to each marginalization. Like PCA, this projection is found by minimizing the reconstruction error when the latent states are projected back up to the original high dimensional space through an encoder matrix $\mathbf{E}_f$. That is, dPCA finds the projection $\mathbf{D}_f$ and $\mathbf{E}_f$ that minimizes the loss function:

$$L = ||\mathbf{Y}_f - \mathbf{E}_f\, \mathbf{D}_f\, \mathbf{Y}||^2, \tag{4.4}$$

where the matrix norm is the Frobenius norm. This essentially amounts to a reduced-rank regression problem which has a known analytical solution. $\mathbf{D}_f$ and $\mathbf{E}_f$ is calculated for all task parameters or parameter combinations $f = 1,...,j$, and the final decoder and encoder matrices are obtained by stacking the $\mathbf{D}_f$s and $\mathbf{E}_f$s. Kobak et al. provides an in-depth walkthrough of the dPCA algorithm and implementation. We applied the Matlab toolbox supplied by Kobak et al. for all our dPCA calculations.

For our data, we chose to use two task parameters: gait cycle percentage, (which in the unnormalized case, corresponds to time), and stride type. Gait cycle percentage varied from 0-100% and stride type varied from three strides before the obstacle stride to two strides after. We combined stride and stride-percentage interaction terms into the stride marginalization, since we expect the stride-related neural activity to also be time-varying. Therefore, the gait percentage marginalization subspace, which we call the stride-invariant subspace, should only vary with the gait cycle percentage and not with the stride type, while the stride marginalization subspace, which varies with

both gait percentage and stride type, will be our stride-dependent subspace. dPCA requires a dimensionality of the model be specified explicitly, so we chose a dimensionality of 10 (5 for the stride-invariant and 5 for the stride-dependent subspace) based on findings of previous studies, but we also tested our results for dimensionality of 6 and 14 (split equally between the two subspaces).

Finally to quantify the timing of the changes in the stride-dependent subspaces in relation to the changes in the movement, we used the same analysis as before during the neural response characterization (figure 4.15k-l). We calculated the cross-correlation between the mean Mahalanobis distance of the kinematics in the obstacle stride and the first component of the stride-dependent subspace in the obstacle stride at varying time lags.

4.4.9 Principle Angles

Within each marginalization subspace, the dimensions are orthogonal (that is, the rows within the encoder matrix $\mathbf{E}_f$ form an orthonormal basis). However, because each marginalization projection is calculated independently, the resulting subspaces do not have to be orthogonal to each other, meaning there could be a significant amount of overlap between the subspaces. To quantify the alignment between two subspaces, we calculated the principle angles between our stride-invarient ($\mathbf{E}_{inv}$) and stride-dependent ($\mathbf{E}_{dep}$) subspaces. The angles range from 0º to 90º which indicates perfectly overlapping or perfectly orthogonal subspaces, respectively. To obtain the principle angles, we performed singular value decomposition on the product of the two encoder matrices:

$$\mathbf{E}_{inv}^{\mathrm{T}}\mathbf{E}_{dep} = \mathbf{U}\Sigma\mathbf{V} \qquad (4.5)$$
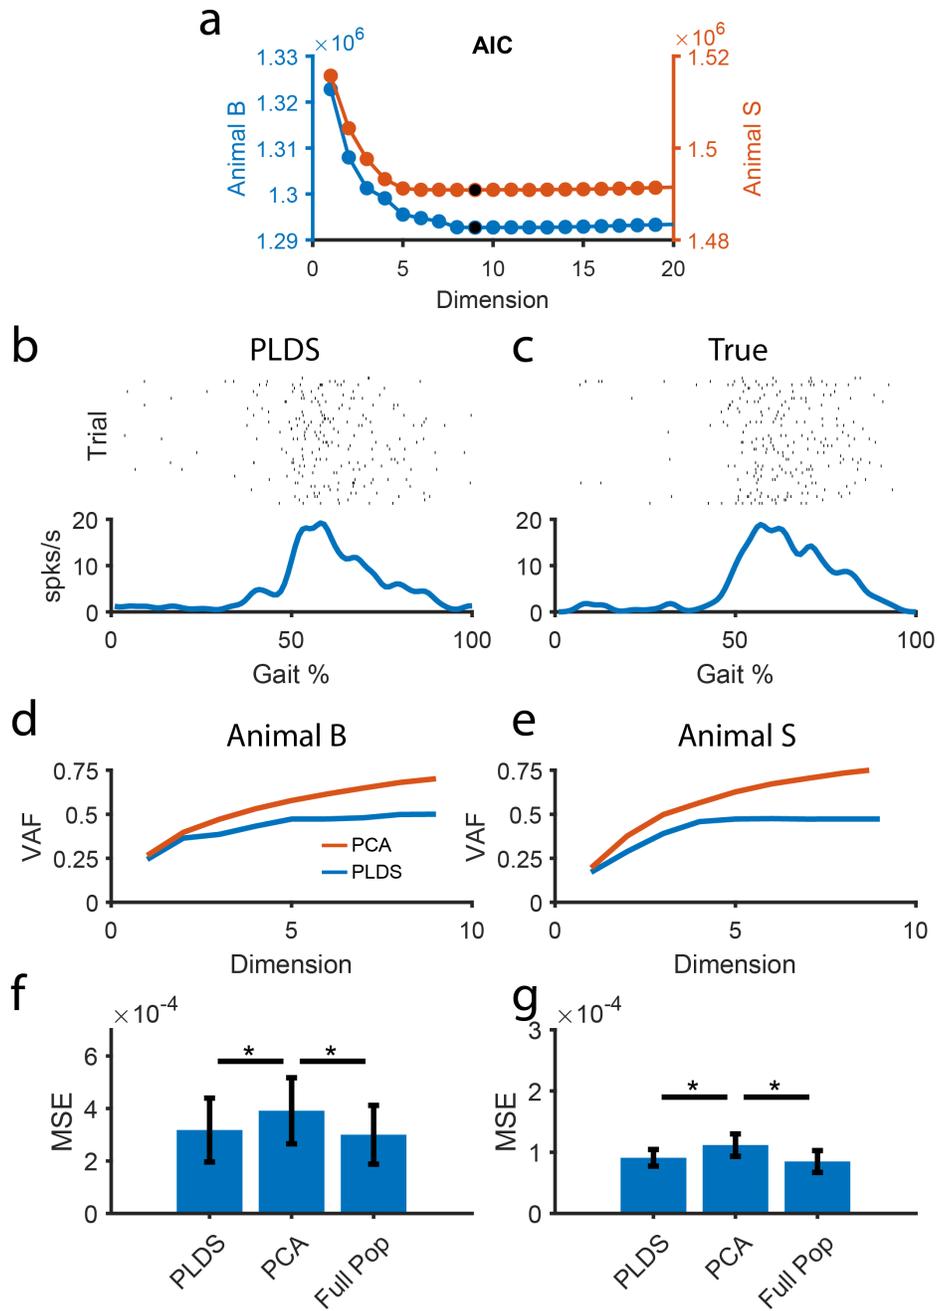
Figure 4.23: **PLDS model validation**. **a)** Akaike information criterion for the PLDS model at varying dimensions. Black dot indicates the minimum. **b)** Example neuron raster and PETH sampled from the PLDS model. The first stride of the trial (three strides before the obstacle stride) is shown. **c** The true raster and PETH of the neuron shown in **b** recorded from Leg-M1. **d,e)** Variance accounted for of the PLDS model (blue) compared to PCA (orange) at various dimensionalities. **f,g)** Mean-squared error for decoding toe height using various input features (PLDS trajectories, PCA trajectories, or the PETHs of the full neural population). **d,f** are for Animal B, **e,g** are for Animal S.

95

The diagonals of the Σ matrix are the cosines of the principle angles, ordered from smallest to largest.

We performed dPCA on both the neural PETHs as well as the trial-averaged kinematics for the obstacle avoidance trials, and calculated the principle angles between the $\mathbf{E}_{inv}$ and $\mathbf{E}_{dep}$ for both. To get a sense of the sampling variance, we performed 500 bootstraps on the trials and calculated the prinpcle angles for each bootstrap resample. The 95% confidence intervals of the bootstrap distribution are shown in figure 4.16d,i.

We hypothesized that the stride-invarient and stirde-dependent subspaces are separate. To build a distribution of the principle angles for the null-hypothesis, we used two subspaces that we know are completely aligned other than from inter-trial variance, the $\mathbf{E}_{inv}$'s (or $\mathbf{E}_{dep}$'s) from two different bootstrap resamples. We performed two bootstrap resamples and calculated the principle angles between $\mathbf{E}_{inv,\ resample1}$ and $\mathbf{E}_{inv,\ resample2}$ as well as the angles between $\mathbf{E}_{dep,\ resample1}$ and $\mathbf{E}_{dep,\ resample2}$. This was repeated 250 times to obtain our null distribution. The 97.5th percentile of the principle angles for this null distribution is shown in figure 4.16d,i.

4.4.10   Rotational structure

To model rotational dynamics within the subspaces, we employed jPCA which is a specialized variant of a linear dynamical system (LDS). The underlying assumption of dynamical systems models is that the current state of the latent components is predictive of future states. LDS models the relationship between states as a difference

equation with time-evolution matrix $\mathbf{A}$:

$$\mathbf{X}_t = \mathbf{A}\,\mathbf{X}_{t-1} \qquad\qquad (4.6)$$

Here, $\mathbf{X}_t$ is the low-dimensional neural activity in the stride-invariant or stride-dependent subspace at gait percentage t. In general-form LDS, $\mathbf{A}$ is unconstrained and can be obtained analytically through least-squares regression. In jPCA, $\mathbf{A}$ is constrained to be skew-symmetric which results in a dynamical system that contains solely rotations. While it is possible to solve for $\mathbf{A}$ analytically in the constrained case, we followed the algorithm implemented in Churchland et al. which uses a gradient-based optimization method. After obtaining $\mathbf{A}$, one can visualize the rotational dynamics by displaying the slope fields of the difference equation in the top two dimensions, as in figure 4.16e,j. The jPCA algorithm was implemented with the Matlab code provided by Churchland et al (Mark M. Churchland et al., 2012). Note that because we are already in a low-dimensional space, we did not perform the preprocessing step of using PCA to project into 6 dimensions before fitting for $\mathbf{A}$, as was done in Churchland et al.

To quantify the strength of rotational structure in each of the subspaces, we calculated the ratio of the jPCA fit R^2 to an unconstrained LDS fit R^2. The unconstrained LDS can contain both rotational and non-rotational dynamics, but if the dynamics in the data are purely rotational, then the best unconstrained LDS $\mathbf{A}$ would be the same as the jPCA $\mathbf{A}$, and the ratio would be 1. The more non-rotational dynamics are present, the more divergent the jPCA fit will be from the best LDS fit and the lower the ratio. We calculated this metric for the stride-invariant data and the stride-dependent data for both neural and kinematic dPCA. As an additional control, to test whether these rotational dynamics are the by-product of single neuron tun-

ing properties, we used the maximum tensor entropy method to generate surrogate datasets that keeps the primary features and correlations of neurons, but random higher-order moments. We ultilized the Matlab code provided by Elsayed and Cunningham to generate the surrogates and calculated the rotational fit ratio on these surrogates (Gamaleldin F Elsayed and John P Cunningham, 2017).

4.4.11   Poisson linear dynamical system model

We employed an unsupervised dimensionality reduction approach to complement our dPCA analysis. Like PCA, PLDS assumes that the recorded neural activity is the result of underlying latent states (Aghagolzadeh and Truccolo, 2014; Macke et al., 2011; Xing et al., 2019). However, in PLDS the spiking of each neuron is modeled by a Poisson process rather than the usually assumed Gaussian distribution. Additionally, PLDS explicitly models temporal dynamics by incorporating a (Gaussian) linear dynamical system similar to what was discussed in the previous section. The full model is described by:

$$
\begin{aligned}
\boldsymbol{x}_{t+1}|\boldsymbol{x}_t &\sim \mathcal{N}(\mathbf{A}\boldsymbol{x}_t, \mathbf{Q}) \\
E\left[\boldsymbol{y}_t|\boldsymbol{x}_t\right] &= \exp([\mathbf{C}\boldsymbol{x}_t + \boldsymbol{d}])
\end{aligned}
\tag{4.7}
$$

where $\boldsymbol{y}_t$ is the vector of recorded spike counts of all neurons at time point t,

$\boldsymbol{x}_t$ and $\boldsymbol{x}_{t+1}$ is the vector of the activity of the latent neural modes at time t and $t+1$ respectively,

$\mathbf{C}$ is the matrix of weights relating the neural modes to the conditional intensity function of the neurons, analogous to the $\mathbf{C}$ in dPCA,

$\boldsymbol{d}$ is the average firing rate of the neurons,

$\mathbf{A}$ is the time evolution matrix governing the temporal dynamics of the neural modes,

analogous to the **A** matrix of the unconstrained LDS model in the previous section, **Q** is the covariance of neural modes after the time evolution, i.e. the covariance of additive Gaussian noise.

Unlike PCA and dPCA, there is no analytic solution for the model, so we used the expectation maximization algorithm to infer the latent state activity and model parameters. Also, due to the non-linearity, there is no closed-form solution to finding the posterior probability of the latent neural modes given the spiking activity observations and estimated parameters in the E-step, so we employ the Laplace approximation which formulates the latent state posterior density at each time step as a Gaussian conditioned on the corresponding observations. We then estimate the values of the neural modes as the mean of this Gaussian which is calculated for each time point of each trial. A detailed description of the inference algorithm can be found in (Aghagolzadeh and Truccolo, 2014). We ran the EM algorithm for a total of 60 iterations.

To estimate the dimensionality of the underlying neural subspace, we employed the Akaike information criterion (AIC):

$$\text{AIC}(d) = 2(d^2 + d(d+1)/2 + nd) - 2\log(\mathcal{L}) \tag{4.8}$$

where d is the number of neural modes and n is the number of neurons and $\mathcal{L}$ is likelihood function. The value that minimized the AIC was chosen as the final model dimensionality.

To check that the PLDS neural modes modeled the recorded neural activity well, we reconstructed the PETHs of the neurons from the activity of the neural modes.

Spike trains for each trial was sampled from a Poisson distribution with a conditional intensity function calculated as in equation 4.7. The reconstructed PETH was then obtained by averaging across trials and smoothing.

4.4.12   Spread metric

The PLDS trajectories shown in figure 4.17 were created by time-normalizing the activity in the latent states of each trial to 0-100% of the gait cycle and then averaging across trials. We displayed the first three PLDS dimensions to visualize the time-varying activity of the neural modes throughout the 6 gait cycles as the animal stepped over the obstacle. Within this three-dimensional space, certain 2-dimensional projections resulted in large amounts of overlap in the neural trajectories across the different strides. We used the Matlab function *viewmtx()* to calculate the projections at various azimuth and elevation viewing angles. We then chose the projection that contained the most overlap as our stride-invariant subspace. A different viewing angle which highlights the changes in the neural trajectories was subjectively chosen as the stride-dependent subspace.

To quantify the amount of overlap in the low-dimensional PLDS projections between different strides, we created a spread metric which determines how divergent the neural trajectories become relative to a reference baseline stride. Because we expect the obstacle stride to have the greatest divergence from the other strides, we chose that stride as the baseline.

Let the vector $\boldsymbol{a}_{\text{ref,t}}$, be the coordinates of the reference stride in the projected 2D space at gait percentage t. For each of the other strides $i = 1...5$, we calculated the

euclidean distance of the closest point to $\boldsymbol{a}_{\text{ref,t}}$ within 10 gait percentages,

$$d_{i,t} = \min_{-10 \leq k \leq 10} ||\boldsymbol{a}_{ref,t} - \boldsymbol{a}_{i,t+k}||^2 \tag{4.9}$$

We chose to use all points within 10 gait percentages to account for any "slippage" due to imperfect time-normalization between the strides. We then took the maximum distance across all of the strides, $\max_i d_{i,t}$ to find the spread metric at gait percentage t. Finally, we obtained the final spread metric by taking the 90th percentile across gait percentages $20 \leq t \leq 80$. We excluded the first and last twenty gait percentages to avoid artifacts from edge effects and used the 90th percentile in order to exclude outliers.

4.4.13 Variance calculations

For PCA and dPCA, we measured how much neural variance the low-dimensional neural modes were able to explain by calculating the variance acounted for (VAF):

$$\text{VAF} = \frac{||\mathbf{X}||^2 - ||\mathbf{X} - \mathbf{EDX}||^2}{||\mathbf{X}||^2} \tag{4.10}$$

Where $\mathbf{X}$ is the matrix of PETHs of all neurons, $\mathbf{E}$ and $\mathbf{D}$ are the encoder and decoder matrices respectively, and the indicated norm is the Frobenius norm. For PLDS, we can not reconstruct the PETHs using encoder and decoder matrices. However, we used the approach described in the PLDS section by reconstructing the neural conditional intensity functions and sampling from the resultant Poisson distribution. We took the average of 200 resamples to account for sampling variance. Let $\widetilde{\mathbf{X}}$ be the reconstructed PETH from this sampling procedure, we then defined the VAF as

$$\text{VAF}_{PLDS} = \frac{||\mathbf{X}||^2 - ||\mathbf{X} - \widetilde{\mathbf{X}}||^2}{||\mathbf{X}||^2} \tag{4.11}$$

4.4.14 Decoding

We utilized a Wiener decoder of order 10 to predict kinematics at each gait percentage point, denoted by $\boldsymbol{y}_t$, using various neural inputs.

$$\boldsymbol{y}_t = \sum_{n=0}^{9} \mathbf{A}_n \, \boldsymbol{x}_{t-n} \tag{4.12}$$

Here $\boldsymbol{x}_{\text{t-n}}$ is the vector of input features time-shifted by n gait percentages and $\mathbf{A}_{\text{n}}$ is the matrix of decoding weights, which we computed through least-squares regression on training data. For all of our decoding analyses, we used the toe height as the decoded kinematic variable.

We tested whether a decoder trained to predict kinematics during basic locomotion could generalize to obstacle avoidance movements. We used the firing rates of the full neural population as the input vector, and calibrated the decoder weights using the neural data and kinematics from the stride three strides before the obstacle stride. We then used those weights to predict the kinematics during the subsequent strides, including the obstacle stride. We also tested the opposite, whether a decoder trained on the obstacle avoidance stride could decode kinematics during basic locomotion. In this case, the weights were trained on the neural and kinematic data during the obstacle stride and applied to each of the other strides. Note that for testing decoder generalization, we did not need to employ cross-validation since the training set for the decoder was completely separate from the testing set. However, we also measured how well the decoders trained on either the basic walking or the obstacle stride could decode the kinematics during that same stride. In these cases,

we employed leave-one-out cross validation by removing one trial from the training set and using it as the testing set. As a control, we built a decoder trained on all of the strides, and calculated the decoding performance across all of the strides using leave-one-out cross validation as well.

When testing decoder performance using low-dimensional neural modes as inputs, we utilized the same procedure, but instead of using the neural firing rates as $\boldsymbol{x}_t$, we used the neural trajectories obtained with dPCA. We measured the decoding performance using just the latent dimensions within the stride-invariant subspace, or the dimensions in both the stride-invariant subspace. In both cases, all of the strides were included in the training set, and leave-one-out cross validation was employed.

We measured decoder performance by calculating the mean squared error between the decoded kinematics, $\widetilde{\boldsymbol{y}}$, and the real kinematics across all gait percentages:

$$\mathbf{MSE} = \frac{\sum_{t=1}^{100}(\boldsymbol{y}_t - \widetilde{\boldsymbol{y}}_t)^2}{100} \tag{4.13}$$

4.5 Appendix

4.5.1 Obstacle avoidance apparatus
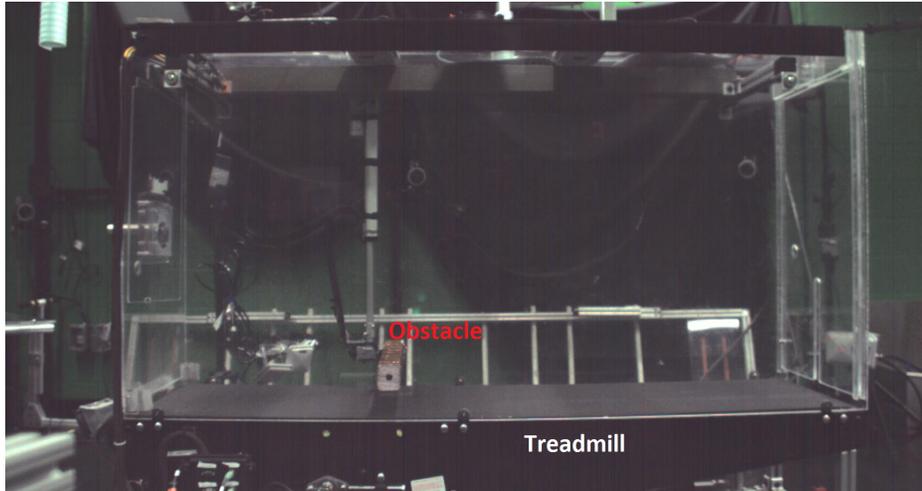
Figure 4.1: **Treadmill Enclosure**. Image showing the obstacle avoidance behavioral apparatus. Plexiglass enclosure housed animals while they walked on the treadmill and stepped over the obstacle bar.

The custom behavioral apparatus was constructed manually in-house. John Murphy, a local Brown University machinist, assisted with the construction of the enclosure walls and installation of the obstacle rig. We purchased a commercially available animal treadmill from Jog A Dog. The plexiglass enclosure was constructed above the treadmill and features a sliding, slotted, and lockable door in the back panel. This allowed the handler to keep the door open while bringing the animal into the enclosure with a pole attached to a neck-collar, before closing the door with the pole still attached. Once the door was closed and locked, the pole can then be detached, allowing the animal to freely move within the enclosure.

The obstacle component consisted of a Styrofoam bar attached to a stepper motor, allowing for rotation into and out of the path of the animal. This rotor motor was attached to a linear actuator which adjusted the height of the obstacle (Polulu

Corporation, Nevada), which was itself attached to a belt linear actuator which could move the whole obstacle arm along the length of the enclosure with programmable velocity and acceleration profiles (Igus, Rhode Island). The Polulu linear actuator was controlled by an accompanying controller board, which interfaces with a host PC via USB serial communication. The rotor motor and belt motor were both driven by Applied Motions STAC stepper drives, which communicates with the host PC via serial using their proprietary SCL protocol. Additionally, a 2-phase quadrature encoder attached to the belt linear actuator was used to accurately track the distance the obstacle has moved, while limit switches were installed at the ends of the enclosure to ensure the obstacle arm would safely stop before reaching the full travel distance of the belt actuator.
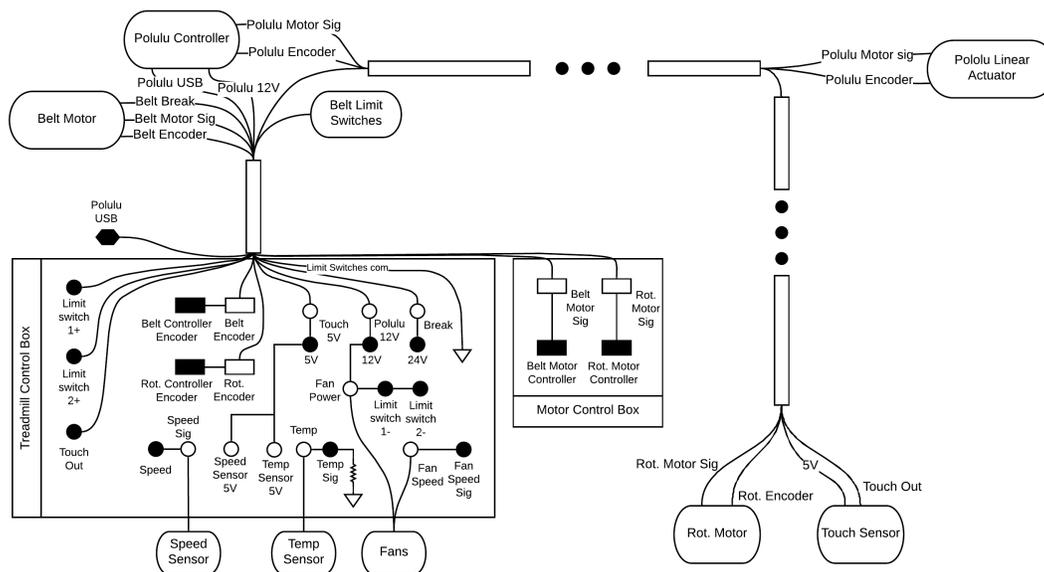
Figure 4.2: **Treadmill-Obstacle Apparatus Schematic**. Overview of the components and connections that comprised the treadmill-obstacle system. Filled black shapes indicates outgoing connections/signals from the system, while white shapes indicate incoming connections to the system. Lines represent individual wires while long rectangles represent bundles of wires, circles represent single voltage channels, rectangle blocks represent multi-pin connectors, triangle represents ground.

In addition to the obstacle controllers, audio speakers were driven by amplifiers

connected to the host PC, and an electret microphone was used to accurately determine the exact time a tone was played. Temperature sensors and fans were installed to ensure the enclosure did not get too hot. A hole was drilled into the ceiling of the enclosure to provide access for a treat dispenser to drop food reward (peanuts, Starbursts candy, raisins) into the enclosure. The treat dispenser is triggered via voltage pulse and can be controlled using one of the analog output channels in the neural signal acquisition system. Finally, a magnetic hall sensor was installed to measure treadmill speed by converting angular position of one of the treadmill axles to voltage. Most of these components and connections are implemented via a central controller box located at the bottom-left side of the treadmill. Figure 4.2 illustrates the schematic of the controller box. A separate shielded box housed the connectors supplying power to the motors to protect the data lines from noise.
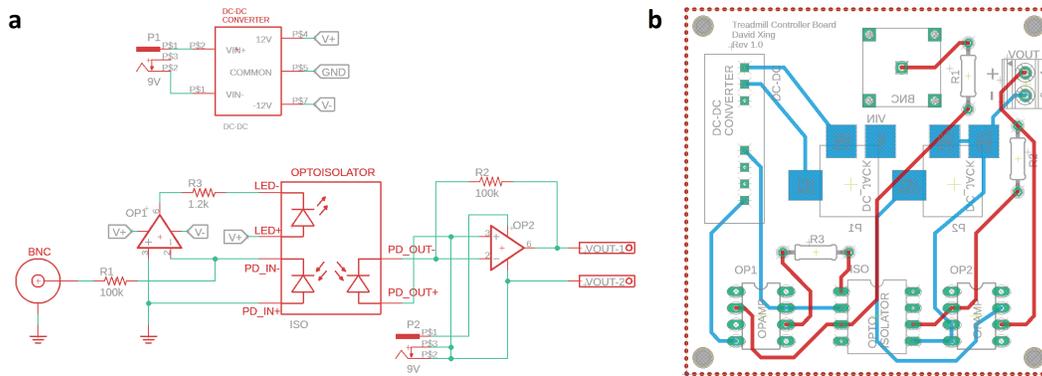
Figure 4.3: **Treadmill Controller Board**. **a)** Circuit schematic of digital treadmill controller board. **b)** Layout of the custom PCB implementing the controller board.

Unfortunately, the commercial treadmill was controlled by an accompanying motor controller box which only came with an analog speed dial. However, the circuit schematic of the controller indicated it was possible to bypass the analog potentiometer dial with an isolated input signal by changing a set of connection pins. In order to control the treadmill via the host PC, we designed and fabricated a custom printed

circuit board (PCB) which allowed for connection to a voltage output channel of the neural acquisition system via coaxial cable. The board ensured electrical isolation from the rest of the controller circuitry by implementing optoisolators and using batteries for supplying power. It also included a switch to allow for flexible swapping between analog and digital control modes. Figure 4.3 illustrates the circuit schematic and layout of the treadmill motor controller board.
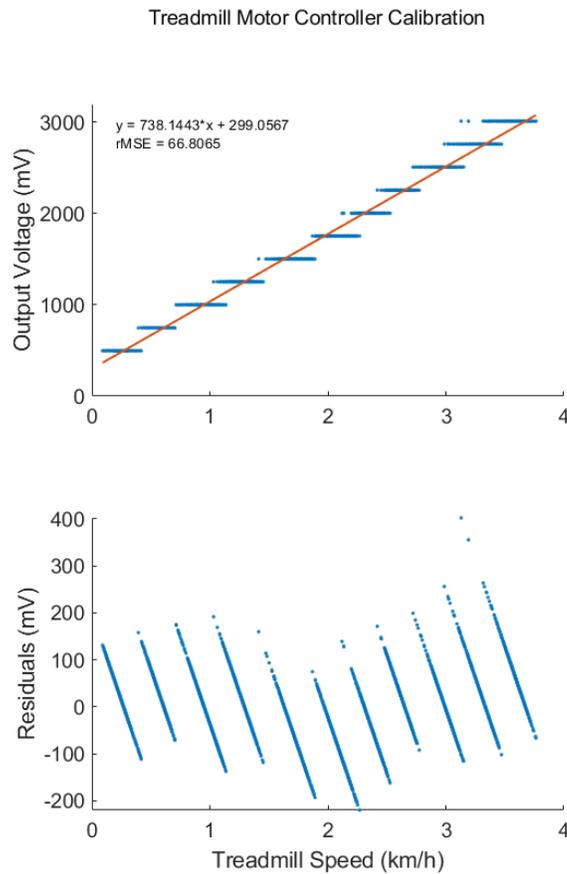
Figure 4.4: **Treadmill Calibration**. Top: calibration of treadmill controller input voltage to measured treadmill speed. Orange line represents line of best fit using least squares regression. Bottom: Residuals from the linear regression.

In order to know what voltage to supply to the controller in order to induce a certain treadmill speed, a calibration was performed prior to experimental record-

ings. Markers were painted on the treadmill belt and varying voltages from 0-3V were supplied to the controller board. Video recordings tracked the marker locations as the treadmill ran, giving a ground truth for treadmill speed. Linear regression was then used to relate treadmill speed with input voltage (figure 4.4).

A similar calibration process was also employed for the obstacle belt motor to determine the relationship between controller SCL values and linear actuator position and velocity. Figure 4.5 illustrates the results of the obstacle position and speed calibration.

a Obstacle Belt Motor Controller Position Calibration   **b** Obstacle Belt Motor Controller Velocity Calibration
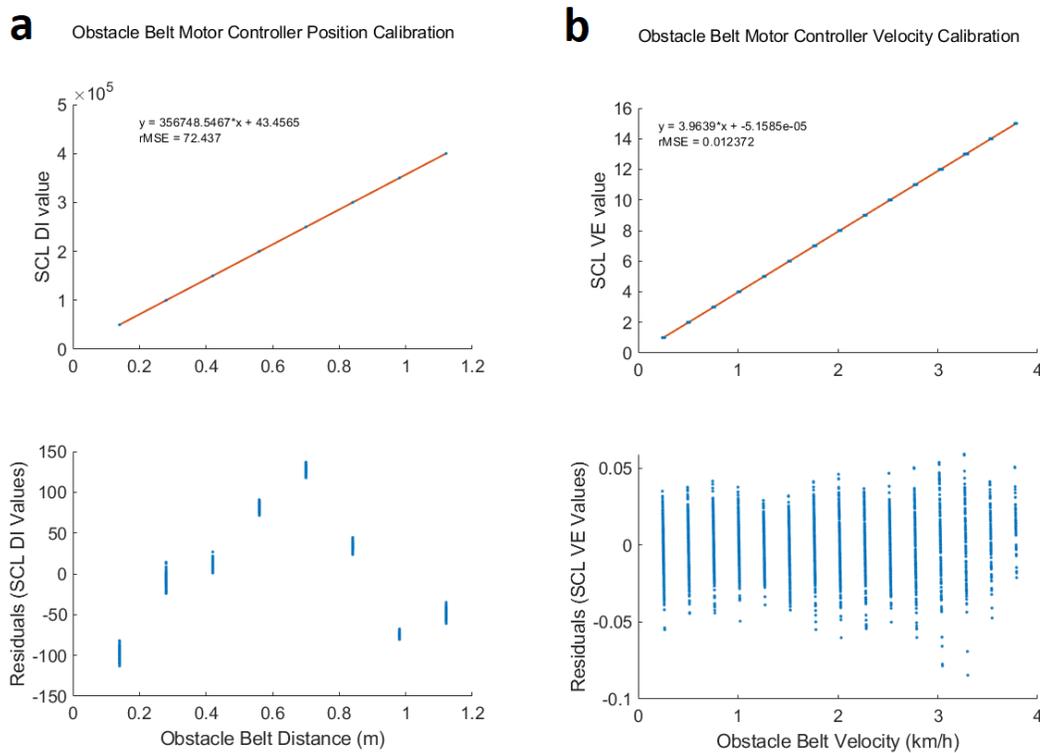
Figure 4.5: **Obstacle Calibration**. Calibration results for the obstacle linear actuator position, **a**, and velocity, **b**. Same format ast in figure 4.4

Finally, all of the apparatus components were controlled on the host PC via a custom command-line program written in C++. The obstacle encoder, treat-dispenser,

speaker, microphone, treadmill motor controller, temperature sensor, and treadmill speed sensor were all interfaced with the Blackrock Microsystems neural signal processor unit, which contains multiple analog input and output channels. Blackrock's cbSDK library was used to interface with the processor unit and provide control signals from the host PC program. Standard Windows libraries were used for serial communication with the rotor and belt motor controllers to rotate and move the obstacle. Sequences of obstacle rotation and movements were constructed and saved as obstacle run profiles within the desktop program. The program mapped specific actions to keyboard button presses. For example, pressing the 'g' key would play a "go" audio tone, while pressing the 'r' key would initiate an obstacle run trial.

4.5.2 Neural recording
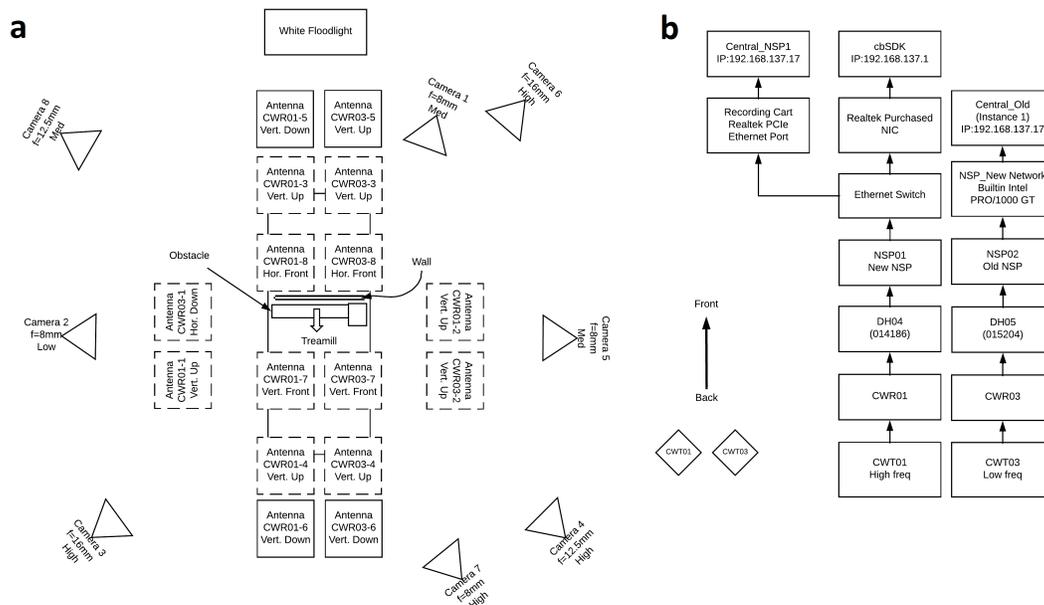
Figure 4.6: **Experimental Recording Setup**. **a)** Diagram depicting the position of neural recording antennas and video cameras in relation to the treadmill enclosure. **b)** Signal recording chain depicting the path of recorded intracortical potentials from headstage to computer.

Recording from two implanted MEAs required using two separate Blackrock neural acquisition systems. Two Cereplex W transmitters were used, tuned to different frequencies to avoid interference (3.2 and 3.5 GHz). Each went to a separate set of 8 antennas strategically placed around the treadmill enclosure to maximize coverage for a total of 16 antennas. Each set of antennas fed into a separate receiver tuned to the different frequencies which eventually fed to two separate recording computers. In addition to running the neural data acquisition, one of the computers also served as the host-PC for controlling the treadmill-obstacle apparatus (this did not appear to strain the system as CPU usage was well below the maximum). 8 high-speed video cameras were also placed around the enclosure, although only three were placed on the left side while 5 were placed on the right side because the limb of interest is the right hind-limb. These were connected to network switch which fed into a third,

110

video-capture dedicated computer. Figure 4.6 illustrates the recording setup and signal processing chain.

Before placing the animal inside the treadmill enclosure, they were first placed in a restraint chair and a measurement of the noise rms voltage of each of the 96 electrode channels were made. This value was used to determine the threshold value for action potential detection for each channel. Figure 4.7 demonstrates the waveforms of the threshold crossing for animal B.
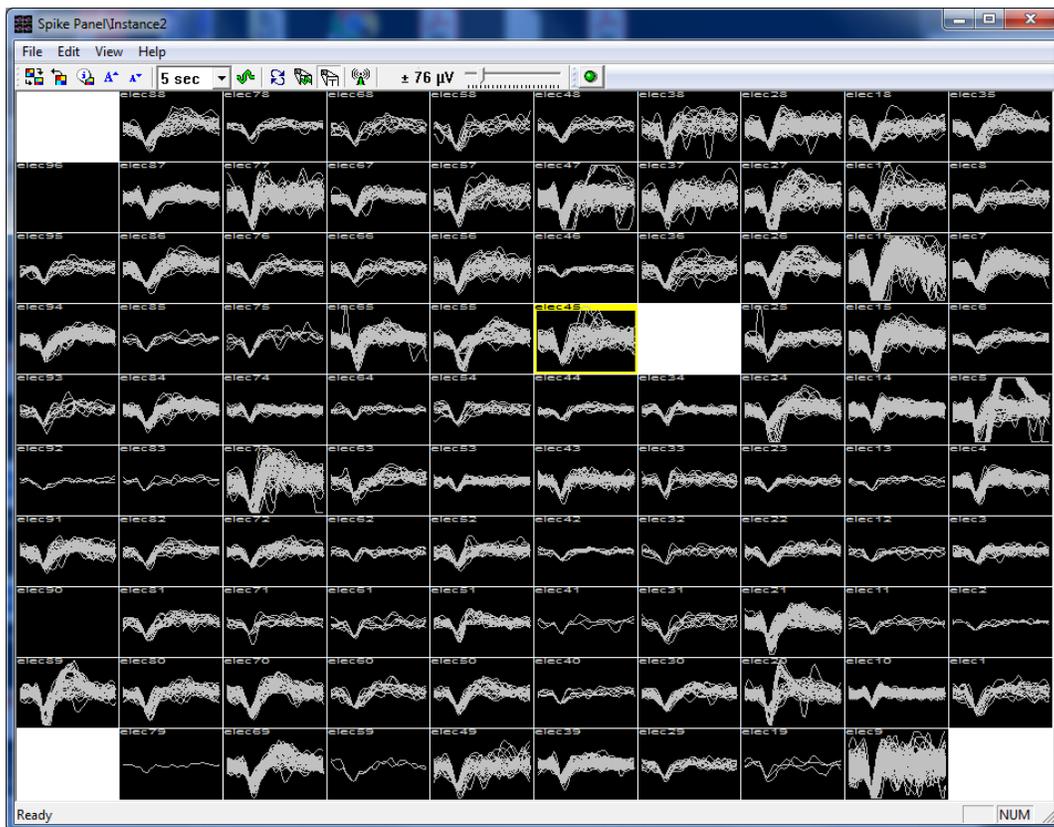
Figure 4.7: **Microelectrode Array Recorded Units**. Threshold crossings for all 96 electrodes in the implanted MEA of Animal B. 44 days after implantation. Threshold was set to 4.0x the standard deviation

Finally, after the neural data is saved, superparamagnetic was used to simi-automatically detect waveform templates for each channel (Quiroga, Nadasdy, and

Ben-Shaul, 2004). Any templates with a signal to noise ratio, defined as the ratio of the waveform amplitude to three standard deviations, above 2 was saved as a putative neuron, and all waveforms which matched the template were assigned to that unit (figure 4.8.
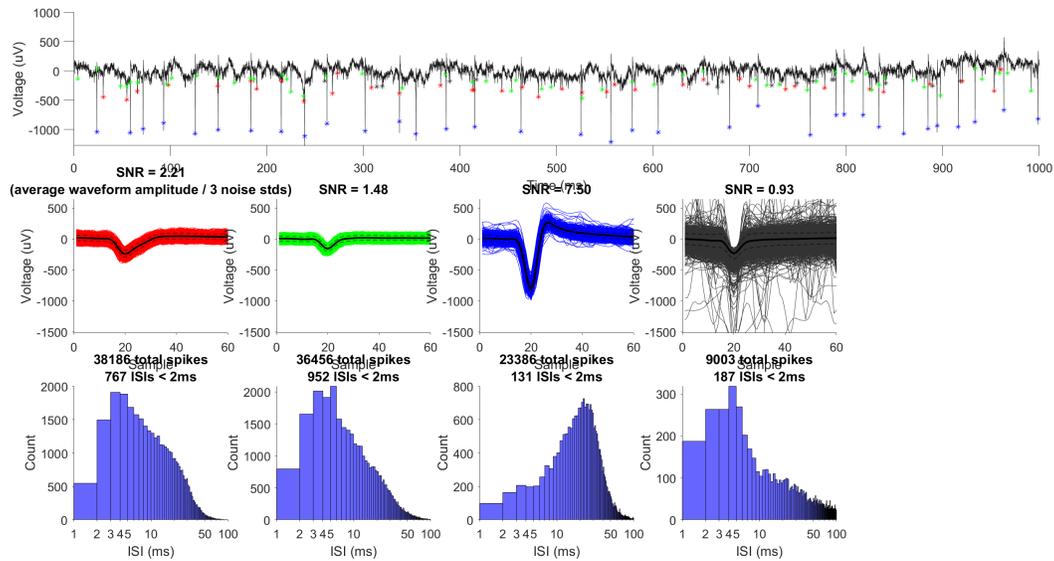
Figure 4.8: **Spike Sorting**. Classified waveforms of single units for an example electrode. Top: intracortical voltage recorded by the electrode. Detected spikes are indicated by a marker of the color of the assigned unit. Bottom: waveforms from all of the classified units along with the inter-spike interval histograms. Black denotes all rejected waveforms (not assigned to any unit).

4.5.3 Wireless EMG recording
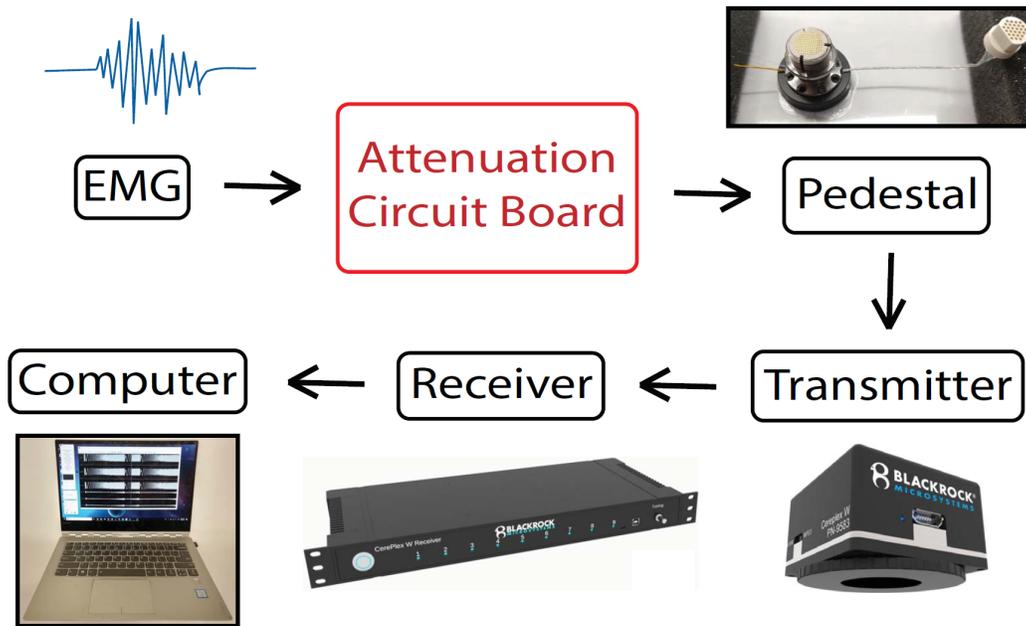
Figure 4.9: **Wireless EMG Approach**. Strategy for recording EMG wirelessly from animals performing obstacle avoidance, leveraging existing commercial hardware.

While the kinematics illustrates one aspect of the movement carried out by the animal during obstalce avoidance, it is useful to also record muscle activity in order measure the direct output of motor neurons. Unfortunately there are currently no commercially available systems designed for wireless EMG acquisition. Here, I will describe our attempt at adapting existing wireless neural recording systems for obtaining unteethered EMG measurements. Much of this work was carried out with the assistance of an (at the time of writing) undergraduate at Brown University, Beatriz Silveira de Arruda.

The over-all strategy for the EMG acquisition system is to exploit the wireless recording capabilities of the Blackrock Microsystems Cereplex W system. This commercially available system is capable of recording from 96 channels of electrophysiological signals, amplifying, sampling, and wirelessly transmitting them to a windows

computer for storage. Currently, the systems is used for recording extracellular neural signals from micro-electrode arrays, and is indeed what we use to obtain our M1 recordings. A percutaneous "pedestal" (see figure 4.9) is fixed to the skull of the subject, which serves as a connection point to route the MEA signals to the amplifier/digitizer/transmitter component (e.g. Cereplex W transmitter). Since the components for wireless electrophysiology are already present in this system, we wanted to modify it to receive signals from intramuscular, in addition to intracortical, recordings.

We have collaborated with Blackrock Microsystems and they have manufactured custom pedestals which routes 16 of the 96 channels to a Omnetics connector rather than to the MEA. We can then connect EMG electrodes (stainless steel Cooner wire) to the Omnetics connector to be used with the Cereplex W system (figure 4.9).

However, one issue is that the signals from the MEA and the Omnetics connector are sent to the same amplifier array, which is designed to take in signals from extracellular neural recordings, and has an input range of -8.191mV to 8.919 mV. EMG amplitudes can be much higher, potentially extending to the hundreds of millivolts range, and would saturate the amplifiers if fed directly to the current system. Therefore, it is necessary to attenuate the signals by the proper amount to be compatible with the Cereplex W headstage and transmitter. We designed and created a custom PCB utilizing an array of voltage dividers reduce the recorded voltages by 100x, or 40dB (figure 4.10a). We were able to use off-the-shelf components along with specific Omnetics connectors to construct this board and miniaturize it to a diameter of 0.9 inches (figure 4.10b).

Unfortunately, because the board must interface with the custom Blackrock pedestal
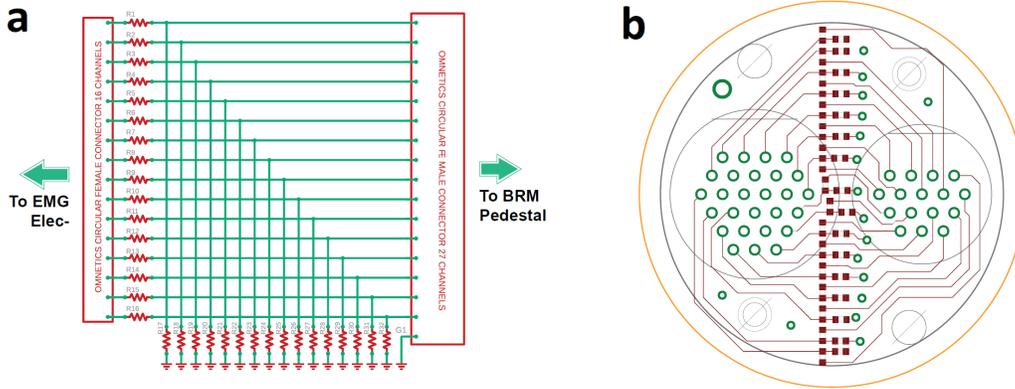
Figure 4.10: **EMG Attenuation Board**. **a)** Circuit schematic of the attenuation board. Circuit consists of basic voltage dividers for each of the EMG channels. **b)** PCB Layout of the attenuation board.

for integration into their data collection chain, it must be placed on top of the head, necessitating an additional head-mounted enclosure. We therefore also designed and fabricated a custom titanium pedestal aimed to house the Omnetics connector attached to the Blackrock pedestal, and the Omnetics connector attached to the intramuscular EMG wires (figure 4.11a-b). This would be chronically implanted, similar to the Blackrock pedestal and would be fixed to the skull with dental cement attached to skull screws. The opening would be covered up and the connectors would be protected by a custom cap that could only be removed only with the use of a screwdriver, to protect from tampering from the animal (figure 4.11e-f). During recording sessions, we would remove the cap and replace it with another component that houses the EMG attenuation board (figure 4.11c-d). The screw-on housing allows for easy connection of the board to the underlying Omnetics connectors and thereby completing the pathway from EMG electrodes, to attenuation circuit, to Blackrock headstage and transmitter.

After manufacturing all the components of the system, we performed a bench-top test validation test to verify that we were able to record electrical signals through the

115
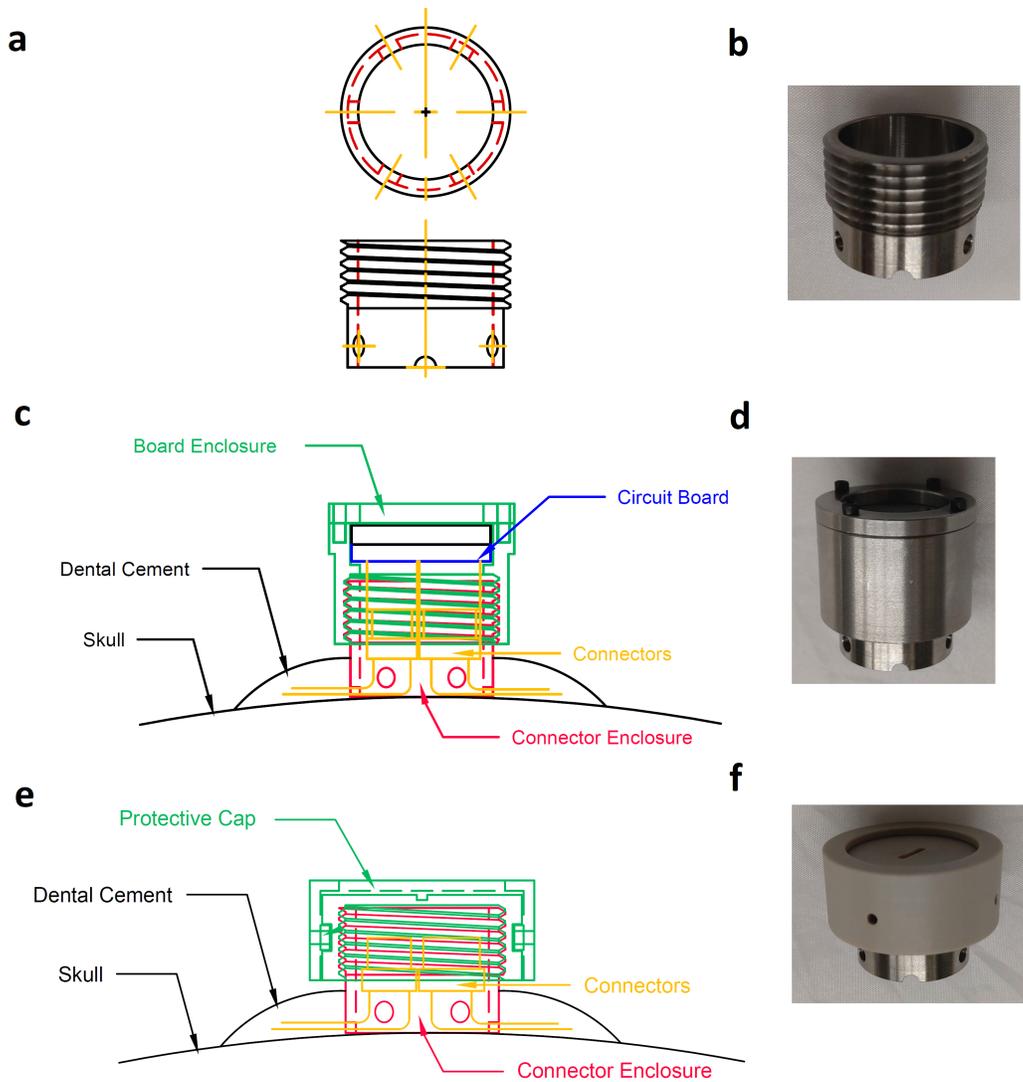
Figure 4.11: **Attenuation Board Implant**. **a-b)** Schematic (left) and picture (right) of the skull mounted titanium pedestal housing the Omnetics connectors attached to the EMG wires and going to the Cereport pedestal. **c-d)** Same as in **a-b** but including the headstage mount housing the EMG attenuation board. **e-f)** Same as in **a-b** but including the protective covering.

EMG wires. We used a function generator to output 10mV and 50mV peak-to-peak sine wave chirps with a frequency sweep of 1Hz to 1kHz (10 second sweep time, log sweep). This output was directly connected to the EMG electrodes and was routed to the Blackrock wireless headstage via the attenuation board. The recorded signals were saved on the computer and compared to the original output of the function gen-

erator, which was also saved using the Blackrock signal acquisition system, but using the analog input channels (rather than ephys frontend channels) instead. Figure 4.12 illustrates the results of the test. The signal recorded via the headstage (orange) follows the real signal (blue) accurately, with an attenuation of 40dB, although there appears to be some phase lag at lower frequencies. The Bode plot confirms these observations, as the phase diagram is non linear, although the deviations appear to be fairly small in the 5-200Hz range. Both the magnitude and phase plots exhibit a large spike at 60Hz, which is most likely due powerline noise. Overall the system performs well and we expect to be able to use it to record muscle activity in behaving animals.

We implanted the modified Blackrock pedestal with the EMG connector in animal B during the MEA implantation surgery. We planned to perform an additional surgery to implant the EMG wires and the second pedestal for housing the connector after the recovery from the first surgery, and so we temporarily placed the EMG Omnetics underneath the scalp near the lambda landmark. However, the skin around the connector appeared to be irritated and the incision eventually dehisced. Unfortunately this allowed the animal to access the connector with his fingers and we had to remove the implanted connector. We were therefore unable to test the EMG recording system *in vivo*, and additional experiments will need to be carried out to fully validate the wireless EMG system.
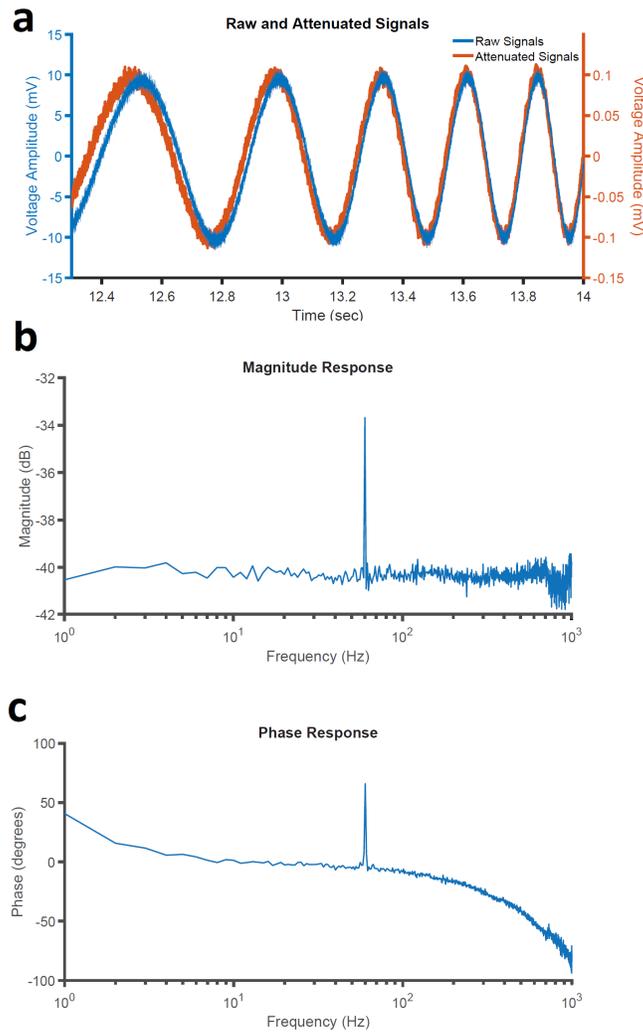
Figure 4.12: **Benchtop System Testing**. **a)** Plots of the sine-wave chirp created by the signal generator (blue) and recorded on the computer though the Blackrock system (orange). **b)** Bode plot attenuation magnitude of the EMG recording system. **c)** Bode plot phase shift of the EMG recording system.

Chapter 5

Genetic manipulation of the corticospinal tract

5.1 Introduction

Over the past few decades, genetic manipulation of protein expression in neurons have developed into a powerful and effective tool for interrogating cell or pathway-specific functions in neural circuits. For example, optogenetics has allowed researchers to activate or inhibit specific cell types or regions in the brain (Deisseroth, 2015), and has found widespread use in perturbation studies in mice (Miri et al., 2017; Deisseroth, 2015). Unfortunately, the development of these tools in primates are greatly lagging behind their rodent counterparts. However, in recent years, researchers have been working on the development of a tet-on/off genetic system for non-human primates. In 2012, it was successfully applied to spinal neurons in rhesus macaques(Kinoshita et al., 2012). Kinoshita et al. leveraged recombinant viral vectors to achieve tract-specfic and temporally reversible shutdown of intersegmental interneurons, and demonstrated that these propriospinal neurons are necessary for the control of independent finger movements. This technique, when applied to other spinal tracts, could be extremely

useful for interrogating the role of specific pathways in different behaviors. In relation to the study of locomotion, it could allow us to causally perturb the descending signals from motor cortex onto spinal circuits and determine whether cortical inputs are an essential component to carrying out locomotion in non-human primates. In this chapter, I will discuss our attempt to apply the genetic shutdown system to the cortical spinal tract.

5.1.1 Tet-on system and double-virus injection
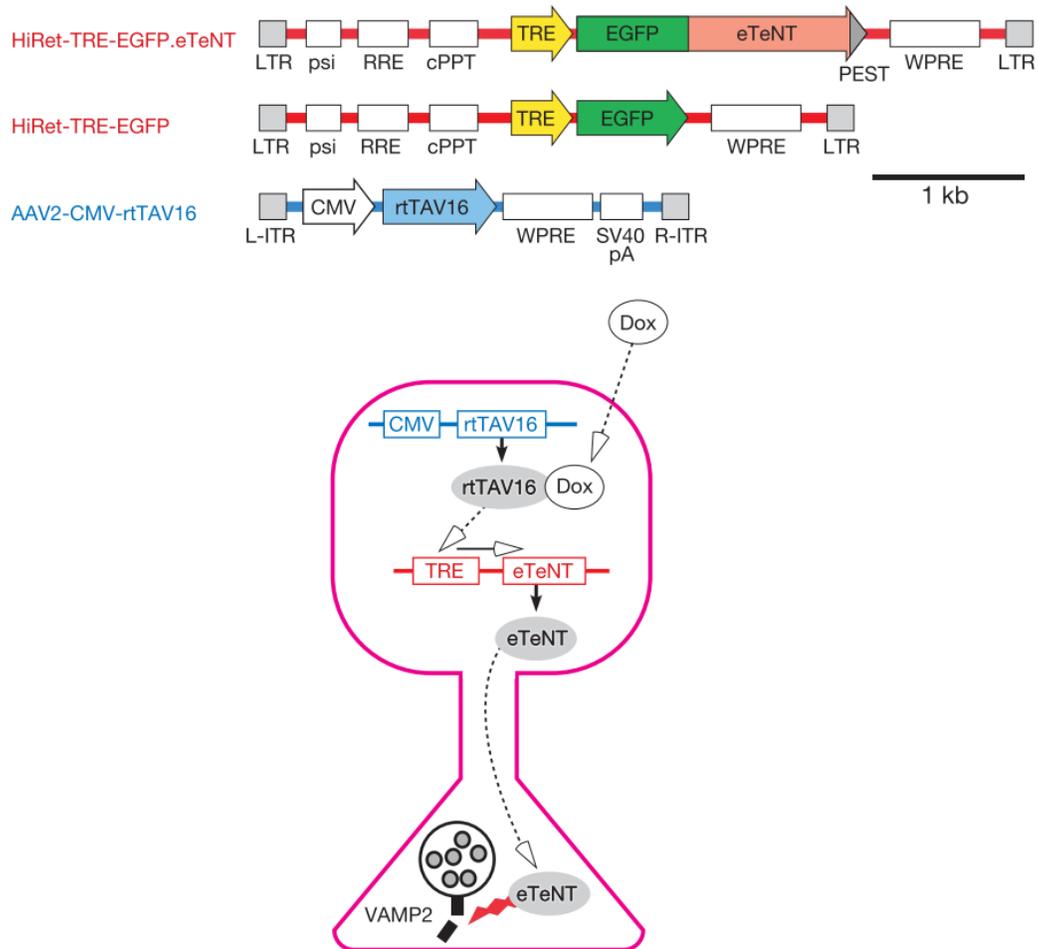
Figure 5.1: Tet-on double virus system. Adapted with permission from (Kinoshita et al., 2012)

The tet-on system allows for specific transcription of a DNA sequence, which in our case, will be the sequence encoding an enhanced version of the tetanus neurotoxin protein (eTeNT). This protein acts by cleaving VAMP2/synaptobrevin2, a protein involved in synaptic neural transmitter release, thereby blocking synaptic transmission of any neurons expressing it(Yamamoto et al., 2003). We want to be able to express eTeNT exclusively in corticospinal tract neurons originating in motor cortex and projecting to the lumbar spinal cord. Additionally, we wish to control when the neurons begin expressing eTeNT and have a way to reverse the lesion and restore connectivity between M1 and lumbar spinal cord. The temporal specificity and reversability can be achieved through the tet-on system. Transcription of the eTeNT sequence is under the control of a tetracycline-responsive element (TRE) promoter, which is only activated when a tetracycline transactivation protein (rtTAV) is bound to it (Gossen and Bujard, 1992). Additionally, rtTAV cannot bind to the TRE unless in the presence of tetracycline (tet) or its analogue, doxycycline (dox)(Zhou et al., 2006). Therefore, a neuron must contain both rtTAV and dox (and of course, the DNA sequence for the eTeNT) in order to express the neurotoxin and shut down synaptic transmission (figure 5.1, bottom). Doxycycline can be administered systemically so only the delivery of the rtTAV and eTeNT sequence via viral vectors is necessary. After the cells have been transfected with the two sequences, the researcher can control when to start eTeNT expression by withholding the dox administration until the desired time. The lesion is also reversible because doxycycline is rapidly eliminated from the body. Pharmacokinetic analyses have found the complete removal of dox from the blood serum within days (Embers et al., 2013; Kinoshita et al., 2012). Therefore, halting dox administration will arrest eTeNT production and reverse the neural shutdown.

The neural pathway specificity is achieved by separating the delivery of the rtTAV sequence and the eTeNT sequence into two separate vectors. One vector is an AAV2

virus which transfects the soma of the cells at the injection site. The second vector is a psudotyped version of the HIV1 lentivirus, modified to have high retrograde transport properties (HiRet)(Kato, Kobayashi, et al., 2011). The second vector is injected at the targeted axonal projections of the cells at the AAV injection site. The cells projecting to the second site will uptake the retrograde virus at the synaptic terminals and the payload is transported up the axons back to the soma(Naldini et al., 1996). Here, the rtTAV sequence (under the strongly expressing CMV promoter) is delivered via the AAV virus while the rtTAV16 is delivered via the lentivirus (figure 5.1, top). With this double virus system, only those cells whose soma is located in the first injection site, and whose axons projects to the second site will become transfected with both viruses. In our case, we wished to shut down the hindlimb corticospinal tract neurons so we targeted the AAV virus injection to leg area of M1 and the lentivirus injection to lamina IX of the lumbar enlargement in the spinal cord.

5.2 Methods

We used the same two rhesus macaques in Chapter 3, animal B and animal S, to test the genetic shutdown system. In both animals, the cortical and spinal injections were carried out in the same surgery. We allowed 1-2 months for the viruses to transfect the neurons after the injection surgery before starting the administration of dox. The behavioral assay used was basic locomotion on the treadmill enclosure described in Chapter 3, along with general observations on the motor dexterity of the targeted limb. We performed a baseline recording of treadmill locomotion the day before the start of dox administration. We then administered doxycycline for one week and performed behavioral analysis throughout. Unfortunately, because these animals were used for other experiments, such as the one described in Chapter 3, we

were unable to perform histology to asses the transfection rate until many years after the initial injection. Figure 5.2 illustrates the experimental timeline.
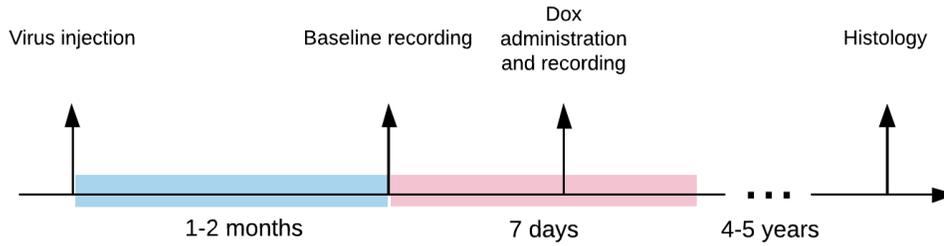
Figure 5.2: Timeline of experimental procedures for the double-virus shutdown system

5.2.1 Virus Injection Surgery

Viral vehicles were received from Dr. Tadashi Isa as a gift. For animal S, we performed many small injection at multiple sites. However, for animal B, we wished to test a novel injection technique that produces a greater volume of coverage from a single injection site, convection enhanced delivery (CED) (Yazdan-Shahmorad et al., 2016). While traditional injection utilizes small injection volumes ($\leq$10 μL), CED uses large injection volumes (50 μL) combined with fast injection rates. This relied on the creation of a custom cannula in order to prevent backflow up out of the injection hole (reflux). Dr. Ilker Ozden assembled and provided the cannulas. Figure 5.3 demonstrates the approximate region that would be exposed to the virus. Almost all of leg area of primary motor cortex is covered by a single injection, greatly reducing the time of the case. Both animals were injected in the left hemisphere of motor cortex and right hemisphere in the spinal cord. The procedure for the CED injection is located in the appendix of this chapter.
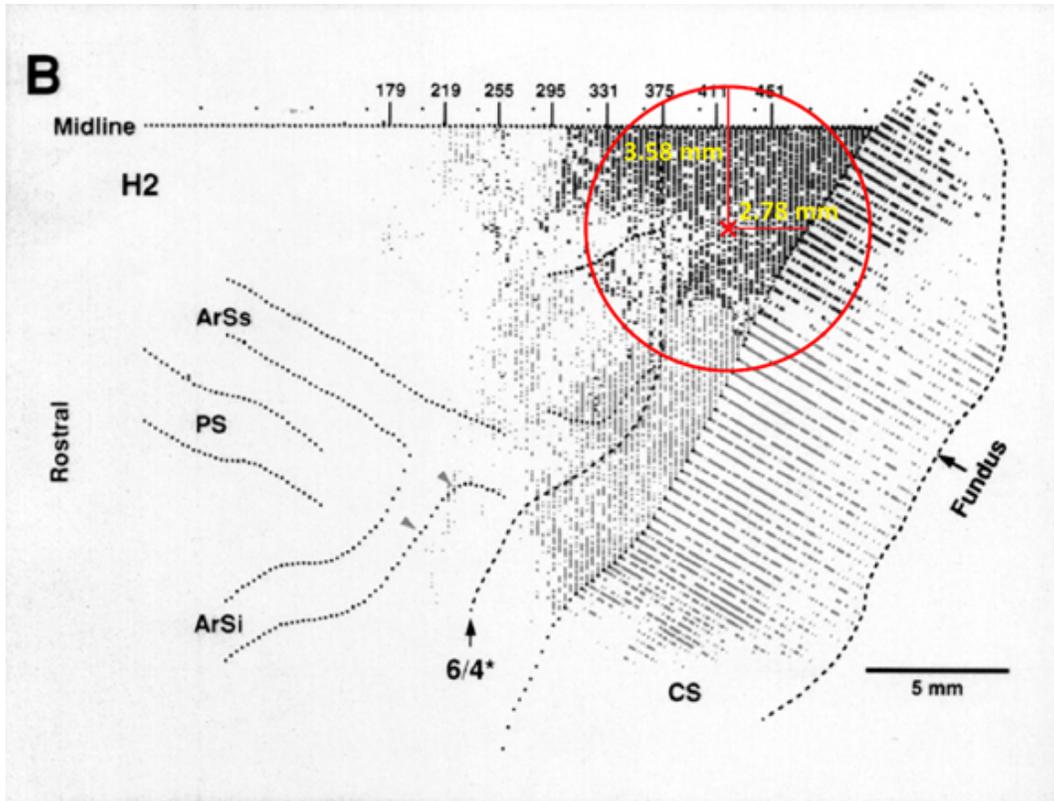
Figure 5.3: Leg area of M1. Red circle indicates the estimated coverage of CED injection at the red cross. Adapted from (He, Dum, and Strick, 1993)

5.2.2 Doxycycline administration and treadmill walking

During doxycycline administration, we provided doxycycline (10mg/kg) orally every morning via a liquid suspension. This suspension was mixed with apple juice for palatablity. In the afternoon, the juice bottles were examined to ensure complete consumption of the dox before animals were taken out of their home cage. Animals were placed in a primate chair, UV reactive markers were applied to the crest, hip, knee, ankle, and toe joints, and then animals were allowed to enter the treadmill enclosure. They would walk on the treadmill at approximately 2.2 km/h for 1-2 minutes while being video recorded by 8 synchronized cameras. After the walking session, we would dispense a food reward and return the animals to their home cage. Food and water was provided ad lib throughout the duration of the recordings.

We obtained the 3D kinematics of the joints from the video recordings using the methods described in section 4.4.5. Like before, the values were time-normalized to 0-100% of the gait cycle with the toe-off event locked to 60%.

5.2.3   Euthanasia and histology

After the doxycycline period, Animals were used for additional projects. At the end of those projects, animals were euthanized (Sodium barbiturate, IV), perfused, and dissected. The perfusion procedure is located in the appendix of this chapter. We extracted the brain and spinal cord and dissected out the injection regions. Unfortunately, we needed to implant microelectrode arrays into leg-M1 in order to complete the other projects, which introduced a confound to our histology assessment. While removing the MEA from the cortex, the tissue around the injection site had adhered to the electrodes and the whole section of tissue was exhumed along with the MEA (figure 5.6). For animal S, this removed almost the entire extent of leg-M1 from the cortical surface; for animal B, some of the grey matter was left intact.

Despite this setback, we went ahead and cryoprotected our tissues in sucrose solution for 4 days before sectioning in a cryostat (40μm thickness). An enhanced green fluorescent protein (eGFP) sequence was upstream of the eTeNT sequence so we stained for anti-GFP to assay eTeNT expression. The tissues were also counterstained for NeuN and DAPI to visualize neuronal cell bodies and cell nuclei respectively. Elaina Atherton and Sophie Brown performed the histology procedure and confocal imaging.

5.3 Results

After commencing dox administration, we did not observe any notable changes in either their behavior or movement. Before we induced the genetic shutdown with dox, the animals would often used their feet to grab at the handler during the painting stage before entering the treadmill. This behavior was not abolished during the the dox administration, and we noticed no deficit in the animals abilities to manipulate their hind-limb extremities. We also did not observe any change in the gait kinematics during treadmill walking (figure 5.4).
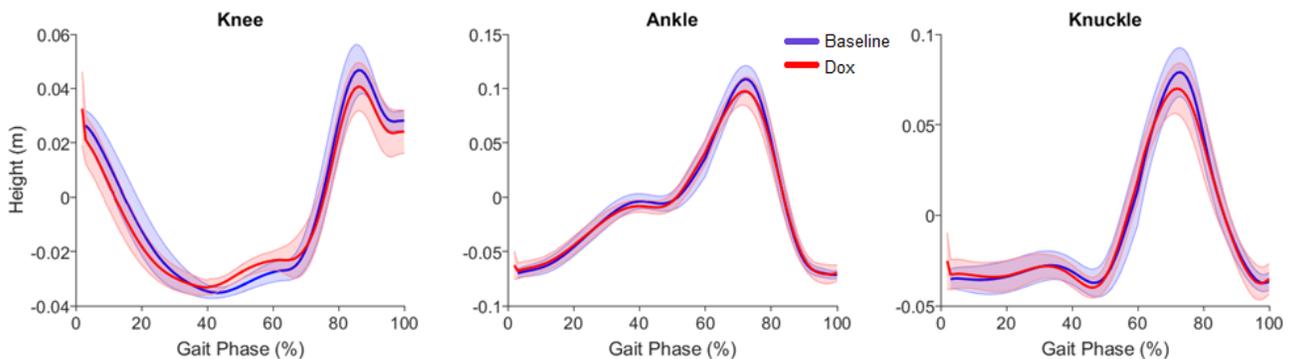
Figure 5.4: Walking kinematics before and during doxycycline administration. Error bars are 1 standard deviation

The histology was inconclusive. In animal S, we did not find any expression of eGFP. In animal B, we were able to detect eGFP in two of cells from two different tissue sections (figure 5.5). However the NeuN stain did not co-localize with anti-GFP. Based on the morphology of the cells, the anti-GFP cells are likely to be astrocytes. We did not detect any cells co-stained for both NeuN and anti-GFP.

5.4 Discussion

Overall, we saw little change in the movements or behavior of the animals after administering the doxycycline and initiating the genetic shutdown. As explained in
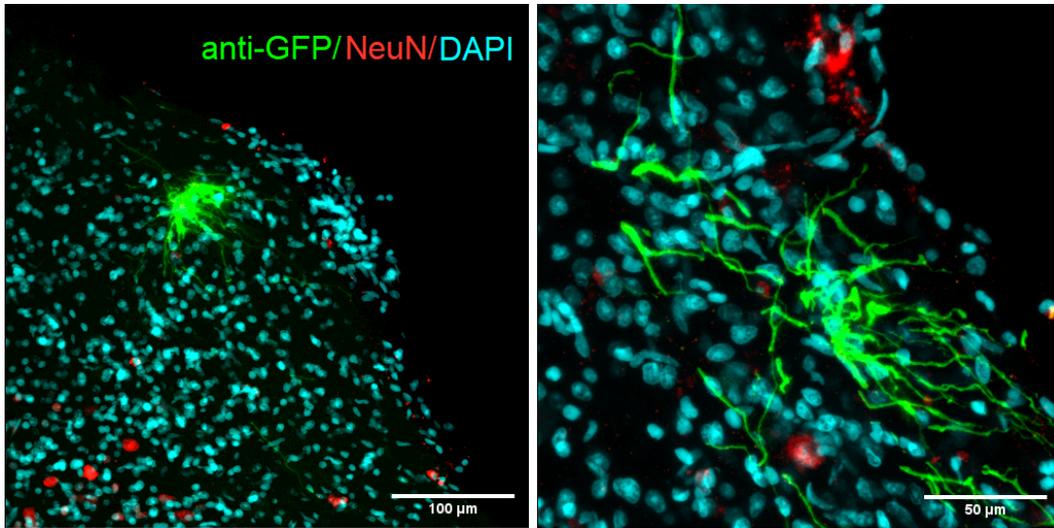
Figure 5.5: Two different coronal sections of leg-M1 from animal B, highlighting the presence of eGFP in motor cortex.

previous chapters, motor cortex may not be necessary for the generation of locomotion movements in non-human primates, and therefore the lack of any observable motor deficit during walking may indicate that the superficial role that motor cortex plays in locomotion for felines, canines, and rodents (Miri et al., 2017; Sten Grillner, Georgopoulos, and Jordan, 1997; G raham Brown, 1911) extends to macaques as well. However, we also observed very little change in the animals ability to perform voluntary, dexterous movements in the targeted hind limb. Because it is well established that M1 plays an essential role in these types of movements, and previous lesion studies have reported a clear motor deficit in voluntary movements post-lesion (Courtine, 2005; Lemon et al., 2012), it is more likely that we either did not transfect a large enough population of neurons to produce an observable effect, or we failed to induce a shutdown of the corticospinal tract altogether. Unfortunately, due to the use of animals in multiple studies, we were unable perform histological analysis until 4-5 years after the injection, and there were confounds which may affect our interpretation of the results.
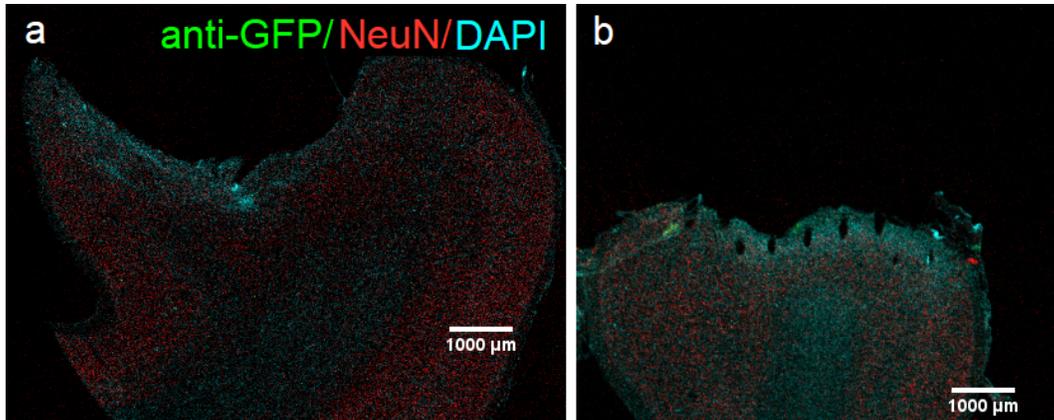
Figure 5.6: Coronal slices of leg-M1 for animal S (**a**) and animal B (**b**) illustrating the tissue damage sustained from removing the implanted MEAs.

There are multiple points of failure that could have been the cause of this negative outcome. One possibility is that the doxycycline may not have been ingested by the animal because we did not directly observe the animal drink the juice containing the doxycycline, and instead relied on ensuring that the juice bottle was empty after a few hours. However, we were able to detect doxycycline in the blood serum post-administration, so this failure point is unlikely. Another possibility is that the viruses were removed from the ice bath for too long during the injection surgery and were not maintained at the appropriate temperature, resulting in poor transfection rates. Yet another possibility may have been that our injection depth was inaccurate and we were not able to target layer V neurons. However, the fact that we were able to detect any eGFP at all in animal B indicates that at least for this subject, the tet-on system was successful. That is, neurons must have received both the eTeNT sequence from the AAV virus as well as the rtTAV sequence from the lentivirus, and have taken up enough doxycycline to initiate the transcription of eGFP and the downstream eTeNT. The lack of behavioral response could be the result of an insufficient number of neurons expressing eTeNT, or it could be that our spinal injection was too far rostral and the M1 neurons are projecting to spinal regions associated with trunk, rather than hind-limb, muscles.

It is also surprising that we did not find any eGFP in the M1 neurons, but in a few astrocytes (figure 5.5). The astrocytes themselves cannot produce eGFP since they do not have axons projecting to the spinal cord and therefore cannot uptake the lentivirus. It has been shown that astrocytes can take on phagocytic functions after tissue damage, and clear up surrounding neural debris (Morizawa et al., 2017; Wakida et al., 2018). Given the proximity of our eGFP-labeled astrocytes to the MEA implant, it is likely that these astrocytes absorbed the remnants of nearby apoptotic neurons expressing eGFP. We should note that it is not altogether too surprising that we were unable to find any eGFP-expressing neurons given the extensive cavity present in the tissue. Figure 5.6 illustrates the area of the missing tissue from M1 in our sections for both animals. In animal S, virtually all the cortical layers have been expunged in the medial aspect of the gyrus. Because this is exactly where we expect our transfected neurons to be located, these histology results cannot be used to determine the efficacy of our viral injections.

Further experiments need to be performed, probably with a dedicated set of animals, in order to comprehensively characterize the genetic shutdown technique. In these future experiments, electrophysiological recordings should compliment behavioral analysis to determine whether the pyramidal neurons were truly blocked from their downstream spinal targets. Additionally, histological assessment should be carried out immediately, without external confounds, to asses the extent of the transfection. Finally, having a well-defined behavioral assay for dexterous movements, such as obstacle avoidance, would serve as a useful control. Nevertheless, the fact that we were able to express eGFP in cortex at all indicates that the application of the tet-on system is viable for corticospinal projections. All previous implementations have been limited to short, locally projecting pathways(Kinoshita et al., 2012; Hioki

et al., 2009; Kato, Inoue, et al., 2007), so it was unknown whether the retrograde transport of the viral payload was able to traverse the long axonal distance spanned by the corticospinal tract. That we were able to successfully induce expression in even a small number of cells is still a promising finding.

5.5 Appendix

5.5.1 CED injection protocol

Virus Drawing:

1. While the animal is being prepped, retrieve one 50 uL aliquot of the AAV and one 50 uL aliquot of the lentivirus from -80 freezer. Put into a Styrofoam transfer box with dry ice, and bring into the operating room.

2. Set up virus drawing station. It will be a base with a segment of stereotactic arm. Connected to the base is a micromanipulator connector to which a micromanipulator arm is attached. The arm will hold the pump for injecting and drawing up the virus (through a connector piece). The pump and syringe will be sterile, so a surgeon will attach the syringe to the pump, and the pump connector to the micromanipulator arm. The virus drawer will then put a tube holder rack right below the cannula of the syringe. The drawer will not touch the pump (except the plug which goes to the pump controller), syringe, or connector.

3. 40 min before the start of the injection, take out an aliquot of the AAV or Lentivirus (whichever is being injected) and place it in the tube rack to warm to room temperature.

4. 10 min before the start of the injection, open the tube containing the virus, and slowly lower the syringe cannula into the tube using the micromanipulator arm being careful not to hit the walls of the tube. The tip of the cannula should be as close to the bottom of the tube as possible, but not touching it. Use a magnifying glass to help visualize the position of the tip in the tube.

5. Once the cannula is in a good position, draw up 50 uL of the virus (the whole aliquot) at a rate of 5 uL/min for 10 min (settings on the pump for 100 uL

syringes: W50000 .0000 5000 MGN (use MHN for 250 uL syringes)). Check the the syringe reading to make sure the correct volume was drawn.

6. Raise the syringe/cannula out of the tube. Then use the pump to slowly push the plunger down (use 100 uL/min rate) until a droplet is visible at the tip (using a magnifying glass). This is to push any air out.

7. The surgeon will now detach the syringe and pump/pump connector from the withdrawing arm and place it onto the injection arm.

Cortical Injection:

1. Perform craniotomy to expose the motor cortex.

2. Position the tip of the cannula 3.58mm lateral from the midline, and 2.78mm rostral from the Central Sulcus. The injection volume should cover a 10mm diameter area (Yazdan-Shahmorad et al., 2016) which will engulf all of the leg area of the motor cortex (see figure 5.3).

3. Slowly lower the micromanipulator arm until the tip of the cannula just touches the cortex. Quickly (if not using the programmable Microdrive) lower the cannula 2 mm down into the cortex. If the cannula could not penetrate the pia and arachnoid, a small incision can be made using a scalpel.

4. Once the cannula is in the cortex, wait 10 min. Then start injecting the virus at 1 uL/min for 4 minutes, 4 uL total, marking the start and end volumes in the syringe to verify that the correct volume was injected (Pump settings for 100 uL syringes: I4000. .0000 1000MGN (use MHN for 250 uL syringes)). Right after, increase the injection volume to 2 uL/min for 5 min, 10 uL total (Pump settings: I10000 .0000 2000MGN (use MHN for 250 uL syringes)). Finally, increase the volume to 5 uL/min for 7 min, 35 uL total (Pump settings: I35000 .0000 5000MGN (use MHN for 250 uL syringes)).

5. After all the virus has been injected, wait 15 min, then pull the cannula out of the cortex.

6. Close the craniotomy.

Spinal Cord Injection:

1. We will want to cover L1-L7 spinal segements of the cord since that is where the motor neuron pools for the hind limbs reside. First expose the spinal cord through a L2 vertebral laminectomy to access the midpoint of the column we want to transfect.

2. Previous CED injection in rhesus spinal cord have shown over 4 cm of perfusion longitudinally with 50 uL injection, albeit with rate of only 0.1 uL/min and injection into the white matter not grey matter (Lonser et al., 1998). However rat studies comparing grey vs white matter injection indicates that the fluid tends to seep out the grey matter and travel along the white matter

3. Move the cannula tip to the center (rostral-caudally) of the L2 vertebral laminec-tomy. To target the ventral horn of the grey matter, move the cannula 1.18 mm lateral to the midline. Then lower the cannula to the top of the cord. Lower the cannula 3.39 mm into the cord. Wait 10 minutes.

5.5.2 Perfusion protocol

1. Administer sodium barbiturate. Check for cessation of corneal reflex, breathing, and QRS complex to confirm death.

2. Make an incision along the midline of the chest to separate the skin and underlying fatty tissue and expose the sternum. Cut the sternum and use retractors to open the chest cavity and expose the heart.

3. Cut and remove the pericardium, locate the aorta and vena cava. Make an incision at the apex and insert the needle connected to the peristaltic pump tubing into the left ventricle. Moved the needle up into the aorta and clamped it in place using a S-shaped hemostat tool. Make an incision in the Vena Cava.

4. Start the pump. Pump heparinzed PBS at 75 ml/min. until the liquid running out of the vena cava appears clear.

5. Switch the line to neutral-buffered formalin (NBF). Apply suction to the area around the heart to avoid any formalin being spilled onto the draft table. The muscles will start to writhe at this time.

6. After about 2L of NBF, stop the pump and let the body sit so it can continue to cross link for an additional 5 minutes.

7. Begin dissection.

Chapter 6

Summary and future steps

In this work, we explored the use of dimensionality reduction methods to to characterize the properties of the primary motor cortex in a variety of locomotor tasks. Many of the findings and techniques developed for studying forelimb stationary reaching movements can be applied to the study of locomotion as well. This includes the dynamical systems perspective of neural control. The idea the neurons do not act as isolated entities but as part of a computational program is rather intuitive. Graham Brown's half-center model provides an excellent example of this view. Two mutually-inhibited ideal neurons would be perfectly (anti)correlated, and therefore can be completely characterized by a one dimensional parameter. Based on our PLDS analysis, the dimensionality of M1 neurons during locomotion appears to be around 10. Using these 10 latent states, we are able to reconstruct walking kinematics just as well as, if not better, than the full neural population in high-dimensional space. As we have shown, this may have benefits for brain-machine interfaces since a high-dimensional feature space needs larger amounts of training data, which could mean longer calibration sessions. Of course, good offline decoding does not necessarily translate to actual use-case gains. Further testing with closed-loop decoders will be required in order to determine the true value of implementing a PLDS locomotion decoder.

Beyond being useful for the development of novel neurotechnologies, these low-dimensional neural modes can also provide insight into how the nervous system carries out goal-directed behavior. Motor cortex is faced with a uniquely challenging task when it is required to perform a dexterous volitional movement during locomotion. While it is still unclear whether motor cortex is playing an active role in driving the muscles during basic unobstructed locomotion, it is unlikely that spinal CPGs do not contribute to the activation of the muscles at all. Therefore, there is a separate entity inducing movements in the limb, while cortex is tasked with carrying out a precise movement with that same limb. From a control systems standpoint, it would certainly be beneficial if cortex is made continuously aware of the activity of this outside entity. Anatomically, there is a pathway from the spinal CPGs to motor cortex, via the dentate nucleus and ventral lateral thalamus. Computationally, we have found the existence of a low-dimensional space within motor cortex activity that is invariant to dexterous volitional movements. Even when the limb is lifted nearly five time higher than normal, these neural modes continue to maintain the same rotational trajectories; all that variance is contained to a separate subspace. Could these stoic oscillators be the reflections of spinal CPGs? Our results seems to support this hypothesis. However, without perturbation studies, it is difficult to definitively make that claim.

Unfortunately, we were unable to successfully implement the genetic shutdown model in our subjects. The tet-on system could have been used to provide the required perturbation: shut down the spinocerebellar tract and observe if the activity in those invariant modes are uniquely affected. Additionally, shutting down the corticospinal tract would have provided some insight into the relevance of M1 activity to ongoing locomotion movements. If we were able to observe a change in locomotion

136

gait after inducing the shutdown, that would be evidence for the engagement of motor cortex during basic locomotion. As it stands, there is still a lot of optimization to do before we can make the double virus technique consistent and robust. Admittedly, a better designed experiment plan would have alleviated many the issues that arose. Nevertheless, we were still able to demonstrate that the system can work. The presence of eTeNT in the motor cortex, however small, indicates that all the individual parts came together. Increasing the transfection rate will be a crucial next step for the development of the this shutdown model.

To conclude. While applying the dynamical systems view of cortical activity to locomotor movements has revealed some insights into its underlying neural processes, I believe that there is still much to be learned about the neural control of locomotion, especially complex locomotion. With the advent of new recording technologies and analytical techniques, we are able to start collecting large datasets across a variety of behaviors. I am optimistic that these developments will lead to greater discoveries and greater advancements for neuroprosthetics.

Bibliography

Afshar, Afsheen et al. (Aug. 2011). "Single-trial neural correlates of arm movement preparation". In: *Neuron* 71.3, pp. 555–64. ISSN: 1097-4199. DOI: `10.1016/j.neuron.2011.05.047`. URL: `http://dx.doi.org/10.1016/j.neuron.2011.05.047%20http://www.ncbi.nlm.nih.gov/pubmed/21835350%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3155684`.

Aghagolzadeh, Mehdi and Wilson Truccolo (2014). "Latent state-space models for neural decoding". In: *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference* 2014, pp. 3033–6. ISSN: 1557-170X. DOI: `10.1109/EMBC.2014.6944262`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/25570630%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4289598`.

— (Feb. 2016). "Inference and Decoding of Motor Cortex Low-Dimensional Dynamics via Latent State-Space Models". In: *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society* 24.2, pp. 272–82. ISSN: 1558-0210. DOI: `10.1109/TNSRE.2015.2470527`. URL: `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7229368%20http://www.ncbi.nlm.nih.gov/pubmed/26336135%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4910090`.

Armer, Madison C. et al. (2013). "Effect of light on the activity of motor cortex neurons during locomotion". In: *Behavioural Brain Research* 250, pp. 238–250. ISSN: 01664328. DOI: 10.1016/j.bbr.2013.05.004. URL: http://dx.doi.org/10.1016/j.bbr.2013.05.004.

Armstrong, D M and T Drew (Jan. 1984a). "Discharges of pyramidal tract and other motor cortical neurones during locomotion in the cat". In: *The Journal of physiology* 346, pp. 471–95. ISSN: 0022-3751. DOI: 10.1113/jphysiol.1984.sp015036. URL: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1199513&tool=pmcentrez&rendertype=abstract%20http://www.ncbi.nlm.nih.gov/pubmed/6699782%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC1199513.

— (1984b). "Locomotor-related neuronal discharges in cat motor cortex compared with peripheral receptive fields and evoked movements." In: *The Journal of physiology* 346, pp. 497–517. ISSN: 0022-3751. DOI: 10.1113/jphysiol.1984.sp015037. URL: http://www.ncbi.nlm.nih.gov/pubmed/6699784%5Cnhttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1199514&tool=pmcentrez&rendertype=abstract.

Armstrong, D. M. (Jan. 1986). *Supraspinal contributions to the initiation and control of locomotion in the cat.* DOI: 10.1016/0301-0082(86)90021-3.

Armstrong, DM and T Drew (1985). "Forelimb electropyographic responses to motor cortex stimulation during locomotion in the cat". In: *J Physiol* 367, pp. 327–351.

Arshavsky, Y I et al. (Aug. 1972). "Origin of modulation in neurones of the ventral spinocerebellar tract during locomotion". In: *Brain research* 43.1, pp. 276–9. ISSN: 0006-8993. DOI: 10.1016/0006-8993(72)90296-X. URL: http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Origin+of+modulation+in+neurones+of+the+ventral+spinocerebellar+tract+during+locomotion#0%20http://www.ncbi.nlm.nih.gov/pubmed/5050196.

Athalye, Vivek R et al. (Feb. 2017). "Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control". In: *Neuron* 93.4, pp. 955–970. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2017.01.016. URL: http://dx.doi.org/10.1016/j.neuron.2017.01.016%20http://www.ncbi.nlm.nih.gov/pubmed/28190641.

Beloozerova, I N and M G Sirota (Feb. 1993). "The role of the motor cortex in the control of accuracy of locomotor movements in the cat". In: *The Journal of physiology* 461, pp. 1–25. ISSN: 0022-3751. DOI: 10.1007/s00415-010-5533-4. URL: http://www.ncbi.nlm.nih.gov/pubmed/8350266%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC1175243%20http://www.ncbi.nlm.nih.gov/pubmed/8350259%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC1175242.

— (1988). "Role of Motor Cortex in Control of Locomotion". In: *Stance and Motion*. Boston, MA: Springer US, pp. 163–176. ISBN: 978-1-4899-0821-6. DOI: 10.1007/978-1-4899-0821-6{_}15. URL: http://books.google.com/books?hl=en&lr=&id=89kUuheAexAC&oi=fnd&pg=PA163&dq=Role+of+motor+cortex+in+control+of+locomotion&ots=Bqdm9UrJSC&sig=LBZPH_sKvFe6AKDo0V5RGHVFFS8%20http://link.springer.com/10.1007/978-1-4899-0821-6_15.

Beloozerova, Irina N and Mikhail G Sirota (2003). "Integration of motor and visual information in the parietal area 5 during locomotion." In: *Journal of neurophysiology* 90.2, pp. 961–71. ISSN: 0022-3077. DOI: 10.1152/jn.01147.2002. URL: http://www.ncbi.nlm.nih.gov/pubmed/12904498.

— (Nov. 1998). "Cortically controlled gait adjustments in the cat". In: *Annals of the New York Academy of Sciences* 860, pp. 550–3. ISSN: 0077-8923. DOI: 10.1111/j.1749-6632.1998.tb09101.x. URL: http://www.ncbi.nlm.nih.gov/pubmed/9928360.

Bishop, Christopher (2006). *Pattern Recognition and Machine Learning.* 1st ed. Springer-Verlag New York. ISBN: 978-0-387-31073-2. URL: `https://www.springer.com/gp/book/9780387310732%20http://www.springer.com/us/book/9780387310732%5Cnhttps://www.springer.com/us/book/9780387310732`.

Bouton, Chad E. et al. (2016). "Restoring cortical control of functional movement in a human with quadriplegia". In: *Nature* 533.7602, pp. 247–50. ISSN: 1476-4687. DOI: 10.1038/nature17435. URL: `http://www.nature.com/doifinder/10.1038/nature17435%5Cnhttp://www.ncbi.nlm.nih.gov/pubmed/27074513%20http://www.ncbi.nlm.nih.gov/pubmed/27074513`.

Brandman, David M., Sydney S. Cash, and Leigh R. Hochberg (2017). "Review: Human Intracortical recording and neural decoding for brain-computer interfaces". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 4320.c, pp. 1–1. ISSN: 1534-4320. DOI: 10.1109/TNSRE.2017.2677443. URL: `http://ieeexplore.ieee.org/document/7869404/`.

Bruno, Angela M, William N Frost, and Mark D Humphries (Aug. 2017). "A spiral attractor network drives rhythmic locomotion." In: *eLife* 6. ISSN: 2050-084X. DOI: 10.7554/eLife.27342. URL: `https://elifesciences.org/articles/27342%20http://www.ncbi.nlm.nih.gov/pubmed/28780929%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5546814`.

Capogrosso, Marco et al. (2016). "A brain-spine interface alleviating gait deficits after spinal cord injury in primates". In: *Nature* 539.7628, pp. 284–288. ISSN: 14764687. DOI: 10.1038/nature20118. URL: `http://dx.doi.org/10.1038/nature20118%20http://www.ncbi.nlm.nih.gov/pubmed/27830790%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5108412`.

Churchland, Mark M. et al. (2012). "Neural population dynamics during reaching". In: *Nature* 487.7405, pp. 51–56. ISSN: 0028-0836. DOI: 10.1038/nature11129. URL: `http://dx.doi.org/10.1038/nature11129`.

Collinger, Jennifer L. et al. (Feb. 2013). "High-performance neuroprosthetic control by an individual with tetraplegia". In: *Lancet (London, England)* 381.9866, pp. 557–64. ISSN: 1474-547X. DOI: `10.1016/S0140-6736(12)61816-9`. URL: `http://dx.doi.org/10.1016/S0140-6736(12)61816-9%20http://www.ncbi.nlm.nih.gov/pubmed/23253623%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3641862`.

Courtine, G. (2005). "Performance of locomotion and foot grasping following a unilateral thoracic corticospinal tract lesion in monkeys (Macaca mulatta)". In: *Brain* 128.10, pp. 2338–2358. ISSN: 0006-8950. DOI: `10.1093/brain/awh604`. URL: `http://www.brain.oxfordjournals.org/cgi/doi/10.1093/brain/awh604`.

Cunningham, John P and Byron M Yu (Nov. 2014). "Dimensionality reduction for large-scale neural recordings". In: *Nature neuroscience* 17.11, pp. 1500–9. ISSN: 1546-1726. DOI: `10.1038/nn.3776`. URL: `http://www.nature.com/doifinder/10.1038/nn.3776%20http://www.ncbi.nlm.nih.gov/pubmed/25151264%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4433019`.

Deisseroth, Karl (Aug. 2015). *Optogenetics: 10 years of microbial opsins in neuroscience*. DOI: `10.1038/nn.4091`. URL: `/pmc/articles/PMC4790845/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4790845/`.

DiGiovanna, Jack et al. (Oct. 2016). "Engagement of the Rat Hindlimb Motor Cortex across Natural Locomotor Behaviors". In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 36.40, pp. 10440–10455. ISSN: 1529-2401. DOI: `10.1523/JNEUROSCI.4343-15.2016`. URL: `http://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.4343-15.2016%20http://www.ncbi.nlm.nih.gov/pubmed/27707977`.

Donati, Ana R. C. et al. (Aug. 2016). "Long-Term Training with a Brain-Machine Interface-Based Gait Protocol Induces Partial Neurological Recovery in Paraplegic Patients". In: *Scientific reports* 6.April, p. 30383. ISSN: 2045-2322. DOI: `10.1038/`

srep30383. URL: `http://www.nature.com/articles/srep30383%5Cnhttp://` `www.ncbi.nlm.nih.gov/pubmed/27513629%20http://www.ncbi.nlm.nih.gov/` `pubmed/27513629%20http://www.pubmedcentral.nih.gov/articlerender.` `fcgi?artid=PMC4980986`.

Drew, T. (1988). "Motor cortical cell discharge during voluntary gait modification". In: *Brain Research* 457.1, pp. 181–187. ISSN: 00068993. DOI: `10.1016/0006-` `8993(88)90073-X`.

Drew, Trevor, Jacques-Etienne Andujar, et al. (Jan. 2008). "Cortical mechanisms involved in visuomotor coordination during precision walking". In: *Brain research reviews* 57.1, pp. 199–211. ISSN: 0165-0173. DOI: `10.1016/j.brainresrev.` `2007.07.017`. URL: `http://linkinghub.elsevier.com/retrieve/pii/` `S0165017307001397%20http://www.ncbi.nlm.nih.gov/pubmed/17935789`.

Drew, Trevor and Simon Doucet (July 1991). "Application of circular statistics to the study of neuronal discharge during locomotion." In: *Journal of neuroscience methods* 38.2-3, pp. 171–81. ISSN: 0165-0270. DOI: `10.1016/0165-0270(91)` `90167-x`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/1784121`.

Drew, Trevor, Wan Jiang, and Witold Widajewicz (Oct. 2002). "Contributions of the motor cortex to the control of the hindlimbs during locomotion in the cat". In: *Brain research. Brain research reviews* 40.1-3, pp. 178–91. ISSN: 01650173. DOI: `10.1016/S0165-0173(02)00200-X`. URL: `http://www.ncbi.nlm.nih.gov/` `pubmed/12589916`.

Drew, Trevor and Daniel S. Marigold (2015). "Taking the next step: Cortical contributions to the control of locomotion". In: *Current Opinion in Neurobiology* 33, pp. 25–33. ISSN: 18736882. DOI: `10.1016/j.conb.2015.01.011`. URL: `http:` `//dx.doi.org/10.1016/j.conb.2015.01.011`.

Elsayed, Gamaleldin F and John P Cunningham (Sept. 2017). "Structure in neural population recordings: an expected byproduct of simpler phenomena?" In: *Nature*

neuroscience 20.9, pp. 1310–1318. ISSN: 1546-1726. DOI: `10.1038/nn.4617`. URL: `http://www.nature.com/doifinder/10.1038/nn.4617%20http://www.ncbi.nlm.nih.gov/pubmed/28783140%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5577566`.

Elsayed, Gamaleldin F. et al. (Dec. 2016). "Reorganization between preparatory and movement population responses in motor cortex". In: *Nature Communications* 7.1, pp. 1–15. ISSN: 20411723. DOI: `10.1038/ncomms13239`. URL: `www.nature.com/naturecommunications`.

Embers, Monica E. et al. (Apr. 2013). "Pharmacokinetic analysis of oral doxycycline in rhesus macaques". In: *Journal of Medical Primatology* 42.2, pp. 57–61. ISSN: 00472565. DOI: `10.1111/jmp.12031`. URL: `/pmc/articles/PMC3593953/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3593953/`.

Fedirchuk, B. et al. (1998). "Pharmacologically evoked fictive motor patterns in the acutely spinalized marmoset monkey (Callithrix jacchus)". In: *Experimental Brain Research* 122.3, pp. 351–361. ISSN: 00144819. DOI: `10.1007/s002210050523`. URL: `https://link.springer.com/article/10.1007/s002210050523`.

Fitzsimmons, Nathan A et al. (2009). "Extracting kinematic parameters for monkey bipedal walking from cortical neuronal ensemble activity". In: *Frontiers in integrative neuroscience* 3.13, p. 3. ISSN: 1662-5145. DOI: `10.3389/neuro.07.003.2009`. URL: `http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=19404411%20http://www.ncbi.nlm.nih.gov/pubmed/19404411%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2659168`.

Flesher, Sharlene N. et al. (Oct. 2016). "Intracortical microstimulation of human somatosensory cortex". In: *Science Translational Medicine* 8.361. ISSN: 19466242.

DOI: `10.1126/scitranslmed.aaf8083`. URL: `https://pubmed.ncbi.nlm.nih.gov/27738096/`.

FORSSBERG, H. et al. (1980). "The locomotion of the low spinal cat. II. Inter-limb coordination". In: *Acta Physiologica Scandinavica* 108.3, pp. 283–295. ISSN: 1365201X. DOI: `10.1111/j.1748-1716.1980.tb06534.x`. URL: `https://pubmed.ncbi.nlm.nih.gov/7376923/`.

Foster, Justin D et al. (Aug. 2014). "A freely-moving monkey treadmill model". In: *Journal of neural engineering* 11.4, p. 046020. ISSN: 1741-2552. DOI: `10.1088/1741-2560/11/4/046020`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/24995476`.

Friel, K M, T Drew, and J H Martin (May 2007). "Differential activity-dependent development of corticospinal control of movement and final limb position during visually guided locomotion." In: *Journal of neurophysiology* 97.5, pp. 3396–406. ISSN: 0022-3077. DOI: `10.1152/jn.00750.2006`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/17376849%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2740651`.

Frigon, Alain et al. (May 2013). "Split-belt walking alters the relationship between locomotor phases and cycle duration across speeds in intact and chronic spinalized adult cats". In: *Journal of Neuroscience* 33.19, pp. 8559–8566. ISSN: 02706474. DOI: `10.1523/JNEUROSCI.3931-12.2013`. URL: `https://www.jneurosci.org/content/33/19/8559%20https://www.jneurosci.org/content/33/19/8559.abstract`.

G raham Brown, By T (Dec. 1911). "The intrinsic factors in the act of progression in the mammal". In: *Proceedings of the Royal Society of London. Series B, Containing Papers of a Biological Character* 84.572, pp. 308–319. ISSN: 0950-1193. DOI: `10.1098/rspb.1911.0077`. URL: `https://royalsocietypublishing.org/doi/10.1098/rspb.1911.0077`.

Gallego, Juan A., Matthew G. Perich, Lee E. Miller, et al. (June 2017). "Neural Manifolds for the Control of Movement". In: *Neuron* 94.5, pp. 978–984. ISSN: 0896-6273. DOI: `10.1016/J.NEURON.2017.05.025`. URL: `https://www.sciencedirect.com/science/article/pii/S0896627317304634?via%3Dihub`.

Gallego, Juan A., Matthew G. Perich, Stephanie N. Naufel, et al. (Dec. 2018). "Cortical population activity within a preserved neural manifold underlies multiple motor behaviors". In: *Nature Communications* 9.1, p. 4233. ISSN: 2041-1723. DOI: `10.1038/s41467-018-06560-z`. URL: `http://www.nature.com/articles/s41467-018-06560-z`.

Ganguly, Karunesh and Jose M. Carmena (July 2009). "Emergence of a stable cortical map for neuroprosthetic control". In: *PLoS biology* 7.7, e1000153. ISSN: 1545-7885. DOI: `10.1371/journal.pbio.1000153`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/19621062%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2702684`.

Ganguly, Karunesh, Dragan F Dimitrov, et al. (May 2011). "Reversible large-scale modification of cortical networks during neuroprosthetic control". In: *Nature neuroscience* 14.5, pp. 662–7. ISSN: 1546-1726. DOI: `10.1038/nn.2797`. URL: `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3389499&tool=pmcentrez&rendertype=abstract%20http://www.ncbi.nlm.nih.gov/pubmed/21499255%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3389499`.

Georgopoulos, A P and S Grillner (Sept. 1989). "Visuomotor coordination in reaching and locomotion." In: *Science (New York, N.Y.)* 245.4923, pp. 1209–10. ISSN: 0036-8075. DOI: `10.1126/science.2675307`. URL: `http://www.sciencemag.org/cgi/doi/10.1126/science.2675307%20http://www.ncbi.nlm.nih.gov/pubmed/2675307`.

Gerasimenko, Yury P. et al. (Oct. 2006). "Spinal cord reflexes induced by epidural spinal cord stimulation in normal awake rats". In: *Journal of Neuroscience Methods* 157.2, pp. 253–263. ISSN: 01650270. DOI: `10.1016/j.jneumeth.2006.05.004`. URL: `https://pubmed.ncbi.nlm.nih.gov/16764937/`.

Gossen, Manfred and Hermann Bujard (1992). "Tight control of gene expression in mammalian cells by tetracycline-responsive promoters". In: *Proceedings of the National Academy of Sciences of the United States of America* 89.12, pp. 5547–5551. ISSN: 00278424. DOI: `10.1073/pnas.89.12.5547`. URL: `/pmc/articles/PMC49329/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC49329/`.

Grillner, Sten, A P Georgopoulos, and L M Jordan (1997). *Selection and initiation of motor behavior.* URL: `papers3://publication/uuid/D61F4D45-A92F-4427-B81C-F9870CAF5F00`.

Guertin, Pierre A. (2013). *Central pattern generator for locomotion: Anatomical, physiological, and pathophysiological considerations.* DOI: `10.3389/fneur.2012.00183`. URL: `/pmc/articles/PMC3567435/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3567435/`.

He, S Q, R P Dum, and P L Strick (Mar. 1993). "Topographic organization of corticospinal projections from the frontal lobe: motor areas on the lateral surface of the hemisphere." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 13.3, pp. 952–80. ISSN: 0270-6474. URL: `http://www.jneurosci.org/content/15/5/3284.long%5Cnpapers://9d5f8f0b-0f0b-459c-b067-2a21bce198bb/Paper/p2495%20http://www.ncbi.nlm.nih.gov/pubmed/7680069`.

Hioki, Hiroyuki et al. (2009). "High-level transgene expression in neurons by lentivirus with Tet-Off system". In: *Neuroscience Research* 63.2, pp. 149–154. ISSN: 01680102. DOI: `10.1016/j.neures.2008.10.010`.

Hochberg, Leigh R et al. (July 2006). "Neuronal ensemble control of prosthetic devices by a human with tetraplegia". In: *Nature* 442.7099, pp. 164–71. ISSN: 1476-4687. DOI: 10.1038/nature04970. URL: http://www.ncbi.nlm.nih.gov/pubmed/16838014.

Hochberg, Leigh R. et al. (May 2012). "Reach and grasp by people with tetraplegia using a neurally controlled robotic arm". In: *Nature* 485.7398, pp. 372–5. ISSN: 1476-4687. DOI: 10.1038/nature11076. URL: http://dx.doi.org/10.1038/nature11076%20http://www.ncbi.nlm.nih.gov/pubmed/22596161%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3640850.

Jarosiewicz, Beata et al. (Dec. 2008). "Functional network reorganization during learning in a brain-computer interface paradigm". In: *Proceedings of the National Academy of Sciences of the United States of America* 105.49, pp. 19486–19491. ISSN: 00278424. DOI: 10.1073/pnas.0808113105. URL: /pmc/articles/PMC2614787/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2614787/.

Kao, Jonathan C et al. (2015). "Single-trial dynamics of motor cortex and their applications to brain-machine interfaces". In: *Nature communications* 6.May, p. 7759. ISSN: 2041-1723. DOI: 10.1038/ncomms8759. URL: http://www.ncbi.nlm.nih.gov/pubmed/26220660.

Kato, Shigeki, Ken-ichi Inoue, et al. (2007). "Efficient gene transfer via retrograde transport in rodent and primate brains using a human immunodeficiency virus type 1-based vector pseudotyped with rabies virus glycoprotein." In: *Human gene therapy* 18.11, pp. 1141–1151. ISSN: 1043-0342. DOI: 10.1089/hum.2007.082.

Kato, Shigeki, Kenta Kobayashi, et al. (Feb. 2011). "A lentiviral strategy for highly efficient retrograde gene transfer by pseudotyping with fusion envelope glycoprotein". In: *Human Gene Therapy* 22.2, pp. 197–206. ISSN: 10430342. DOI: 10.1089/hum.2009.179. URL: www.liebertpub.com.

149

Kim, Sung-Phil et al. (2008a). "Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia". In: *Journal of Neural Engineering* 5.4, pp. 455–476. ISSN: 1741-2560. DOI: `10.1088/1741-2560/5/4/010`.

— (2008b). "Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia". In: *Journal of Neural Engineering* 5.4, pp. 455–476. ISSN: 1741-2560. DOI: `10.1088/1741-2560/5/4/010`. URL: `http://stacks.iop.org/1741-2552/5/i=4/a=010?key=crossref.6fe01f9b36a651dd79f75663970059ae`.

Kinoshita, Masaharu et al. (July 2012). "Genetic dissection of the circuit for hand dexterity in primates." In: *Nature* 487.7406, pp. 235–8. ISSN: 1476-4687. DOI: `10.1038/nature11206`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/22722837`.

Kobak, Dmitry et al. (Apr. 2016). "Demixed principal component analysis of neural population data". In: *eLife* 5.APRIL2016. ISSN: 2050084X. DOI: `10.7554/eLife.10989`.

Krouchev, Nedialko and T Drew (2013). "Motor cortical regulation of sparse synergies provides a framework for the flexible control of precision walking." In: *Frontiers in computational neuroscience* 7.July, p. 83. ISSN: 1662-5188. DOI: `10.3389/fncom.2013.00083`. URL: `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3708143&tool=pmcentrez&rendertype=abstract`.

Lafreniere-Roula, Myriam and David A. McCrea (Aug. 2005). "Deletions of rhythmic motoneuron activity during fictive locomotion and scratch provide clues to the organization of the mammalian central pattern generator". In: *Journal of Neurophysiology* 94.2, pp. 1120–1132. ISSN: 00223077. DOI: `10.1152/jn.00216.2005`. URL: `https://pubmed.ncbi.nlm.nih.gov/15872066/`.

Lara, A. H., J. P. Cunningham, and M. M. Churchland (Dec. 2018). "Different population dynamics in the supplementary motor area and motor cortex during reach-

ing". In: *Nature Communications* 9.1, p. 2754. ISSN: 2041-1723. DOI: `10.1038/s41467-018-05146-z`. URL: `http://www.nature.com/articles/s41467-018-05146-z`.

Laurens, Jean, Hui Meng, and Dora E Angelaki (Nov. 2013). "Computation of linear acceleration through an internal model in the macaque cerebellum". In: *Nature neuroscience* 16.11, pp. 1701–8. ISSN: 1546-1726. DOI: `10.1038/nn.3530`. URL: `http://www.nature.com/doifinder/10.1038/nn.3530%20http://www.ncbi.nlm.nih.gov/pubmed/24655651%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4126411%20http://www.ncbi.nlm.nih.gov/pubmed/24077562%20http://www.pubmedcentral.nih.gov/articlerender.fcg`.

Lemon, Roger N. et al. (July 2012). "Lawrence and Kuypers (1968a, b) revisited: copies of the original filmed material from their classic papers in Brain". In: *Brain* 135.7, pp. 2290–2295. ISSN: 1460-2156. DOI: `10.1093/brain/aws037`. URL: `https://academic.oup.com/brain/article-lookup/doi/10.1093/brain/aws037`.

Lonser, Russell R. et al. (1998). "Direct convective delivery of macromolecules to the spinal cord". In: *Journal of Neurosurgery* 89.4, pp. 616–622. ISSN: 00223085. DOI: `10.3171/jns.1998.89.4.0616`. URL: `https://pubmed.ncbi.nlm.nih.gov/9761056/`.

Lundberg, A. and C. G. Phillips (1973). "T. Graham Brown's film on locomotion in the decerebrate cat." In: *The Journal of physiology* 231.2, 90P–91P. ISSN: 0022-3751. URL: `https://pubmed.ncbi.nlm.nih.gov/4720963/`.

Macke, J H et al. (2011). "Empirical models of spiking in neural populations". In: *Advances in Neural Information Processing Systems 24: 25th Conference on Neural Information Processing Systems*, pp. 1350–1358. URL: `http://discovery.ucl.ac.uk/1328983/`.

McCrea, David A. and Ilya A. Rybak (Jan. 2008). "Organization of mammalian loco-motor rhythm and pattern generation". In: *Brain Research Reviews* 57.1, pp. 134–146. ISSN: 01650173. DOI: `10.1016/j.brainresrev.2007.08.006`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0165017307001798`.

Middleton, Frank A. and Peter L. Srrick (1997). "Cerebellar output channels". In: *International Review of Neurobiology* 41, pp. 61–82. ISSN: 00747742. DOI: `10.1016/s0074-7742(08)60347-5`. URL: `https://pubmed.ncbi.nlm.nih.gov/9378611/`.

Miller, S. and J. Van der Burg (1973). "The Function of Long Propriospinal Pathways in the Co-Ordination of Quadrupedal Stepping in the Cat". In: Springer, Boston, MA, pp. 561–577. DOI: `10.1007/978-1-4613-4547-3{\_}44`. URL: `https://link.springer.com/chapter/10.1007/978-1-4613-4547-3_44`.

Miri, Andrew et al. (Aug. 2017). "Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex." In: *Neuron* 95.3, pp. 683–696. ISSN: 1097-4199. DOI: `10.1016/j.neuron.2017.06.042`. URL: `http://linkinghub.elsevier.com/retrieve/pii/S0896627317305949%20http://www.ncbi.nlm.nih.gov/pubmed/28735748`.

Morizawa, Yosuke M. et al. (Dec. 2017). "Reactive astrocytes function as phagocytes after brain ischemia via ABCA1-mediated pathway". In: *Nature Communications* 8.1. ISSN: 20411723. DOI: `10.1038/s41467-017-00037-1`. URL: `https://pubmed.ncbi.nlm.nih.gov/28642575/`.

Naldini, Luigi et al. (1996). "In vivo gene delivery and stable transduction of post mitotic cells by a lentiviral vector". In: *Science* 272.5259, pp. 263–267. URL: `http://science.sciencemag.org/content/272/5259/263`.

Orlovsky, Grigori, T. G. Deliagina, and Sten Grillner (Sept. 1999). *Neuronal Control of LocomotionFrom Mollusc to Man.* Oxford University Press. ISBN: 9780198524052. DOI: `10.1093/acprof:oso/9780198524052.001.0001`. URL: `https://oxford.`

universitypressscholarship.com/view/10.1093/acprof:oso/9780198524052.
001.0001/acprof-9780198524052.

Pandarinath, Chethan et al. (Oct. 2018a). "Inferring single-trial neural population dynamics using sequential auto-encoders". In: *Nature Methods* 15.10, pp. 805–815. ISSN: 15487105. DOI: 10.1038/s41592-018-0109-9. URL: http://www.nature.com/articles/s41592-018-0109-9.

— (Oct. 2018b). "Inferring single-trial neural population dynamics using sequential auto-encoders". In: *Nature Methods* 15.10, pp. 805–815. ISSN: 1548-7091. DOI: 10.1038/s41592-018-0109-9. URL: http://www.nature.com/articles/s41592-018-0109-9.

Pearson, Keir and Rod Gramlich (2010). "Updating neural representations of objects during walking." In: *Annals of the New York Academy of Sciences* 1198, pp. 1–9. ISSN: 1749-6632. DOI: 10.1111/j.1749-6632.2009.05422.x. URL: http://www.ncbi.nlm.nih.gov/pubmed/20536915.

Perret, Claude and Jean Marie Cabelguen (Apr. 1980). "Main characteristics of the hindlimb locomotor cycle in the decorticate cat with special reference to bifunctional muscles". In: *Brain Research* 187.2, pp. 333–352. ISSN: 00068993. DOI: 10.1016/0006-8993(80)90207-3.

Prilutsky, Boris I et al. (2005). "Quantification of Motor Cortex Activity and Full-Body Biomechanics During Unconstrained Locomotion Quantification of Motor Cortex Activity and Full-Body Biomechanics During Unconstrained Locomotion". In: *Journal of Neurophysiology*, pp. 2959–2969. ISSN: 0022-3077. DOI: 10.1152/jn.00704.2004.

Quiroga, R. Quian, Z. Nadasdy, and Y. Ben-Shaul (Aug. 2004). "Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering". In: *Neural Computation* 16.8, pp. 1661–1687. ISSN: 08997667. DOI: 10.1162/089976604774201631. URL: https://pubmed.ncbi.nlm.nih.gov/15228749/.

Raposo, David, Matthew T. Kaufman, and Anne K. Churchland (Jan. 2014). "A category-free neural population supports evolving demands during decision-making". In: *Nature Neuroscience* 17.12, pp. 1784–1792. ISSN: 15461726. DOI: `10.1038/nn.3865`. URL: `https://pubmed.ncbi.nlm.nih.gov/25383902/`.

Rigosa, J et al. (Oct. 2015). "Decoding bipedal locomotion from the rat sensorimotor cortex". In: *Journal of neural engineering* 12.5, p. 056014. ISSN: 1741-2552. DOI: `10.1088/1741-2560/12/5/056014`. URL: `http://stacks.iop.org/1741-2552/12/i=5/a=056014?key=crossref.ebf70cb2619cded3198a1ff0558020af%20http://www.ncbi.nlm.nih.gov/pubmed/26331532`.

Rivers, T. J. et al. (2014). "Gaze shifts and fixations dominate gaze behavior of walking cats". In: *Neuroscience* 275, pp. 477–499. ISSN: 18737544. DOI: `10.1016/j.neuroscience.2014.06.034`. URL: `http://dx.doi.org/10.1016/j.neuroscience.2014.06.034`.

Russo, Abigail A. et al. (Feb. 2018). "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response". In: *Neuron* 97.4, pp. 953–966. ISSN: 10974199. DOI: `10.1016/j.neuron.2018.01.004`. URL: `/pmc/articles/PMC5823788/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5823788/`.

Sadtler, Patrick T. et al. (2014). "Neural constraints on learning". In: *Nature* 512.7515, pp. 423–426. ISSN: 0028-0836. DOI: `10.1038/nature13665`. URL: `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4393644&tool=pmcentrez&rendertype=abstract`.

Sathyamurthy, Anupama et al. (Feb. 2018). "Massively Parallel Single Nucleus Transcriptional Profiling Defines Spinal Cord Neurons and Their Activity during Behavior". In: *Cell Reports* 22.8, pp. 2216–2225. ISSN: 22111247. DOI: `10.1016/j.celrep.2018.02.003`.

Serruya, Mijail D et al. (Mar. 2002). "Instant neural control of a movement signal." In: *Nature* 416.6877, pp. 141–2. ISSN: 0028-0836. DOI: `10.1038/416141a`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/11894084`.

Sherrington, C. S. (Mar. 1906). "Observations on the scratch-reflex in the spinal dog". In: *The Journal of Physiology* 34.1-2, pp. 1–50. ISSN: 14697793. DOI: `10.1113/jphysiol.1906.sp001139`. URL: `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1465804/`.

Shik, M. L. and G. N. Orlovsky (1976). *Neurophysiology of locomotor automatism.* DOI: `10.1152/physrev.1976.56.3.465`. URL: `https://pubmed.ncbi.nlm.nih.gov/778867/`.

Song, Weiguo et al. (Nov. 2009). "Multiple types of movement-related information encoded in hindlimb/trunk cortex in rats and potentially available for brain-machine interface controls." In: *IEEE transactions on bio-medical engineering* 56.11 Pt 2, pp. 2712–6. ISSN: 1558-2531. DOI: `10.1109/TBME.2009.2026284`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/19605313%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2883457`.

Srivastava, Kyle H et al. (Jan. 2017). "Motor control by precisely timed spike patterns." In: *Proceedings of the National Academy of Sciences of the United States of America* 114.5, pp. 1171–1176. ISSN: 1091-6490. DOI: `10.1073/pnas.1611734114`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/28100491%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5293088`.

Stavisky, Sergey D. et al. (2017). "Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions". In: *Neuron.* ISSN: 10974199. DOI: `10.1016/j.neuron.2017.05.023`.

Stuart, Douglas G. and Hans Hultborn (2008). "Thomas Graham Brown (1882-1965), Anders Lundberg (1920-), and the neural control of stepping". In: *Brain Research*

Reviews 59.1, pp. 74–95. ISSN: 01650173. DOI: `10.1016/j.brainresrev.2008.`
`06.001`. URL: `http://dx.doi.org/10.1016/j.brainresrev.2008.06.001`.

Taylor, Dawn M, Stephen I Helms Tillery, and Andrew B Schwartz (2002). "Direct
Cortical Control of 3D Neuroprosthetic Devices". In: 1829.2002, pp. 3–7. DOI:
`10.1126/science.1070291`.

Truccolo, Wilson, Uri T Eden, et al. (Feb. 2005). "A point process framework for
relating neural spiking activity to spiking history, neural ensemble, and extrinsic
covariate effects". In: *Journal of neurophysiology* 93.2, pp. 1074–89. ISSN: 0022-
3077. DOI: `10.1152/jn.00697.2004`. URL: `http://courses.cs.washington.`
`edu/courses/cse528/13sp/truccolo05.pdf%20http://www.ncbi.nlm.nih.`
`gov/pubmed/15356183`.

Truccolo, Wilson, Leigh R Hochberg, and John P Donoghue (Jan. 2010). "Collective
dynamics in human and monkey sensorimotor cortex: predicting single neuron
spikes". In: *Nature neuroscience* 13.1, pp. 105–11. ISSN: 1546-1726. DOI: `10.1038/`
`nn.2455`. URL: `http://www.nature.com/doifinder/10.1038/nn.2455%20http:`
`//www.ncbi.nlm.nih.gov/pubmed/19966837%20http://www.pubmedcentral.`
`nih.gov/articlerender.fcgi?artid=PMC2820252`.

Vargas-Irwin, Carlos E et al. (Jan. 2015). "Spike train SIMilarity Space (SSIMS): a
framework for single neuron and ensemble data analysis". In: *Neural computation*
27.1, pp. 1–31. ISSN: 1530-888X. DOI: `10.1162/NECO{\_}a{\_}00684`. URL: `http:`
`//arxiv.org/abs/1512.05463%20http://www.ncbi.nlm.nih.gov/pubmed/`
`25602775%20http://www.ncbi.nlm.nih.gov/pubmed/25380335%20http:`
`//www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4276520`.

Wakida, Nicole M. et al. (Apr. 2018). "Phagocytic response of astrocytes to damaged
neighboring cells". In: *PLoS ONE* 13.4. ISSN: 19326203. DOI: `10.1371/journal.`
`pone.0196153`. URL: `/pmc/articles/PMC5927416/?report=abstract%20https:`
`//www.ncbi.nlm.nih.gov/pmc/articles/PMC5927416/`.

Williams, Alex H. et al. (June 2018). "Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis". In: *Neuron* 98.6, pp. 1099–1115. ISSN: 0896-6273. DOI: `10.1016/J.NEURON.2018.05.015`. URL: `https://www.sciencedirect.com/science/article/pii/S0896627318303878?via%3Dihub`.

Wodlinger, B et al. (Feb. 2015). "Ten-dimensional anthropomorphic arm control in a human brain-machine interface: difficulties, solutions, and limitations". In: *Journal of neural engineering* 12.1, p. 016011. ISSN: 1741-2552. DOI: `10.1088/1741-2560/12/1/016011`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/25514320%5Cnhttp://stacks.iop.org/1741-2552/12/i=1/a=016011?key=crossref.5778415d2c773d0c18afd19da30b4962%20http://www.ncbi.nlm.nih.gov/pubmed/25514320`.

Xing, David et al. (Oct. 2019). "Low-Dimensional Motor Cortex Dynamics Preserve Kinematics Information During Unconstrained Locomotion in Nonhuman Primates". In: *Frontiers in Neuroscience* 13, p. 1046. ISSN: 1662-453X. DOI: `10.3389/fnins.2019.01046`. URL: `https://www.frontiersin.org/articles/10.3389/fnins.2019.01046/full`.

Yakovenko, Sergiy and Trevor Drew (2009). "A Motor Cortical Contribution to the Anticipatory Postural Adjustments That Precede Reaching in the Cat". In: 2007, pp. 853–874. DOI: `10.1152/jn.00042.2009.`.

— (2015). "Similar Motor Cortical Control Mechanisms for Precise Limb Control during Reaching and Locomotion". In: *Journal of Neuroscience* 35.43, pp. 14476–14490. ISSN: 0270-6474. DOI: `10.1523/JNEUROSCI.1908-15.2015`. URL: `http://www.jneurosci.org/content/35/43/14476%5Cnhttp://www.jneurosci.org/content/35/43/14476.full%5Cnhttp://www.jneurosci.org/content/35/43/14476.full.pdf%5Cnhttp://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.1908-15.2015`.

Yamamoto, Mutsuya et al. (July 2003). "Reversible suppression of glutamatergic neurotransmission of cerebellar granule cells in vivo by genetically manipulated expression of tetanus neurotoxin light chain". In: *Journal of Neuroscience* 23.17, pp. 6759–6767. ISSN: 02706474. DOI: `10.1523/jneurosci.23-17-06759.2003`. URL: `/pmc/articles/PMC6740723/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6740723/`.

Yazdan-Shahmorad, Azadeh et al. (Mar. 2016). "A Large-Scale Interface for Optogenetic Stimulation and Recording in Nonhuman Primates". In: *Neuron* 89.5, pp. 927–939. ISSN: 10974199. DOI: `10.1016/j.neuron.2016.01.013`. URL: `https://pubmed.ncbi.nlm.nih.gov/26875625/`.

Yin, Ming et al. (Dec. 2014). "Wireless neurosensor for full-spectrum electrophysiology recordings during free behavior". In: *Neuron* 84.6, pp. 1170–82. ISSN: 1097-4199. DOI: `10.1016/j.neuron.2014.11.010`. URL: `http://www.ncbi.nlm.nih.gov/pubmed/25482026`.

Yu, Byron M et al. (July 2009). "Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity". In: *Journal of neurophysiology* 102.1, pp. 614–35. ISSN: 0022-3077. DOI: `10.1152/jn.90941.2008`. URL: `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2712272&tool=pmcentrez&rendertype=abstract`.

Zelenin, Pavel V et al. (2011). "Contribution of different limb controllers to modulation of motor cortex neurons during locomotion." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 31.12, pp. 4636–4649. ISSN: 0270-6474. DOI: `10.1523/JNEUROSCI.6511-10.2011`.

Zhou, X. et al. (Oct. 2006). "Optimization of the Tet-On system for regulated gene expression through viral evolution". In: *Gene Therapy* 13.19, pp. 1382–1390. ISSN: 09697128. DOI: `10.1038/sj.gt.3302780`. URL: `https://pubmed.ncbi.nlm.nih.gov/16724096/`.