

Using neural-machine interfaces to explore skill learning in central and peripheral nervous systems

Paul Botros



Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2022-191

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2022/EECS-2022-191.html>

August 10, 2022

Copyright © 2022, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Using neural-machine interfaces to explore skill learning in central and peripheral nervous
systems

by

Paul Abraham Botros

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

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jose M. Carmena, Co-chair

Professor Jan Rabaey, Co-chair

Professor David Foster

Professor Chunlei Liu

Summer 2022

Abstract

Using neural-machine interfaces to explore skill learning in central and peripheral nervous systems

by

Paul Abraham Botros

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Jose M. Carmena, Co-chair

Professor Jan Rabaey, Co-chair

Understanding the process of how the brain learns a new skill, whether a motor skill such as riding a bike or a cognitive skill such as playing chess, is tantamount to understanding dysfunction in this process during neurological disease. For instance, patients with Parkinson's disease are less able to initiate and/or learn new movements, and obsessive-compulsive disorder (OCD) is theorized to stem from dysfunctions in circuits connecting behavior and reward. Elucidation of these neural circuits can inform nascent stimulation-based therapies and improve the next generation of pharmacological and behavioral interventions for many neurological diseases. Furthermore, outside of the brain, understanding the capabilities of the peripheral nervous system to accommodate new skills can contribute to efforts attempting to restore lost motor function to amputees and paralyzed patients. Throughout this thesis, we will explore the neural basis of skill learning across both the central and peripheral nervous systems by utilizing neural-machine interfaces, which link neural activity to effectors such as a computer cursor or a prosthetic device and thus provide the experimenter a principled modality in which to study the learning of a new skill.

The process of skill learning can span a wide timeframe, but this thesis focuses on the "early" part of learning, in which an animal initially acquires and begins to refine a skill. Here, we define a skill as any set of intended behaviors – whether covert or overt – that improve in consistency over time as a result of training and reward. Within this thesis, we'll first explore which skills can be learned in the human peripheral nervous system by investigating the level of flexibility latent in our muscles, with an eye towards designing a high-performance, non-invasive neural-machine interface. In particular, classical experiments suggest motor unit activity within a single muscle lies along a low-dimensional manifold due to a common descending input, theoretically severely constraining the throughput of an assistive device utilizing motor unit activity. However, by utilizing a neural-machine interface paradigm across six days of training in

eight adults, we found that motor unit dimensionality is higher than previously theorized, showing promise that a motor-unit-based neural machine interface could provide clinical benefit.

Second, we'll ask the "how" of skill learning: how do we learn skills, whether motor or cognitive, in the brain? Here, we focus on studying a core mechanism of learning, known as "credit assignment", which enables activity in the cortex to become more stereotyped over training. Utilizing a neural-machine interface paradigm in rats, we investigate the role of the striatum, an area well-known for housing associations between behavior and reward, in credit assignment. We found that activity of neurons in the striatum form an internal model of cortical activity, continuously estimating both the proximity of current cortical activity to reward and the change in this proximity, with this model of cortical activity differing in complementary yet distinct ways between dorsal and ventral striatum. Such a model could be critical to implementing credit assignment in the brain.

Finally, we'll conclude this thesis on a practical note by discussing the closed-loop, performant software built for the aforementioned studies that have largely been open-sourced to benefit the broader neuroscience community. Neuroscience is an increasingly complex, performance-sensitive discipline, with neural data volume exploding over the past few decades. Software developed throughout this thesis work empowers researchers to build more flexible, more scalable systems with relative ease.

Therefore, by utilizing neural-machine interfaces to interrogate the functionality of the nervous system, the results described in this thesis contribute to our understanding of skill learning, demonstrating novel levels of flexibility in the peripheral nervous system with applications in clinical translation and suggesting a critical role of the striatum in skill learning as a part of credit assignment, a core neural mechanism of learning.

Contents

Contents	i
Acknowledgements	ii
Chapter 1 Introduction	1
Neural-Machine Interfaces as a Clinical Tool	2
Brain-Machine Interfaces as a Neuroscience Tool for Exploring Skill Learning	6
Scalable Software for Neuroscience	8
Chapter 2 Skilled independent control of individual motor units via a non-invasive neuromuscular-machine interface	10
2.1 Introduction	10
2.2 Results	12
2.3 Discussion	27
2.4 Methods	34
2.5 Supplementary Material	47
Chapter 3 A neural model of proximity to reward	50
3.1 Main	50
3.2 Materials and Methods	60
3.3 Supplementary Figures and Data	75
Chapter 4 Scalable Software for Neuroscience	83
4.1 River: a cross-platform, high-throughput, structured streaming framework	83
4.2 DuraPy: a service container for service-oriented neuroscience experiments	85
Chapter 5 Conclusion	89
5.1 Future Work	89
Chapter 6 Bibliography	92

Acknowledgements

To Jose: thank you for taking a chance on someone coming from a non-traditional background and with limited research experience, for conveying your visions with contagious enthusiasm, and for connecting me with talented mentors to guide my PhD journey.

To Ema: thank you for your mentorship, patience, and wisdom. Whether you know it or not, you taught me the value of precision and thoroughness; almost everything I've learned during this PhD has come from you or our work together. I can confidently say without your guidance my PhD would have been drastically worse off.

To Nuria: your perseverance and work ethic – especially as a simultaneous researcher and mother – are truly admirable. It's been my pleasure to work with you these past couple years, and I've learned a lot from your guidance.

To former and current labmates, Ellen, Gabby, David, Tanner, Albert, Maki, Will: thank you all for your camaraderie, bright attitudes, and hangouts over the years. To Ellen and Gabby in particular: thanks for the company during the long hours in the basement, and cheers to us for sticking it out until the very end.

To the Berkeley veterinarians and the veterinarian staff: thank you for your hard work and (sometimes last-minute) accommodations over these years. To Li Ka Shing staff: thank you for keeping everything running smoothly.

And of course, to my partner, Brittany. You have made this grueling and isolating PhD a more bearable journey, injecting bits of sunshine into my life when I needed it the most. Thank you for regularly commiserating with me through both the good and bad, and for reminding me what's important. And to Brittany, my friends, and my family: thank you for all for having patience with me as I finished this journey, and providing me perspective, motivation, and hope.

Chapter 1

Introduction

Watching the fingers of a virtuoso pianist fly amongst their 88 keys is, at first, almost an unbelievable experience. The spatial and temporal precision with which keys are played seemingly contrasts with the often effortless demeanor exuded by the pianist, who often does not visually focus on the keys. Decades of training affords the pianist this ability: to consistently and accurately recall intricate bimanual motor sequences. However, this pianist was once naive to their skill, begging the question: how did this person learn this skill? More specifically, what changes in the nervous system occurred as they progressed from naive to virtuoso?

Understanding the process of skill learning is tantamount to understanding and treating its dysfunction during neurological disease. For instance, patients with Parkinson's disease are less capable of initiating existing movements and learning new movements (Graybiel and Grafton 2015), and, among other psychological disorders, obsessive-compulsive disorder (OCD) is theorized to stem from dysfunctions in circuits connecting behavior and reward (Burguière et al. 2015). In addition to revealing the neurophysiological basis of existing pharmacological and behavioral treatments of these disorders, elucidation of these neural circuits can inform nascent stimulation-based therapies that utilize electrical stimulation to selectively activate and deactivate relevant neural circuits leading to behavioral dysfunction. Such stimulation-based therapies, including deep brain stimulation (DBS), have been popular for treatment of Parkinson's disease for decades, but have recently shown early promise for treatment of anxiety, OCD, and PTSD (Shanечи 2019) as our understanding of relevant neural circuits improves. In addition, outside of the brain, understanding the capabilities of the peripheral nervous system to accommodate new skills can contribute to efforts attempting to restore lost motor function to amputees and paralyzed patients.

Of the many possible aspects of this learning process to interrogate, this thesis focuses on the "early" part of learning, in which an animal initially acquires and begins to refine a skill. Here, we define a skill as any set of intended behaviors – whether covert or overt – that improve in consistency over time as a result of training and reward (Krakauer and Shadmehr 2006). Notably, we focus on the phase of skill learning that is goal-directed, in which the behavior is dependent on perceived reward value (Graybiel and Grafton 2015). This contrasts with the later phase of learning, habitual learning, in which behavior becomes more automatic (and likely more so the type of learning that underlies the virtuoso pianist's playing) (Graybiel and Grafton 2015).

Within this thesis, we'll first explore which skills can be learned in the human peripheral nervous system by investigating the level of flexibility latent in our muscles, with an eye towards designing a high-performance, non-invasive assistive device. In particular, we explored the

dimensionality of individual motor units within a single muscle, hypothesizing that such dimensionality might be higher than many classical motor experiments suggest. Second, we'll ask "how": how do we learn skills, whether motor or cognitive, in the brain? Here, we focus on studying a core mechanism of learning, known as "credit assignment", which enables activity in the cortex to become more stereotyped over training. We investigate the role of the striatum, an area well-known for housing associations between behavior and reward, in representing cortical activity in a manner conducive to credit assignment. Finally, we'll conclude this thesis on a practical note by discussing the closed-loop, performant software built for the aforementioned studies that have largely been open-sourced to benefit the broader neuroscience community. Neuroscience is an increasingly complex, performance-sensitive discipline, with neural data volume exploding over the past few decades, and software developed throughout this thesis work can empower researchers to build more flexible, scalable systems with relative ease.

Neural-Machine Interfaces as a Clinical Tool

All research performed in this thesis was united in methodology: through the use of neural-machine interfaces. Neural-machine interfaces create a direct control pathway between neural activity and an external actuator via a decoder (Shanechi 2019) (Figure 1.1). The external actuator can be physical, such as an exoskeleton or robotic arm, or virtual, such as a computer cursor. The subject then utilizes sensory and task feedback to adapt their neural activity to the decoder in order to optimize performance (Orsborn and Pesaran 2017), forming a closed-loop system. Neural-machine interfaces that source neural activity input directly from the brain are commonly referred to as brain-machine interfaces (BMIs) (Shanechi 2019), and we'll denote those that source neural activity from the peripheral nervous system as neuromuscular-machine interfaces (NMIs).




Figure 1.1. Figure adapted from (Shanechi 2019). Brain-machine interfaces record neural activity directly from the brain via invasive or non-invasive modalities and feed neural activity into a decoder, which is responsible for translating input neural activity to output commands to an actuator. Sensory feedback, e.g. visual feedback of a computer cursor on a screen, closes the learning loop, allowing the brain to adapt its neural activity to better control the actuator.

Historically, BMIs utilizing invasive recordings of the brain have demonstrated great clinical promise for restoring lost function to patients with neurological disorders, most commonly targeting paralyzed patients who have lost significant motor function. These BMIs most commonly utilize recordings of individual or small groups of neurons at a high temporal resolution. Proof-of-concept BMIs have restored independence in people with severe paralysis, enabling tetraplegic people to control robotic arms and exoskeletons (Collinger, Wodlinger, et al. 2013; Benabid et al. 2019), navigate computers (Hochberg et al. 2006), and even regain control of their own paralyzed limbs through electrical stimulation (Bouton et al. 2016). Representing state-of-the-art performance, a recent BMI study even demonstrated an impressive communication throughput of 90 characters per minute via a handwriting decoding approach (Willett et al. 2021). BMIs utilizing epidural brain signals, which are more stable than subdural, single-neuron recordings, enabled a tetraplegic patient to control an exoskeleton (Benabid et al. 2019). Unfortunately, despite demonstrated clinical promise, BMIs have not yet enjoyed widespread usage. Surgical risk,

implant aesthetics, and usability concerns have been cited as reasons for their unpopularity, even for patients that are severely paralyzed (Blabe et al. 2015).

Alternatively, BMIs can also be designed using entirely non-invasive recording modalities, most commonly utilizing the electroencephalogram (EEG). Non-invasive BMIs work similarly in principle to their invasive counterparts in which neural activity is transformed via a decoder into actuator commands; however, unlike invasive modalities, recorded neural activity consists of averages across millions of neurons due to spatiotemporal filtering by the skull and scalp (Grover and Venkatesh 2017). EEG BMIs have demonstrated 1 (N. A. Bhagat et al. 2016), and with ample training 2 (Wolpaw and McFarland 2004), controllable degrees of freedom and have enabled patients to communicate digitally at moderate speeds (Nakanishi et al. 2018) and steer robots (Tonin and Millán 2021). However, the severely compromised signal-to-noise ratio of non-invasive modalities poses significant challenges to building a performant BMI, hindering throughput and ease-of-use for patients (Millán and Carmena 2010). Other non-invasive modalities besides EEG, such as fMRI or magnetoencephalography (MEG), similarly suffer from either poor temporal or spatial resolution when recording brain activity.

While there appears to be a tradeoff between performance and invasiveness within the BMI field, the brain is not the only access point for the nervous system. As we will explore in more detail in Chapter 2, activity in muscles can be interpreted as the direct output of the spinal cord, serving as an alternative access point for the nervous system (Farina et al. 2017). Muscle activity arises when motor neurons descending via the spinal cord receive sufficient upstream activation, either through direct cortical control or summed activity within the spinal cord itself (Kandel 2013). Once sufficiently depolarized, a motor neuron fires, and its action potential propagates to the various muscle fibers it innervates via the neuromuscular junction. Neuromuscular junctions typically have large “safety factors”, in that the resulting end-plate potential (EPP; neuromuscular analog to an excitatory postsynaptic potential in nerves) tends to be significantly larger than the minimum required to depolarize a muscle fiber, and as a result, the neuromuscular junction reliably propagates action potentials. Combined with the fact that a given muscle fiber is only innervated by a single motor neuron and that the action potentials in these muscle fibers occur approximately at the same time, a motor neuron and its innervated muscle fibers are known as a “motor unit”, where a motor neuron and its innervated muscle fibers are all typically envisioned as coincidentally activating as one (Kandel 2013).

Electrodes placed on the surface of the skin can record electrical activity arising from motor unit activation, known as electromyography (EMG). EMG signals are relatively high signal-to-noise ratio as compared to non-invasive recordings of the brain, as the intervening tissue between the muscle and recording grids does not introduce as significant of a distortion as the skull during EEG recordings. The coincident activation of the various innervated muscle fibers for each motor neuron creates a unique spatiotemporal footprint for each motor unit, enabling blind-source separation techniques to identify activations of individual motor units (Negro et al. 2016) in real-time (Barsakcioglu et al. 2020).

Neuromuscular-machine interfaces (NMIs) can exploit muscle activity to allow disabled patients to control external devices. Classic NMI design utilizes spatially-broad “global” EMG activity, where each channel of data originates from a large electrode and thus represents the average muscle activity across a swath of motor units in a single muscle. Typically feeding global EMG across multiple muscles into a biomimetic decoder, this paradigm enabled intuitive control of external effectors, such as hand prostheses for amputees (Zhuang et al. 2019). Recent technological advances allowed NMIs to utilize real-time activity of individual motor units rather than global EMG activity, providing greater spatial selectivity and consequently a finer-grained level of control (Farina et al. 2017; Barsakcioglu et al. 2020). A comparison of spatial and temporal resolutions of various neural recording modalities can be seen in Figure 1.2.




Figure 1.2. Illustration of spatial and temporal resolutions of various neural recording modalities. Invasive, high-performance brain-machine interfaces utilize recording modalities with high spatial and temporal resolution, whereas traditional non-invasive modalities such as EEG suffer in either spatial or temporal resolution, resulting in decreased performance. EMG (bolded) is the modality under study in Chapter 2 of this thesis, as it theoretically retains both high spatial and temporal resolution while remaining non-invasive.

However, even if utilizing spatially-specific individual motor unit activity, NMIs assume that the maximum dimensionality of control extracted from a single muscle is bounded by the number of mechanical degrees of freedom of the muscle, which commonly is just 1 (De Luca and Erim 1994). This assumption stems from prevailing theories that muscles are driven by a low-dimensional “common drive”: descending cortical or spinal input that activates motor units of a

given muscle simultaneously (Fuglevand et al. 1993). If true, this would severely constrain the dimensionality of control, especially in cases where the overall number of muscles available for control is limited (i.e. tetraplegics), and consequently constrain the usefulness of an NMI to a clinical population. Chapter 2 of this thesis tests this prevailing theory thoroughly, concluding that motor unit dimensionality may in fact be higher than that predicted by the “common drive” theory, thus providing promising evidence towards the use of NMIs as a clinical device.

Brain-Machine Interfaces as a Neuroscience Tool for Exploring Skill Learning

While brain-machine interfaces were originally ideated for use as an assistive device for patients, recent approaches in the past decade have utilized BMIs as a tool to study neural circuits themselves. By creating a link between neural activity and an external effector, BMIs effectively appear to the brain as yet another tool to be learned, with neural representations of natural motor skills appearing remarkably similar to those of neuroprosthetic skill (Ganguly and Carmena 2009; Peters et al. 2014; Athalye et al. 2017). With more classical experimental paradigms, one problem when studying neural circuits underlying motor skills is that the experimenter does not know how the cortical activity under study relates to some given behavior; for example, as a rat reaches for a pellet of food, it is undefined how exactly some given set of neurons in motor cortex under study might relate to the dynamics of that reach, and how that relationship might vary trial-by-trial or day-to-day. Utilizing a BMI, the experimenter can instead explicitly define the relationship between cortical activity and the task through the decoder, establishing a causal link between the cortical activity under study and behavior (Athalye et al. 2020). Therefore, since it both appears similar to motor skill learning and affords the experimenter additional control over their study, BMI experiments can be utilized to more precisely interrogate the role of neural circuits during learning (Gulati et al. 2017; Athalye et al. 2018; Sadtler et al. 2014) and consequently make general conclusions about skill learning.

Recent studies have utilized BMI experiments to study a particular mechanism underlying skill learning, known as “credit assignment” (Athalye et al. 2020). Credit assignment is the process by which the brain biases particular sets of neurons in order to change behavior to increase the rate of reward, and is one key mechanism underlying reinforcement learning (Sutton and Barto 2018). By changing either the underlying network connectivity of the neurons or by modifying the inputs to the neurons, the brain can theoretically enable neural activity to be more consistent from trial to trial, which, assuming that reinforced neural activity results in successful behavior, would enable increasing success rates (Athalye et al. 2020).

However, despite it playing a fundamental role in learning, the neural basis of credit assignment is poorly understood. In particular, the challenge of implementing a credit assignment mechanism is often split into two aspects. First, there exists a temporal aspect: reward is often delayed relative to the behavior that led to reward, where this delay can span a wide temporal range – for example, reaching to grab a pellet of food separates behavior and reward by a few seconds,

but the moves taken in a chess game eventually leading to a win has delays spanning hours. In computational models of reinforcement learning, temporal credit assignment is often solved via “value functions”, which are continuously updateable functions that estimate the current expected value of the animal’s state or actions (Sutton and Barto 2018). Value functions provide continuous feedback to the animal as it performs actions in a task, where the animal tends towards taking actions that increase its expected reward according to these functions. While these value functions are initially noisy and inaccurate, when a (temporally delayed) reward is achieved, these value functions are appropriately updated to increase the value of recent actions or states, thus putatively increasing the accuracy of the value functions. Subsequently, via continuous improvements to value functions, the behavior of the animal should eventually increase in reward rate. However, the relationship between computational and biological implementations of temporal credit assignment remains unclear. The second aspect of credit assignment is the structural (i.e. spatial) aspect: how does the brain select the relatively few neurons to bias from its highly interconnected network of billions? Due to the brain’s highly interconnected nature, selective biasing of particular neurons would both be critical to changing the output of the network in a precise way and be more energetically efficient than globally increasing gains (Sutton and Barto 2018).




Figure 1.3. Figure from (Averbeck and Murray 2020) depicting a theorized dorsal/ventral organization of various neural circuits involved in goal-directed learning in the primate brain. Dorsal and ventral streams stay topographically aligned within cortex, striatum, pallidum, thalamus, and the resulting cortico-thalamic loop, but have shared inputs and outputs in the amygdala, temporal lobes, and hypothalamus. dlPFC: dorsolateral prefrontal cortex; DS: dorsal striatum, GP: globus pallidus, MDI: medial-dorsal thalamus; vmPFC: ventro-medial prefrontal cortex; cOFC: caudal orbito-frontal cortex; VS: ventral striatum, VP: ventral pallidum;

MDm: medial dorsal thalamic nucleus; BLA: basolateral amygdala, EC: entorhinal cortex; BNST: bed nucleus of the stria terminalis, PVT: paraventricular thalamus; Hypo: hypothalamus.

While computational treatments may provide some clues, a holistic model of credit assignment to cortical neurons as implemented in the mammalian brain remains elusive. As one key ingredient to credit assignment (and arguably the most definitive conclusion on credit assignment to date), dopamine in the brain has been shown to represent the reward prediction error (RPE) of the animal, i.e. the difference between expected and actual reward (Schultz et al. 1997; Dabney et al. 2020). This RPE signal could be sufficient to drive changes to the animal's value functions (Sutton and Barto 2018). However, much more surrounding credit assignment remains a mystery. The striatum has been a primary focus of research, as it is both anatomically well-positioned to mediate between behavior and reward and has been theorized to house action-value and state-value functions in the brain along its dorsal and ventral aspects, respectively (Shin et al. 2021; Lee et al. 2012), though the nature of this representation in the striatum has been questioned (Peters et al. 2021; Elber-Dorozko and Loewenstein 2018) (Figure 1.3). Indeed, the striatum seems to have a complicated relationship with behavior, differing across probabilistic vs deterministic tasks (V. D. Costa et al. 2016), learning timescales (Yin et al. 2009), and flexibility in the task (Nicola 2010), and results sometimes differ depending on the analysis methodology or spatial scale of the neural recordings (Sales-Carbonell et al. 2018; Peters et al. 2021). Furthermore, studies have historically compared striatal activity either to behavior or to broad swaths of cortex, leaving it unclear how the striatum might more specifically represent task-relevant cortical activity to underlie temporal credit assignment. Finally, while a mechanism that enables dopamine to selectively bias recently active synapses has been described (Yagishita et al. 2014), it remains unclear how striatal activity observed during tasks relates to selective modulation of particular cortical neurons as a part of structural credit assignment.

In Chapter 3, we utilize a BMI experiment in rats to comment on the striatum's role in credit assignment. As described above, a BMI experiment is well-suited to study the relationship between cortical and striatal activity: activity from experimenter-defined cortical neurons can drive the task and thus the process of credit assignment can be observed through these particular cortical neurons. In particular, through analyzing simultaneous recordings of the striatum during a BMI task, we shed light on how cortical activity relates to striatal activity through an internal, continuous model of proximity to reward, demonstrating differential roles across dorsal and ventral axes of the striatum.

Scalable Software for Neuroscience

Modern neuroscience experiments, especially closed-loop experiments such as those described here, are increasingly complex and heterogeneous. Experiments often measure behavior across multiple modalities simultaneously, and increasing interest in ethological behaviors make behaviors less defined *a priori* compared to rigid task structures (Markowitz et al. 2018; Miller et al. 2022). Concomitantly, neural data volume has exploded in volume in recent decades (Steinmetz

et al. 2021; Ota et al. 2021), enabling experiments to even combine neural recording modalities to maximize the number of neurons under study (Peters et al. 2021). Adding a real-time requirement to the above, utilizing closed-loop methodologies are necessary for implementing brain-machine interfaces (Shanechi 2019) and have also been suggested to increase the efficiency of neuroscience experiments (Z. S. Chen and Pesaran 2021).

An oft-neglected aspect of modern neuroscience is the role of – and time spent on – the software that underlies modern experiments and analytical methods. The experimental demands described above are challenging to implement in a robust, correct, and scalable manner, especially for neuroscientists who, by nature, are not often experienced software engineers. The penalty of mistakes is grave: often months to years in iteration. What results from these challenging experimental demands, then, is a slow iteration cycle for neuroscience research, in part due to unfamiliarity with existing software tools and a lack of knowledge of optimal software practices leading to implementations of limited potential for reuse.

Fortunately, many of these demands of modern neuroscience have been solved in the software industry, at least to varying degrees, and thus tools can leverage these existing software tools or practices to improve the development of neuroscience experiments. Such tools, if open-sourced and properly documented, could then increase the speed and correctness at which neuroscience proceeds. For example, open-source tools for data acquisition (Siegle et al. 2017) and for performing spike sorting on state-of-the-art electrophysiological probes (Steinmetz et al. 2021) have already enabled rapid, reproducible advances, and the neuroscience community heavily leverages existing open-source libraries for data analysis and machine learning. Given this, Chapter 4 of this thesis focuses on contributions made to the open-source neuroscience ecosystem aimed at two problems increasingly common in experiments. In particular, we focus first on a solution for utilizing and analyzing heterogeneous streams of data arising from possibly disparate devices, simplifying setup and powering closed-loop experiments. Secondly, we highlight a solution for general neuroscience experimental design, leveraging a type of process architecture (“microservices”) commonly used in the software industry to achieve scalability and developer efficiency. While much more can be done, we hope these contributions can accelerate the neuroscience community.

Chapter 2

Skilled independent control of individual motor units via a non-invasive neuromuscular-machine interface

Brain-machine interfaces (BMIs) have the potential to augment human functions and restore independence in people with disabilities, yet a compromise between non-invasiveness and performance limits their relevance. Here, we hypothesized that a non-invasive neuromuscular-machine interface (NMI) providing real-time neurofeedback of individual motor units within a muscle could enable independent motor unit control to an extent suitable for high-performance BMI applications. Over 6 days of training, 8 participants progressively learned to skillfully and independently control three biceps brachii motor units to complete a two-dimensional center-out task. We show that neurofeedback enabled motor unit activity that largely violated recruitment constraints observed during ramp-and-hold isometric contractions thought to limit individual motor unit controllability. Finally, participants demonstrated the suitability of individual motor units for powering general applications through a spelling task. These results illustrate the flexibility of the sensorimotor system and highlight individual motor units as a promising source of control for BMI applications.

2.1 Introduction

Brain-machine interfaces (BMIs) aim to create an artificial link between intentions and actions. By detecting user intent from neural activity, BMIs can enable symbiotic human-machine interactions that are independent of the motor system and thus have great potential to augment human functions. Proof-of-concept clinical studies have tapped into this potential to restore independence in people with severe paralysis, demonstrating systems that allowed tetraplegic people to control robotic arms and exoskeletons (Collinger, Wodlinger, et al. 2013; Benabid et al. 2019), navigate computers (Hochberg et al. 2006), and even regain control of their own paralyzed limbs through electrical stimulation (Bouton et al. 2016). However, despite decades of advances, the reach of brain-machine interfaces remains relatively limited, largely caused by the current trade-off between BMI invasiveness and performance (Millán and Carmena 2010; Blabe et al. 2015). Intracortical BMIs demonstrate outstanding performances but present significant associated risks (Hochberg et al. 2006; Collinger, Wodlinger, et al. 2013; Bouton et al. 2016; Pandarinath et al. 2017; Benabid et al. 2019); non-invasive BMIs, such as those based on electroencephalography (EEG), have a low barrier-to-entry, but their poor spatial resolution and vulnerability to noise

artifacts have so far limited them to specialized use-cases and to information transfer rates too slow to control complex devices (Millán and Carmena 2010).

Alternatively, user intent can be accessed at the level of the muscles via surface electromyography (EMG), a non-invasive technology whose high temporal resolution recordings of single motor units support broadly applicable yet powerful applications. Neuromuscular-machine interfaces (NMIs) built on surface EMG seek to detect intended actions from muscle activity, as in the case of myoelectric-controlled prostheses where intended hand movements are decoded from the activity of a range of upper-limb muscles (Hahne et al. 2018; Zhuang et al. 2019). Existing NMIs can effectively link predefined sets of muscle functions to device control signals, but typically require interfacing at least as many muscles as the number of actions to be controlled. Therefore, the bandwidth of these devices is often insufficient to enable effective control of assistive devices in people with paralysis, in which only a limited number of muscles can be used as sources of control. In addition, in decoding intended movements, these NMIs do not differentiate between movements aimed at device control or for interacting with the environment, making their utility for augmentative applications limited. To effectively enable these applications typically restricted to BMIs, NMIs should thus seek to increase the number of degrees of freedom that can be extracted per muscle.

Henneman's size principle theorizes that individual motor units within a muscle are recruited in a fixed order (Elwood Henneman 1957, 19) and thus cannot be controlled independently from one another. This long-standing theory of orderly recruitment has been primarily supported by experiments assessing motor unit recruitment properties during isometric, slow-ramping contractions within controlled laboratory conditions (Elwood Henneman 1957; E Henneman et al. 1974; Milner-Brown et al. 1973; ter Haar Romeny et al. 1982b; De Luca and Mambrito 1987), and is consistent with the prevailing view that motor units within a muscle are controlled by a common descending neural drive (De Luca and Erim 1994). According to this understanding, the bandwidth of existing NMIs would be neurophysiologically constrained. However, out-of-order motor unit recruitments have been observed during complex motor behaviors (Grimby and Hannerz 1970; Nardone et al. 1989; Marshall et al. 2021), and multidimensional descending neural drives have been shown to contribute to this flexibility (Marshall et al. 2021). In addition, pioneering studies in neurofeedback reported that people can learn to volitionally control individual motor units belonging to the same muscle when provided with visual and/or auditory feedback linked to the units' activity (Virginia F. Harrison and Mortensen 1962; J. V. Basmajian 1963; V. F. Harrison and Koch 1972; Illyés 1977; John V. Basmajian and De Luca 1985). In particular, Harrison and Mortensen reported a subject that was able to learn, within an hour of training, to isolate and produce predetermined patterns of activity in 4 motor units of the tibialis anterior muscle (Virginia F. Harrison and Mortensen 1962). The Basmajian group expanded on these findings with a series of exploratory studies, reporting selective control of up to 6 motor units across a variety of participant demographics and muscles, including the tibialis anterior, biceps brachii, and the abductor pollicis brevis (John V. Basmajian and De Luca 1985). These results suggest that individual motor units within a single muscle could

potentially provide the per-muscle bandwidth required for powering BMI applications. However, because of the qualitative and observational nature of these neurofeedback studies, the extent to which individual motor units can be volitionally controlled remains largely unclear. Indeed, previous studies did not quantify control quality or independence, test whether motor units can be recruited both exclusively and simultaneously of one another, assess the ability to skillfully modulate firing rates, examine learning capabilities over time, or compare motor unit activities between periods of neurofeedback control with relevant isometric contractions of the same muscle. Answering these questions, in addition to significantly enhancing our understanding of flexibility in the neuromuscular system, could assess the suitability of using individual motor units as the source of control for a new class of NMIs for both translational and augmentative applications.

In order to address these questions, here we utilized a neurofeedback paradigm coupled with an operant learning task to interrogate the emergence and execution of skilled, independent control of individual motor units. We devised an NMI that provides visual and auditory feedback of biceps brachii motor units in real-time using neuromuscular signals recorded from a high-density grid of surface EMG electrodes and trained 8 participants over 6 consecutive days to use this system on a center-out task requiring skilled and independent control of three motor units. Participants increased in proficiency in this skilled motor unit control both within and across days of training, with modest control of at least 2 motor units demonstrated even on the first day. Through comparisons to isometric, ramp-and-hold contractions, we provide evidence that neurofeedback enabled participants to control individual motor units outside of recruitment constraints thought to limit motor unit controllability. We then demonstrated an application of such fluent motor unit control through a speller task, in which participants used motor unit activity to navigate a virtual keyboard to spell sentences. These results characterize skillful control of individual motor units, contributing to our understanding of sensorimotor flexibility and suggesting motor unit control can be a viable paradigm for both clinical translation and human augmentation applications.

2.2 Results

We devised an NMI capable of providing real-time visual and auditory neurofeedback of biceps brachii motor unit action potentials (**Figure 2.1A**). This NMI measured neuromuscular signals using a high-density grid of surface EMG (HD sEMG) electrodes and used previously validated blind source separation and classification techniques to decompose these signals into individual motor unit action potentials in real-time (Negro et al. 2016; Barsakcioglu et al. 2020). After a brief initialization period for the decomposition model, we first instructed participants to use the NMI's neurofeedback to explore covert strategies to control individual motor units independently from one another. The goal of participants during this exploration procedure was to find and sort in order of controllability the three motor units they felt had the highest potential for independent control (**Figure 2.1B**). A motor unit selection algorithm highlighted motor units with potential for independent control and guided participants in this task. After this exploration period, participants'

ability to control their selected motor units was tested in a center-out task (**Figure 2.1C, D; Supplementary Video 1**). A population-coding strategy was used to map motor unit activity into the 2D position of a computer cursor, and participants had to operate this cursor to achieve the displayed targets. 12 peripheral targets were used to evaluate whether participants could recruit the selected motor units exclusively of one another (T1, T2, and T3 targets) and simultaneously in combinations of two (T4 targets) and could regulate the firing rate of the recruited units (close and far targets, **Figure 2.1C**). T1, T2, and T3 targets were ordered such that T1 corresponded to exclusive recruitment of the subjectively easiest motor unit to activate independently and T3 the subjectively hardest. A center target requiring participants to coactivate all the selected units at a similar intensity (T5 target) was also used. These targets were grouped into 3 difficulty levels, with targets of increasing difficulties becoming available after reaching an average success rate greater than 3 targets per minute on the tested targets (**Figure 2.1D**). We used this paradigm to train 8 participants over 6 consecutive days. Participants' arms were constrained to fixed elbow and wrist angles via a sensorized orthosis for the entirety of each session. Additionally, while we did not explicitly track motor units across days, we used markings on skin to ensure consistent electrode positioning.

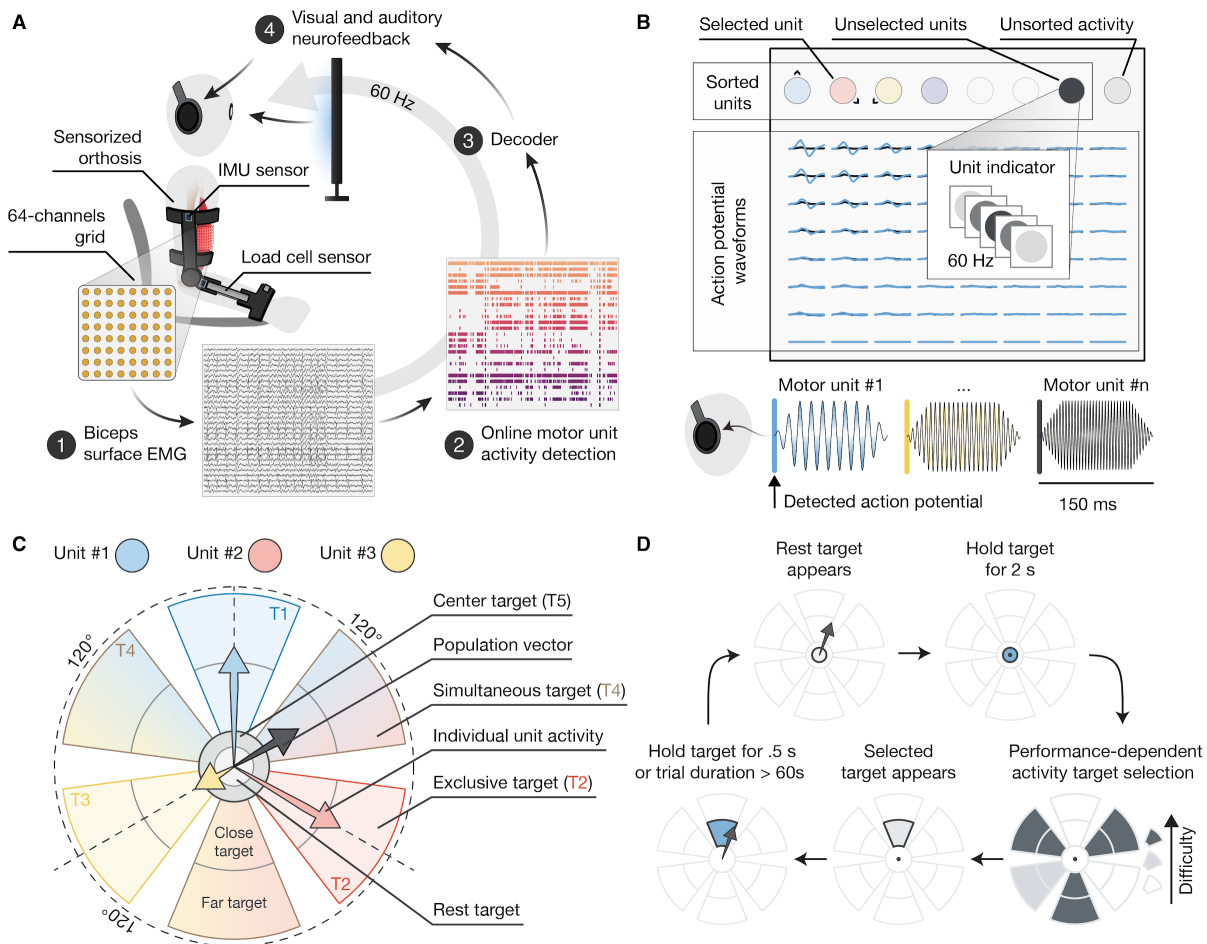


Figure 2.1 | Experimental setup. **A**, Schematic of the neuromuscular-machine interface (NMI) used to enable individual motor unit control of the biceps brachii. Participants are seated on a chair wearing a sensorized orthosis constraining the elbow joint at

100 degrees and the wrist at its neutral position. Load sensors are used to measure the isometric elbow-flexion and forearm-supination forces. IMU sensors are used to track arm movements. The NMI control loop is divided into 4 steps. First, biceps brachii neuromuscular signals are measured using a high-density grid of 64 surface EMG electrodes. Second, an online decomposition model is used to detect motor unit action potentials from the measured signals. Third, a decoder transforms the detected motor unit activity into task-dependent neurofeedback signals. Last, auditory and visual neurofeedback signals are delivered to the participants via headphones and a computer monitor. **B**, Schematic of the user interface and neurofeedback signals used during the exploration procedure. Multi-channel waveforms of the detected motor unit activity are displayed and updated at 60Hz. Neurofeedback of the detected motor unit activity is also provided by LED-like indicators flashing when an action potential is detected. Both waveforms and unit indicators are color coded. Colored signals indicate the activity of a subset of selected individual motor units. Black signals indicate the activity of unselected motor units. Finally, light-grey signals indicate detected events that have not been categorized as motor unit activity, i.e. unsorted activity. Auditory neurofeedback signals followed the same categorization between selected, unselected, and unsorted units and consisted of 150 ms pitch-coded stimuli. **C**, Center-out task neurofeedback, decoder, and targets. The activity of three selected motor units is transformed into cursor position using a population coding schema. The cursor position is indicated by a grey arrow originating at the center of the screen and represents the population vector. The same unit-specific visual indicators and auditory stimuli employed in the exploration period are used here. A total of 12 peripheral targets (T1, T2, T3, and T4), 1 center target (T5), and 1 rest target were included. **D**, Center-out task protocol. The task is divided into trials. To start a trial participants need to hold the cursor within the rest target for a minimum of 2 seconds. A target is then selected from a performance-dependent pool of available targets. At first, only T1, T2, and T5 targets were available. T3 and T4 targets were progressively added depending on participants' performance within that day. The trial's target is displayed and the participant has 60 seconds to achieve it before the trial is declared unsuccessful.

Skilled independent control of individual motor units on day 1

We found that participants displayed independent control over selected motor units already at day one (**Figure 2.2**). In particular, participants successfully completed an average of 95.6% and 79.2% of the presented T1 and T2 targets on day one, demonstrating independent control of motor unit #1 and #2, respectively (**Figure 2.2A-C**, $p < 0.001$ when testing for % successful trials > 0). All but one participant surpassed the 3 targets per minute threshold in success rate required to enable T3 targets, and half of the participants subsequently reached sufficient proficiency to also enable T4 targets (**Figure 2.2C**). Participants encountered no difficulty in performing T5 targets, succeeding in all the corresponding trials. We also found no statistically significant difference in the percentage of correct trials between targets with different distances ($p > 0.05$ for each target category, **Figure 2.2D**). These results demonstrate that participants, without any prior training, can gain skilled independent control of 2 or 3 motor units within a single session, suggesting some level of latent flexibility in the sensorimotor system.




Figure 2.2 | Independent control of individual motor units during the first day of training. **A**, Representative traces of center-out task signals for one participant during the first training day. First row, smoothed, normalized firing rate of the selected motor units used to control the cursor position. Second row, bipolar surface EMG signals from the three channels that best discriminate the activity of the selected motor units and relative raster plot of the detected motor unit firings. Third row, cursor position (r and θ , solid black traces) and targets (colored boxes) displayed in polar coordinates; grey dotted lines indicate interpolated values when θ is undefined ($r=0$). Bottom, arm position and angular velocity about the two axes of largest variation (PC #1 and #2). Grey-shaded areas crossing the different plots indicate ongoing trials and the relative target; empty spaces between these areas indicate rest targets. **B**, Median (lines) and 95% confidence interval (shaded areas) of the selected motor unit waveforms measured from the EMG channels in **A**. **C**, Summary statistics of the first training day. Left, box-plots representing the percentage of correct trials for each of the performed targets and participants. * indicate a significant difference from 0, $p < 0.0001$. Middle, box-plots representing the number of trials performed for each of the performed targets and participants. Right, medians (black lines) and 95% confidence intervals (shaded areas) of the number of participants that successfully performed at least one trial for each target category. **D**, Effect of target distance on percentage of correct trials. Colored point plots indicate the medians and 95% confidence intervals of the percentage of correct trials for close and far targets, for each color-coded target category. Light-grey scatter plot and box-plots report the raw data points and their distribution, respectively. No significant difference was found between targets of the same category but different distance ($p > 0.5$, n for each target category is indicated in **C**).

Learning over time

We next evaluated how participants' performance evolved over time. For this, a trial performance metric was first computed, which embedded information regarding the average distance of the cursor from the target, trial duration, and participants' ability to selectively recruit target-specific motor units (**Figure 2.3A and B**). A linear mixed-effect model was used to predict trial performance as a function of time, while controlling for possible variations between participants, days, and targets. Participants' performance increased both within ($p < 0.001$) and across days ($p < 0.006$), with fixed effects equivalent to an increase in performance of 1.4 standard deviations over 100 trials and of 0.4 standard deviations over the 6 days, respectively (**Figure 2.3C and D**). The fixed-effect for the interaction between the within- and the across-day time variables was non-significant ($p = 0.094$). The model intercept corresponded to an average successful trial rate of roughly 95% (standardized performance of -0.44, **Figure 2.3A**), confirming the previous analyses indicating successful task performances already at day 1.

Target-specific models were then built to better evaluate the effect of training on participants' ability to control the three selected motor units exclusive of one another (T1, T2, T3 targets). Results showed significant across-day learning for all 3 targets, but only significant within-day learning for the first two motor units, highlighting the importance of multi-day training to enable the emergence of skilled control of multiple individual motor units (**Figure 2.3E**). The interaction between learning within and across days was significant for T1 targets ($p = 0.028$) but not for T2 and T3 targets ($p = 0.2$ and $p = 0.67$, respectively). In addition, model intercepts showed that participants' performance on T1 and T3 targets were respectively higher and lower than average (T1 intercept > 0 , T3 intercept < 0 , $p < 0.05$), indicating that participants accurately ranked their motor units in order of controllability following the exploration procedure.

We finally analyzed how participants' performances on the simultaneous targets (T4) evolved over time. Since every participant did not reach these targets every day, only across-day learning was analyzed. Specifically, a generalized linear mixed-effect model was used to evaluate how the rate of successful trials evolved across days (**Figure 2.3F**). The fixed effect was significant, indicating an overall increase in the success rate across all participants ($p = 0.016$).

These analyses demonstrate that by the end of the 6 days of training all participants gained skilled independent control of the selected motor units (**Figure 2.3C-E**). The increase in performance across days also shows that learning is robust to changes in recording setups, suggesting strong potential for an NMI that would exploit this strategy to extract volitional control signals.

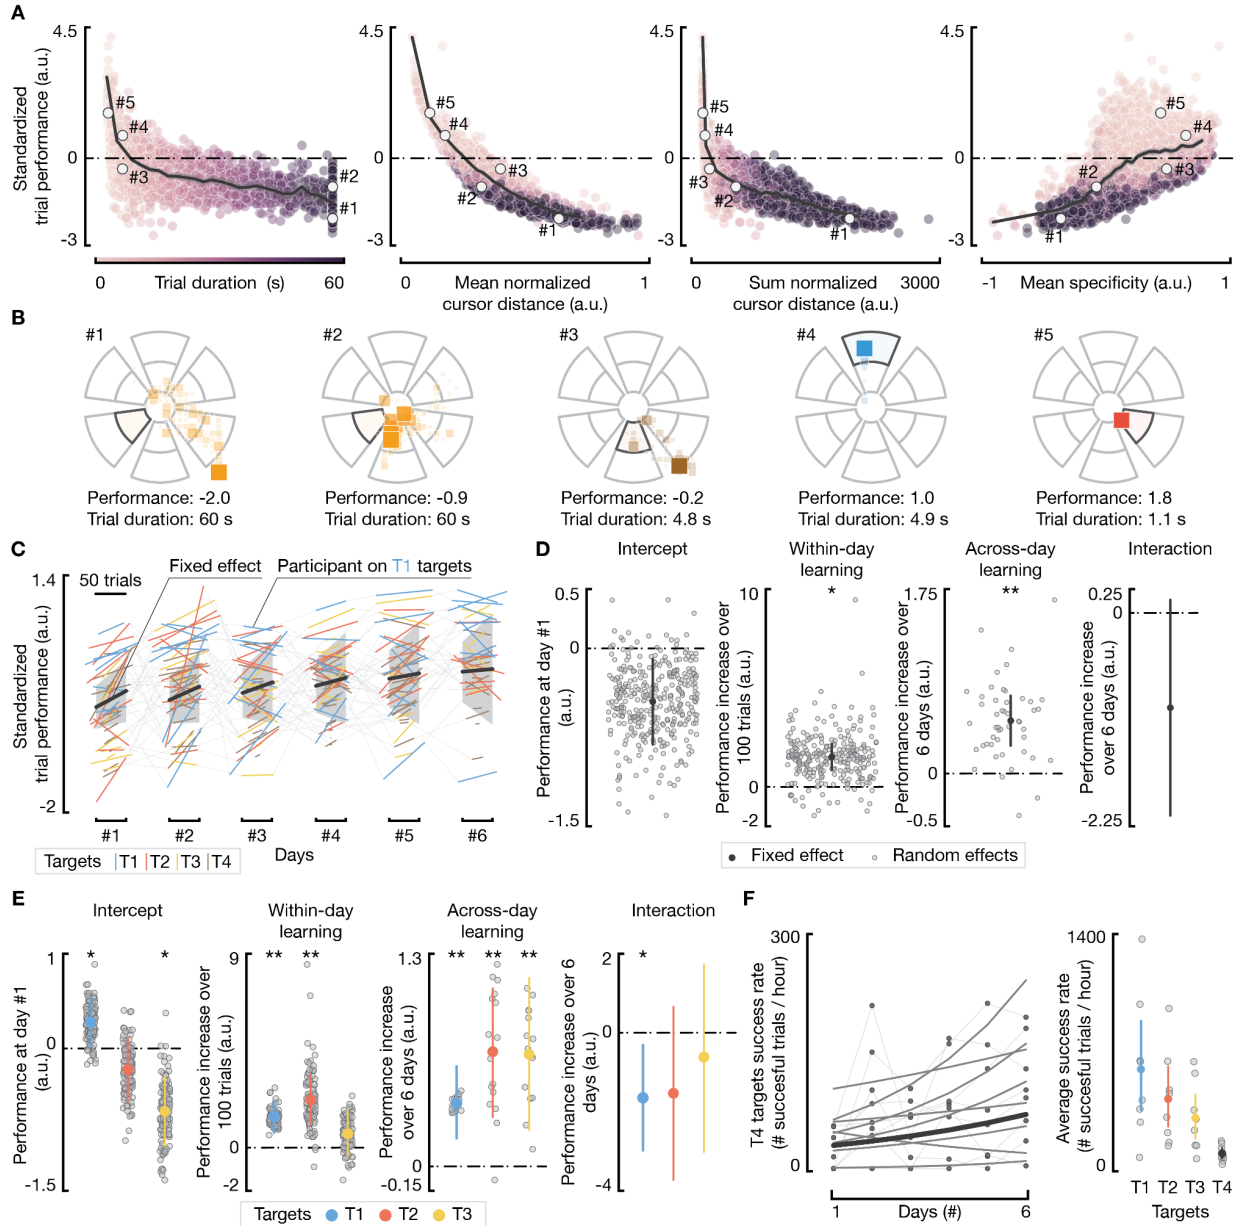


Figure 2.3 | Learning to control individual motor units independently over time. **A**, pair-plots showing the relationship between the holistic performance metric used to evaluate participants' proficiency in the center-out task and 4 metrics measuring specific behavioral characteristics: trial duration, mean and sum of the normalized cursor distance from target center, and mean specificity. White dots indicate 5 examples further displayed in **B**. Grey lines show the average performance for each metric; shaded areas indicate the 95% confidence interval. **B**, Scatter plots representing the temporal distribution of cursor position during the 5 trials depicted in **A**. Color alpha and square dimensions are proportional to the time spent in a given position. Trial #1 and #2 are both examples of unsuccessful trials. While the trial duration is the same for both (60 seconds), the holistic metric indicate better performance for trial #2, properly capturing differences in cursor trajectories between these two trials. Similarly, trials #3 and #4 are similar in duration but different in performance. Trial #5 reports an example of a high performance trial. **C**, Regression lines of the linear mixed-effect model used to evaluate overall learning within- and across-day (n samples = 5249). Thick black lines represent the regression lines of the within- and across-day fixed-effects, i.e., the effects that are generalized across participants, sessions, and targets; shaded grey areas indicate the 95% confidence intervals. Thinner, colored lines represent the fitted regression lines for each participant and target category. **D**, Fixed and random effects for key model parameters. The intercept indicates the performance at day #1. The interaction is between the within- and the across-day time variables. **E**, Fixed and random effects for key parameters of the models used to evaluate unit-specific learning behaviors (n samples = 1311, 1230, 1050, for the T1, T2, and

T3 models, respectively). **F**, Success rate of T4 targets across days fitted using a Poisson generalized linear mixed-effect model (n samples = 48, fixed effect $p=0.013$, left) and average success rates across all days for T1, T2, T3, and T4 targets reported for qualitative comparisons (right). Left, the thick line indicates the fixed-effect regression line; thinner lines indicate the regression lines for each participant; dots indicate the raw values. Stars indicate a statistically significant difference from 0: * indicates a $p<0.05$, ** indicates a $p<0.01$. Right, small grey dots indicate the averages for each participant; colored dots and bars represent the across-participants medians and corresponding 95% confidence intervals, respectively.

Exploration and acquisition of independent motor unit control

We then evaluated the role of the exploration period, occurring immediately before the center-out task, in the acquisition of independent motor unit control. Due to the unstructured nature of the exploration period, motor unit firing rates were first decomposed into separate components via non-negative matrix factorization (NMF) to identify groups of units that were often mutually active. The number of components was fixed to 3, aligning with the instructions given to the participant to ultimately select 3 representative motor units. Then, the cumulative independent firing time (CIFT) was computed for each component as the fraction of time a component was independently active relative to the overall time in which it was active (**Figure 2.4A** and **B**). The three components were then ordered in descending order by the CIFT value 2 minutes into the exploration period, and CIFTs were compared between this initial point and their final values (**Figure 2.4B**).

The CIFT increased significantly over the course of the exploration period (**Figure 2.4B** and **C**). The overall mean CIFT across the three components increased from 0.40 after the second minute of exploration to 0.51 at the end ($p<0.0001$, **Figure 2.4C**). The first component (C1) was activated nearly completely independently at the beginning of the exploration period, emphasizing the level of ease in attaining independent control in one set of motor units. However, C1 then began to co-activate more throughout the exploration period as the participant explored strategies for activating other sets of units, as indicated by a decreasing CIFT ($p<10^{-5}$; **Figure 2.4C**). On the other hand, the other two components (C2 and C3) increased in independent activation over time ($p<0.05$; **Figure 2.4C**), illustrating a progressive learning process.

We next asked whether participants' motor unit control in the exploration period improved across days. There was a significant increase in the mean CIFT for exploration periods across a participant's 6 days of training ($p=0.017$; **Figure 2.4D**). Participants thus demonstrated across-day improvements in independent motor unit control in both the center-out task and the exploration period. Finally, the mean CIFT displayed during the exploration period was found to have a strong positive correlation with the center-out task performance of the same day (fixed-effect slope: 1.61; $p<10^{-6}$; **Figure 2.4E**). Taken together, these results characterize the within-day and across-day processes by which participants acquired independent control of motor units.

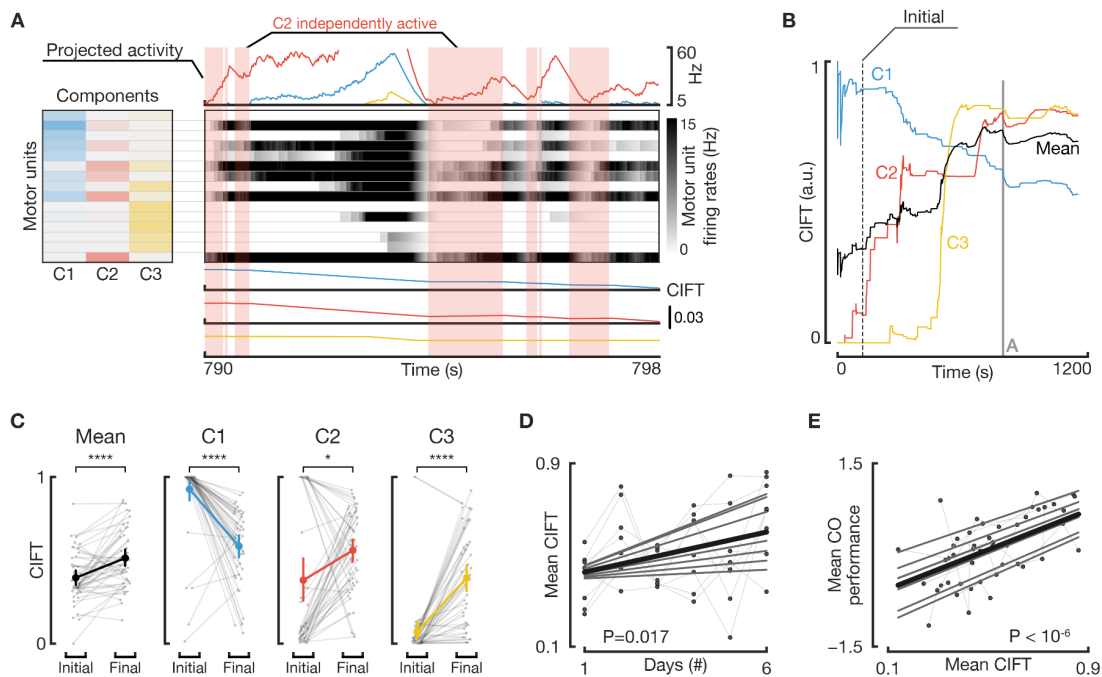


Figure 2.4 | Exploration and acquisition of independent motor unit control. **A**, Representative 8-second example for the extraction of components via non-negative matrix factorization (NMF) and the computation of CIFT. Three components are extracted from firing rates (center, grayscale heatmaps; white = 0, black = max) via NMF, yielding component-wise weights for each motor unit (left) and their corresponding projected activities (top). Then, CIFT is computed for each of the 3 components as the fraction of time spent independently active versus time spent active (displayed in the 3 bottom rows). In this example, C2 (red) has periods where it increases in CIFT (red shaded blocks) since it was independently active and where it decreases when C1 or C3 are also active. **B**, Data from the full, 20-minute exploration period from which data from **A** originated. For comparison of time courses for CIFT, we take values at an initial point (dotted line, left; 2 minutes into period) and at the period's final point. Black trace represents the mean CIFT across the three components. **C**, Changes in CIFT between initial and final points for the mean CIFT, C1, C2, and C3 (left to right). Faded black dots and lines are individual exploration periods, Stars indicate: * $p < 0.05$, **** $p < 0.0001$ from a paired t-test, $n = 48$. **D**, Mean CIFT at the end of the exploration period compared across 6 days of training and relevant regression lines from a linear mixed model fit on this data. Thin gray lines indicate participant-specific regression lines, while the thick black line represents the regression line for the fixed effect (linear mixed model, $p = 0.017$, $n = 48$). **E**, Relationships between the mean center-out task performance and the mean CIFT at the end of the preceding exploration period. Definitions of dots and lines are the same as in **D**. Fixed-effect slope of 1.61, $p < 10^{-6}$, $n = 48$.

Muscle activity dimensionality

Participants' success in the center-out task required independent motor unit control, indicating that the activity of the selected motor units lay along a multi-dimensional manifold. To evaluate how this differed from isometric motor behaviors typically used to study motor unit recruitment properties, each day participants were tested in a force-control task, using the same experimental setup as in the rest of the session. Here, participants were instructed to match displayed force profiles by performing isometric, ramp-and-hold contractions in the two primary movement directions of the biceps, elbow flexion and forearm supination (Naito 2004) (**Figure 2.5A**). Participants accurately reproduced the target forces (mean normalized $r > 0.95$, **Figure 2.5B**).

To analyze the dimensionality of motor unit activity between tasks, we computed the participation ratio of motor unit firing rates (**Figure 2.5C**). The participation ratio measures the spread of the variance captured by each of the dimensions of the eigen-spectrum (Recanatesi et al. 2019; Gao et al. 2017a) and can be envisioned as a continuous interpolation of the number of principal components needed to explain 80-85% of variance (83% in our study's data).

We found that motor unit firing rates had a higher average participation ratio during the center-out task than during the force-control task ($p < 0.0001$; **Figure 2.5D-E**). Participation ratio between tasks significantly increased whether considering solely the 3 selected motor units or the remaining unselected motor units, signifying an increase in dimensionality across the entire population of motor units ($p < 0.0001$; **Figure 2.5E**). In addition, selected units' firing rates could predict the concurrent firing rates of the unselected motor units fairly well (mean $R^2 > 0.56$ for both tasks) through a linear transform, indicating strong correlations between activities of the two groups ($p < 10^{-10}$ different than zero; **Figure 2.5F-G**). However, for the same population of units, the R^2 metric was lower in the center-out task, indicating a decoupling between selected and unselected units (**Figure 2.5G**; $p < 10^{-10}$). Relatedly, the participation ratio during the center-out task increased over the 6 days of training in the selected motor units but not in the unselected motor units ($p < 0.01$ in selected, $p = 0.68$ in unselected; **Figure 2.5H**).

Finally, we compared motor unit firing rates dimensionality with that computed on the integrated EMG (iEMG), a commonly used feature for measuring overall muscle activity. Participation ratios computed on the iEMG showed similar across-task differences, yet participation ratio increased more for firing rates than for iEMG ($p < 0.01$; **Figure 2.5I-J**).

Taken together, these results reveal the center-out task enabled both a significant, population-level increase in dimensionality relative to during stereotyped, isometric contractions and an increased decoupling between unselected and selected motor unit populations. Differences in dimensionality measured using firing rates and iEMG also suggest that individual motor unit resolution recordings are necessary to fully appreciate muscle activity dimensionality.




Figure 2.5 | EMG dimensionality during stereotyped isometric contractions and neurofeedback tasks. **A**, Overview of force-control task. Participants matched trapezoidal force profiles shown to them on-screen in varying amplitudes and in various combinations of elbow flexion and wrist supination. **B**, Example set of 8 control task trials (gray highlights) in one representative session. Participants performed various trial types (top row) and matched target forces fairly accurately (second row; blue flexion and orange supination). Two features were extracted: iEMG (third row) and motor unit activity (fourth row; each tick is a detected firing of its row's motor unit). **C**, Toy data illustrating 3 clusters with different distributions of variance in order to explain the computation of participation ratio. When variance is equally distributed across the principal components, as in the circular distribution (top-left, blue), participation ratio (PR) is maximized and equals the number of data dimensions. On the other hand, when variance is more unevenly weighted across dimensions, as in the ellipsoidal distributions (green or orange), PR decreases towards 1. **D**, PR for each session's force-control (x axis) and center-out (y axis) tasks. Different colors represent different participants; faded dots represent actual sessions while highlighted dots represent medians within participants. Circles represent PR of firing rates; triangles for iEMG. Dotted line represents line of equal PR between the two tasks, i.e. $y=x$. **E**, PR for firing rates of all units (left), only selected motor units (center), and unselected motor units (right) across the force-control (FC, red) and center-out (CO, green) tasks. Faded dots represent individual sessions. Differences across tasks for all 3 populations were significant ($p < 0.0001$, paired t-test, $n=48$). **F**, Two 10-second representative examples of simultaneous firing rates for the 3 selected motor units (top) and 2 unselected motor units (middle and bottom rows). Dotted lines indicate the predicted firing rates of the unselected motor units from the selected units' firing rates. **G**, Coefficients of determination (R^2) between optimal linear transformation of selected motor units' firing rates and unselected motor units' firing rates. Force-control's R^2 had a mean of 0.815, while the center-out's mean was 0.567 ($p < 10^{-10}$ different from zero for both, $n=48$). The center-out task had a lower mean than the force-control task ($p < 10^{-10}$, $n=48$). **H**, Change in PR in firing rates of selected motor units (top) and unselected motor units (bottom) during the center-out task over each participant's 6 days of training. Dots connected by light gray lines represent a particular participant's sessions, while light gray bold lines represent the subject-specific regression line. Bold black lines represent the fixed-effect regression line from the linear mixed model fit to this data. The fixed-effect slope was significant for selected units ($p=0.01$, $n=48$).

but not significant for unselected units ($p=0.68$, $n=48$). **I**, PR for iEMG between force-control (FC) and center-out (CO) tasks. PR for iEMG increased ($p=0.004$ different from zero, $n=48$). **J**, Changes in PR between force-control and center-out tasks for iEMG (left) and firing rates (right). The across-task increase in PR in iEMG was less than that of firing rates ($p=0.003$, paired t-test, $n=48$).

Motor unit recruitment order

Our results suggest that recruitment order of biceps brachii motor units might be more flexible than previously thought and that neurofeedback may enable motor unit recruitment that significantly differs from those observed during typically-studied, isometric, slow-ramping contractions. To evaluate this divergence in motor behaviors, we compared the stability of motor unit recruitment order across tasks.

We first assessed recruitment thresholds of selected and unselected motor units during the force-control task. Flexion and supination recruitment thresholds for all units spanned a wide range, distributed in agreement with the common model of motor unit frequency distribution skewing towards more lower-threshold units within a muscle (Fuglevand et al. 1993) (**Figure 2.6A-B**). 97% of motor units selected for the center-out task were also detectable during isometric muscle contractions; the remaining 3% were not recruited during flexion or supination contractions possibly due to small changes in postures that often occurred between tasks, and were excluded from the following analysis. 12% of selected motor units were recruited exclusively during either flexion or supination contractions, and 33% had categorically different recruitment thresholds between flexion and supination contractions. This varied recruitment order is in support of existing studies reporting biceps motor units can be recruited selectively for flexion or supination (ter Haar Romeny et al. 1982a; Borzelli et al. 2020; ter Haar Romeny et al. 1984) (**Figure 2.6B**).

We then asked if the joint firing rates of selected motor units observed during the center-out task overlapped with those observed during the force-control task. In order to focus on the same firing rate regime across the two tasks, the firing rates of each of the 3 selected motor units were first normalized to 1.5 times the normalization value used for that unit in the center-out task, discarding firing rates above these values. Then, after binning these normalized firing rates, a control space was defined for each session's force-control and center-out tasks as the set of joint binned firing rates observed for at least 160ms during each task (**Figure 2.6C**). Since a particular point in control space could be populated with only a nominal amount of time, the control space served as a liberal estimate of putatively achievable joint firing rates and, consequently, described possible motor unit recruitment orders achievable in the two tasks. The control space of the force-control task spanned a small fraction of the possible controllable space (8.9%, **Figure 2.6D-E**), suggesting the force-control task activity of the 3 selected motor units were largely stereotyped and low-dimensional, in agreement with the above dimensionality analysis. In contrast to this, the control space of the center-out task occupied significantly more of the possible controllable space (38.1%, $p<10^{-5}$; **Figure 2.6D-E**). Furthermore, if elbow flexion and forearm supination were effective strategies for individual motor unit control, one might expect firing rates of high-performing center-out sessions to overlap in control spaces between tasks. In opposition to this, center-out task performance increased proportionally to the relative increase in control space achieved during the center-out task, further emphasizing that individual motor unit control did not

solely rely on elbow flexion and forearm supination primitives, even for the higher-performing sessions ($p=0.012$, **Figure 2.6F**). This relationship holds even when only considering the easiest two targets (T1/T2) and the control space across the 2 respective units ($p=0.014$, **Figure 2.6F**).

These results demonstrate that, despite known task-dependent recruitment order variation in the biceps brachii, motor unit recruitment during the neurofeedback task significantly differed from that measured during elbow flexion and wrist supination isometric contractions, suggesting independent motor unit control involved an expansion of activity outside of these two existing motor primitives.




Figure 2.6 | Motor unit recruitment orders during neurofeedback and isometric contraction tasks. **A**, motor unit recruitment thresholds for both elbow flexion (x-axis) and forearm supination (y-axis) for all recorded motor units across all sessions and participants. Dots displayed in grayed areas below $y=0$ represent units that were only activated during flexion; dots left of $x=0$ represent units only activated during supination. Blue dots represent motor units that were selected for the center-out task; gray dots otherwise. **B**, table showing the distribution of selected motor units in particular recruitment threshold categories: NR: “not recruited”; LO: motor units with thresholds less than 10% of the maximum force observed that day; HI: remaining motor units with valid thresholds. Motor units selected for the center-out task had a lower average recruitment threshold for both flexion (0.23 for selected units vs 0.33 for unselected units, $p<10^{-5}$) and supination (0.18 for selected units vs 0.26 for unselected units, $p<0.001$) than unselected motor units. **C**, Representative data illustrating the computation of control spaces across selected units, in the case of only 2 units. Normalized firing rates for two selected motor units are shown in the top row, both during a supination trial in the force-control task (left) and during two center-out trials (right). The bottom row shows the force-control (left) and center-out (right) control spaces of these 2 units in purple shading, where the particular control space points visible in the top row’s examples are outlined in black. **D**, Visualization of the control spaces populated during the center-out (CO) task, the force-control (FC) task, and their difference. For visualization purposes here, non-negative matrix factorization (NMF) was first used to reduce the selected units’ firing rate dimensionality from 3 to 2 before computation of the center-out (top-left) and force-control (top-center) control spaces. Component #1 was defined as the component out of the 2 that explains more variance. For each session, points in the session’s center-out control space not seen in that session’s force-control control space were termed “novel”, and then the percent of sessions in which a particular point in control space was novel was computed (top-right). Gray squares indicate control space

points in which fewer than 5% of the sessions had that point populated. Bottom: variance of original firing rates explained by the 2 NMF components during the force-control task (mean 97.5%) was significantly (paired t-test; $p < 10^{-5}$, $n=48$) higher than that during the center-out task (mean 77.6%). **E**, Proportion of populated control space by firing rates during the force-control (FC, left) and center-out (CO, right) tasks. FC control space was significantly lower than CO control space (8.9% FC vs 36.1% CO, paired t-test, $p < 10^{-5}$, $n=48$). **F**, Relationship between center-out performance and the relative fraction of all control space bins populated across both tasks that were novel during the center-out (CO) task. Regression lines had a significant positive slope both when comparing all 3 selected units' targets to a 3-D control space (blue, $p=0.012$ slope > 0 , $n=48$) and when comparing only the first 2 units' targets to their respective 2-D control space (purple, $p=0.014$ slope > 0 , $n=48$). Individual dots represent a session. Stars indicate statistically significant slopes different than zero ($p < 0.05$).

Confound analyses

While participants' elbow and wrist joints were constrained by the orthosis, gross movements at the level of the shoulder or the spine could have affected motor unit detection quality. To control for this potential confound, in addition to instructing participants to minimize movements when trying to control motor unit activity, we recorded arm kinematics and analyzed arm movements throughout the center-out task. The rotational axis of the kinematic sensors was first aligned to the axis of largest variation, with 87% of movement occurring around a single rotational axis (**Figure 2.7A**). Data along this axis were then used to quantify movement amplitudes throughout the center-out task. In particular, the mean absolute velocity (MAV) during trial and inter-trial periods was used to measure the overall movement observed across the different task conditions. For each target, we then computed the within-day median and used this statistic to evaluate possible movement strategies. Results highlight minimal movements across all conditions, with a grand median value of approximately 0.48 deg/s (**Figure 2.7B**). When comparing the statistics of active targets to the rest targets (i.e. the inter-trial periods), we found a statistically significant increase in median MAV only during T5 targets ($p < 0.001$, **Figure 2.7B**), highlighting how the nonspecific motor unit recruitment required by these targets pushed participants to perform vigorous muscle contractions to obtain the target as quickly as possible. In contrast, we found a significant movement reduction during both trials for T1/T2 targets and the rest trials ($p < 0.001$ and $p < 0.01$, respectively) and no significant differences between trials for the T3 target and rest trials, confirming that participants minimized body movements throughout this task. We then evaluated whether participants relied on specific movements to independently control the selected units by trying to decode the cursor position during the center-out task from arm kinematics using a multilayer perceptron. Decoding performances were poor (median $R^2=0.165$, **Figure 2.7C**), indicating weak correlations between arm kinematics and cursor position. These combined results provide compelling evidence that the independent control of single motor units observed throughout the center-out task was neither the result of movement artefacts nor relied on specific movements.

Another confound that could have facilitated independent motor unit control is the presence of crosstalk from neighboring muscles in the recorded neuromuscular signals. Aside from the biceps brachii, the brachialis is the next most likely muscle to be recorded by our electrodes due to its proximity; however, while the biceps brachii is known to participate in both flexion and supination, the brachialis participates only in elbow flexion (van Zuylen et al. 1988; Staudenmann and Taube 2015). In order to assess recordings for brachialis contamination, we computed the

correlation of each channel's iEMG to flexion and supination forces during periods in the isometric contraction task where these task-oriented contractions were tested separately (**Figure 2.7B**). Correlations for flexion and supination were averaged within the three groups of channels most vulnerable to brachialis contamination: the column located most externally (i.e. nearer to the long head of the biceps), the column located most internally, and the distal row of channels. Mean correlations for all channel groups remained relatively high (> 0.7) across both flexion and supination. While spatial differences in correlations are expected even within the biceps brachii, channels primarily recording from the brachialis should display a marked drop in supination correlation during supinating contractions (Staudenmann and Taube 2015). The high correlations for both flexion and supination suggest brachialis contamination in our recordings was minimal and that the recording grid was primarily placed over the biceps brachii. Taken together, these results suggest that movement artifacts and crosstalk contaminations are unlikely to have significantly affected the validity of our results.

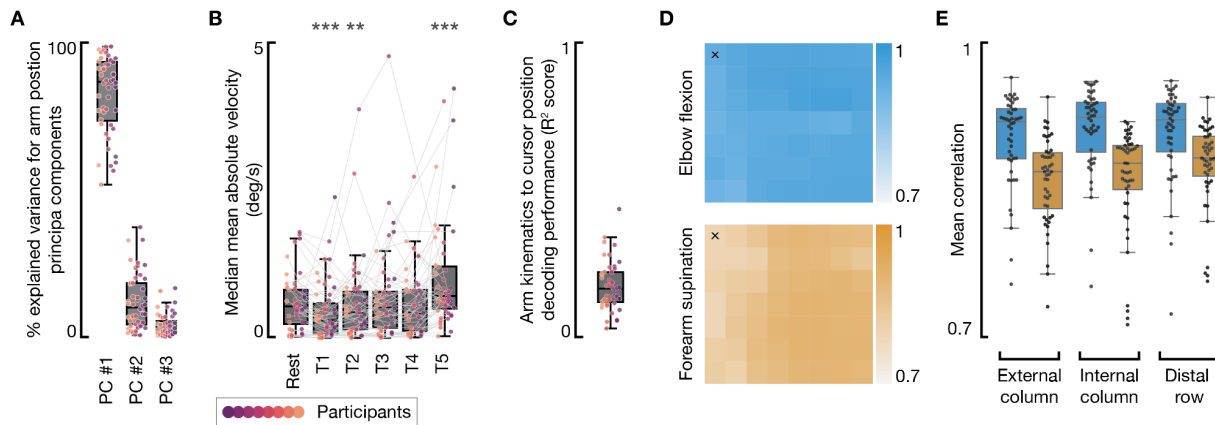


Figure 2.7 | Confound analyses. **A**, Fraction of variance explained by the first three principal components of the recorded kinematic data during the center-out task. **B** Arm movement velocities during the center-out task. Scatter-plot represents the median of the trials' mean absolute value (MAV) velocity for each participant, session, and target. Box-plots represent these statistics' distribution for each target. *** and ** indicate a significant difference between a given target category and the rest target ($p < 0.001$ and $p < 0.01$, respectively, bootstrapping with $n = 10000$ iterations, n samples = 48, 48, 47, 39, 48 for T1, T2, T3, T4, and T5 targets, respectively). Lines indicate data for a single participant in a given day. **C**, Decoding performances of a multilayer perceptron (MLP) trained to predict cursor position from arm velocity and position. Each dot represents the decoding performance (R^2) on the testing set (30% of the data) for one day and participant (color coded). **D**, Representative correlations between iEMG for each of the 56 channels to elbow flexion (blue) or forearm supination (orange) forces during one session. Channels are arranged according to physical position: the cells marked with "x" represent the most external (i.e. closest to the long head of the biceps brachii) and proximal channels recorded on the bicep. **E**, The mean correlations to flexion (blue) or supination (orange) forces within 3 different channel groupings most vulnerable to contamination from the brachialis. Each dot is a session's correlation.

Speller task

We then demonstrated a proof-of-concept of the translational potential of an NMI enabling skilled independent motor unit control. To demonstrate a clinically relevant application, participants were tested on a commonly used copy-typing speller task requiring point-and-click navigation of a virtual keyboard (Pandarinath et al. 2017; P. Nuyujukian et al. 2015; Jarosiewicz et al. 2015)

(**Supplementary Video 2**). This speller task utilized the same selected motor units from the center-out task but, as opposed to the center-out's position decoding, instead translated the normalized motor unit firing rates into the velocity of an on-screen cursor (**Figure 2.8A**). Navigating this cursor on a virtual OPTI-II keyboard displayed on the computer monitor, participants copied sentences by controlling motor units independently for both cursor movement and cursor clicking (Pandarinath et al. 2017; Rick 2010) (**Figure 2.8A**). The keyboard featured wraparound borders, which in combination with the cursor's velocity control allowed for full 2D navigation even with a single motor unit. We reasoned this more permissive control strategy to be better suited for translational applications compared with the control strategy used in the center-out task. Cursor clicking was triggered by simultaneously recruiting all the selected motor units, similar to achieving the center-out T5 target. Participants performed the speller task after at least 30 minutes of center-out task execution on the last 3 days of training, plus on any prior days in which they felt confident with their performance and completed a minimum of 60 minutes of recording.

Information throughput was assessed with the achieved bitrate, a conservative estimate of the true throughput of an assistive device (P. Nuyujukian et al. 2015). Average and peak bitrates on the speller task were promising: the mean average bitrate on the last day was 0.43 bits/s — corresponding to 5.41 correct characters per minute at a 92.2% accuracy — with a mean peak bitrate of 0.55 bits/s (**Figure 2.8C**). Participants significantly increased their average speller bitrates over days of training, echoing similar across-day learning as seen in the center-out tasks and exploration periods ($p=0.005$; **Figure 2.8D**).

Dimensionality as measured by the participation ratio significantly increased ($p<0.0001$) during the speller task relative to the isometric contraction task and was not significantly different than that of the center-out task ($p>0.05$), indicating participants used a strategy based on multidimensional independent motor unit control (**Figure 2.8E**). Participants' strategies for moving the cursor leaned more towards recruiting the 3 individual motor units exclusively of one another than simultaneously in combinations of two (mean CIFT speller > mean CIFT center-out, $p<0.0001$), agreeing with the increased difficulty observed during T4 targets in the center-out task requiring simultaneous unit activation (**Figure 2.8F**). Taken together, these results demonstrate the translational potential of an NMI enabling skilled independent motor unit control.




Figure 2.8 | Speller task. **A**, Overview of the speller. Left: the user interface displayed on a computer monitor. Participants navigated their cursor (black dot) via the activities of the same 3 selected motor units from the center-out task. The sentence to be typed was displayed at the top, with untyped letters grayed out. Any mistakenly typed characters had to be deleted with the “<” key before participants could proceed. Right: motor unit activities were translated into changes in velocity of the cursor, allowing the user to smoothly move the cursor across the screen. Keys were selected by co-contracting all three motor units in the same manner as the “T5” target from the center-out task; instances of this key selection highlighted and labelled in gray bars at top and gray circles on the keyboard. **B**, Smoothed bitrates for one participant’s 3 days of speller task. Dotted lines indicate average bitrate across that day’s speller task. **C**, Bitrates on the last day of training for all participants (dots). **D**, Participants’ speller task performances over days of training. Each line represents a participant-specific regression line, while the bold black line indicates the fixed-effect slope from the linear mixed model for this data ($p=0.005$, $n=24$). Color of each dot represents the duration of the speller task for that day. **E**, The participation ratio of firing rates during the speller task significantly increased relative to that day’s force-control task (paired t-test; $p<0.0001$, $n=24$) but was not statistically different than that of the center-out task (paired t-test; $p>0.05$, $n=24$). **F**, Mean CIFT metric computed within the 3 selected motor units for each speller task period increased relative to the mean CIFT during the center-out task (paired t-test; $p<0.0001$, $n=24$).

2.3 Discussion

Leveraging a non-invasive neurofeedback paradigm, we probed volitional control of individual motor units within the biceps brachii. Over 6 days of training, participants steadily improved performance in a center-out task requiring both exclusive and simultaneous control of three motor units. We found that the dimensionality of motor unit activity during this task exceeded that measured during stereotyped, isometric muscle contractions and provided compelling evidence that this independent motor unit control involved changes in motor unit recruitment order relative to such contractions. Finally, demonstrating an application of this NMI through a speller task, we

showed that participants could use this acquired motor unit control in a task requiring continuous, multi-dimensional control. Here we discuss the significance of these results for motor control theories and translational applications.

Skilled independent control of individual motor units

Volitional control of individual motor units was first reported in pioneering neurofeedback studies in the 1960s and 1970s (Virginia F. Harrison and Mortensen 1962; J. V. Basmajian 1963; V. F. Harrison and Koch 1972; Illyés 1977; John V. Basmajian and De Luca 1985). In these studies, the raw electrical signals measured from intramuscular electrodes were used to provide participants with visual and/or auditory neurofeedback signals on the underlying motor unit activity. Using this neurofeedback system in unstructured tasks similar to our exploration procedure, authors reported that participants were able to selectively activate individual motor units in the abductor pollicis brevis (J. V. Basmajian 1963), extensor digitorum (V. F. Harrison and Koch 1972), and the tibialis anterior muscles (Virginia F. Harrison and Mortensen 1962) and able to vary the firing rate of isolated motor units on command (John V. Basmajian and De Luca 1985). Despite this initial interest, the extent to which individual motor units can be volitionally controlled remained largely unclear, as these initial studies did not report quantitative measures of control quality or independence, assess independence of motor units by testing both exclusive and simultaneous motor unit activation, quantify the ability to modulate firing rates, examine learning capabilities over time, or compare periods of neurofeedback motor unit control with naturalistic, isometric contractions of the same muscle.

Here, we found that individual motor units can be controlled independently from one another and that control proficiency can be improved with training. In particular, we showed that over 6 days of training in a center-out task, participants progressively acquired skilled independent control of three motor units of the biceps brachii. This skilled control was evidenced by participants' ability to control each of the selected units' firing rate both exclusive of (T1, T2, and T3 targets) and in combination with other units (T4 targets) to achieve targets at different distances from the center of the screen. Critically, not shown in previous neurofeedback studies, participants' abilities to both exclusively and simultaneously activate the 3 selected motor units demonstrate the existence of multiple descending neural drives simultaneously converging onto a single motor pool. Through simultaneous recordings of non-selected motor units, dimensionality analysis also revealed that selected units were not activated in complete isolation but instead activated in conjunction with other units. Finally, through comparisons of center-out task firing rates to those observed throughout the comprehensive combinations of forearm supination and elbow flexion of the force-control task, we demonstrated that neurofeedback enabled motor unit activity patterns expanding beyond recruitment constraints previously thought to characterize neural control of motor units. These results present an unprecedented level of control over individual motor units belonging to the same muscle, bolstering the characterization of selective motor unit activation documented in previous studies and expanding our understanding of the flexibility possible in neuromuscular control.

Mechanisms of independent motor unit control

It is currently accepted that the orderly recruitment of motor units within a muscle applies, not to anatomically defined motoneuron pools as originally suggested by Henneman's size principle (Elwood Henneman 1957), but to function-specific motoneuron populations that can innervate multiple muscles and/or compartments within a single muscle (Borzelli et al. 2020; ter Haar Romeny et al. 1984; van Zuylen et al. 1988; John E. Desmedt and Godaux 1977; Duchateau and Enoka 2008; Thomas et al. 1986; Manning et al. 2010; Riek and Bawa 1992; Buchanan and Lloyd 1995; Wakeling 2009). In particular, the biceps brachii, used in this study, is known to have multiple anatomical neuromuscular compartments, with separate subdivisions even within the gross anatomical divide of the short and long head (Segal 1992). Biceps brachii motor unit recruitment has been shown to vary depending on the contraction levels in the flexion and/or supination directions, with motor units distributed across the biceps with no clear spatial distribution relevant to function (ter Haar Romeny et al. 1982a; Borzelli et al. 2020; ter Haar Romeny et al. 1984; Herrmann and Flanders 1998). Other muscles have been shown to have similar task-dependent recruitment order differences, such as in the first dorsal interosseus muscle when performing flexion versus abduction of the index finger and in a variety of non-multifunctional arm muscles (van Zuylen et al. 1988; J. E. Desmedt and Godaux 1981). In agreement with this existing body of literature, we found that biceps motor unit recruitment significantly differed between elbow-flexion and forearm-supination isometric contractions.

Participants' ability to control the selected motor units exclusive of and in combination with one another in the center-out task suggests that a minimum of 3 partially independent neural drives must converge to the biceps motoneuron pool. If the prevailing task-specific orderly recruitment model is correct, these neural drives should be associated with established motor primitives, such as elbow flexion and forearm supination for the biceps (ter Haar Romeny et al. 1984). Therefore, learning to exclusively recruit individual motor units should be equivalent to finding the appropriate motor tasks to perform, and any increases in center-out task performance should be attributed more to learning the association between these particular tasks and the computer cursor than to improvements in performance of these well-established primitives. Radhakrishnan et al. (Radhakrishnan et al. 2008) studied this type of learning in a center-out task similar to that of our study, in which participants learned to control a computer cursor through various arbitrary, non-intuitive combinations of upper-limb motor tasks, e.g., through simultaneous elbow flexion and index finger abduction. These participants learned the task and achieved a high-level performance plateau within 30 minutes. In stark contrast to Radhakrishnan et al.'s study, participants' performance in this study increased throughout several days of training with no significant decrease in overall learning rate over time, suggesting a different mechanism for independent motor unit control than recalling established motor primitives. Moreover, dimensionality of the selected units' firing rates in the center-out task increased throughout the 6 days of experiments, similarly signifying an activation strategy requiring refinement over time and thus not clearly related to executing stereotyped tasks. Observed combinations of firing rates

across the selected units significantly expanded during the center-out task relative to the force-control task, suggesting that individual motor unit control did not solely rely on the execution of elbow flexion and forearm supination motor primitives. Aligned with this observation, participants sometimes reported relying on abstract biceps contraction strategies that they were not able to precisely describe (**Supplementary Discussion 1**). Taken together, these observations suggest that the boundaries between neuromuscular compartments may not be as strictly defined by established motor primitives as previously thought.

Our results suggest the presence of some latent flexibility in motor unit recruitment order that allows for the formation of novel motor patterns. In particular, recruitment analyses showed that neurofeedback enabled participants to expand the available control space of the selected motor units beyond the robust low-dimensional manifold spanned during combinations of forearm-supination and elbow-flexion isometric contractions. Such flexibility could rely on selective pathways that bias motor unit recruitment in neuromuscular compartments otherwise controlled by a single descending neural drive. Selective recruitment mechanisms have been previously hypothesized to account for de-ordered motoneuron recruitment under certain conditions, as for example during ballistic (Grimby and Hannerz 1977) or lengthening (Nardone et al. 1989) muscle contractions or following cutaneous stimulation (Garnett and Stephens 1981). This selective motor unit activation has been hypothesized to arise from heterogeneously distributed excitatory input to the spinal motoneuron pool and/or through excitatory or inhibitory synaptic currents that bias pools of motor units (Kernell and Hultborn 1990). While there has been a lack of empirical evidence suggesting that these pathways are involved during established motor behavior (Bawa et al. 2014), we suggest that such mechanisms, enabled by neurofeedback, might underlie this study's observed flexibility in motor unit recruitment order. However, the presence of selective recruitment mechanisms should not be interpreted as a lack of orderly recruitment. On the contrary, the population-level increases in firing rate dimensionality during the center-out task emphasize the existence of constraints between motor units that could restrict which motor units are able to be selectively recruited, as the ability to selectively recruit every individual motor unit in the nervous system would be computationally infeasible (E Henneman et al. 1974). We, therefore, propose that both these mechanisms can influence neurofeedback-enabled motor unit recruitment and that the orderly recruitment of subgroups of motor units observed during isometric contractions may not be an immutable constraint of unit activation but rather be an emergent property of motor control. These observations are particularly aligned with the parallel work of Marshall et al. on flexible control of individual motor units (Marshall et al. 2021), which provides compelling evidence of de-ordered motor unit activity during complex motor behaviors and suggests that descending neural drives can contribute to flexible motor unit recruitments to adapt muscle activity to task demands.

While orderly recruitment of motor units maximizes the computational efficiency of the central nervous system during the production of a stereotyped output (E Henneman et al. 1974), additional flexibility in motor unit recruitment can enable the neuromuscular system to cope with the wide range of movement conditions needed for everyday life (Kernell and Hultborn 1990). Our

study sheds additional light on the ongoing debate on the generalization of orderly recruitment principles and the ultimate flexibility of the sensorimotor system (Bawa et al. 2014).

Potential translational implications

Our ability to skillfully control different muscles is tied to our ability to sense their state through proprioceptive organs (Tuthill and Azim 2018). Aligned with a motor control framework maximizing robustness and computational efficiency, these organs only provide information at the level of each muscle and not at the level of single motor units, making it impossible to naturally learn to skillfully control different motor units independently. By artificially augmenting the resolution of proprioceptive signals using neurofeedback, here we demonstrated that skilled and independent motor unit control is possible. In particular, the developed NMI enabled participants to expand the control space of individual motor units within a single muscle beyond that observed during isometric contractions typically studied in the biceps brachii. This increased number of independent degrees of freedom (DoFs) per muscle could enable some applications typically reserved for BMIs. In particular, in people with cervical spinal cord injury, where in the most severe cases only the muscles innervated by cranial nerves remain functional, existing NMIs provide limited utility compared to BMIs: most muscles with residual control perform critical motor functions and would not be useful for a traditional myoelectric application in which only natural motor primitives are mapped to the controlled end-effectors. However, by instead expanding the number of controllable DoFs within a single muscle, the proposed NMI could more effectively repurpose residual muscles with dispensable functions, such as the auricular muscles (Schmalfuß et al. 2015), or even the functionally-paralyzed muscles themselves if sufficient volitional motor unit control remains (Ting, Vecchio, et al. 2021).

The proposed NMI could also be advantageous for applications in able-bodied users. Indeed, by increasing the number of control signals that can be extracted per muscle, fewer muscles are required to operate the same number of actions, thereby potentially minimizing interference with functions needed for daily life. However, it remains to be assessed whether independent control of individual motor units can be achieved while multitasking and whether the added control dimensionality is sufficient to compensate for the increased cognitive demand and training required to operate such a NMI.

In addition to expanding the control dimensionality of individual motor units, the developed NMI enabled precise control of selected individual motor units both exclusive and in combination with other units. As illustrated in the speller task, this ability enables motor units to produce continuous, multi-dimensional control from a single muscle. Within the BMI domain, this form of control has been required to skillfully and intuitively operate complex devices and has enabled not only typing but also more general BMI applications, such as point-and-click navigation of a computer (Paul Nuyujukian et al. 2018) and control of multi-DoF robotic effectors (Collinger, Wodlinger, et al. 2013; Benabid et al. 2019). Since most non-invasive BMI implementations rely on discrete, task-specific control schemas due to the difficulty in non-invasively decoding continuous control signals (Tonin and Millán 2021), such control paradigms

have historically been confined to invasive BMIs, enabling invasive BMIs to far outstrip performance over their non-invasive counterparts (Benabid et al. 2019; Wang et al. 2013; Jarosiewicz et al. 2015; Pandarinath et al. 2017; Willett et al. 2021). However, recent surveys indicated 40% of surveyed patients with tetraplegia or paraplegia would not undergo implantation even if the implant restored daily function (Blabe et al. 2015; Collinger, Boninger, et al. 2013), indicating the dire need for a class of generally-applicable non-invasive devices that remains comparable to their invasive alternatives. Furthermore, while best-in-class EEG speller implementations have demonstrated much higher bitrates (3+ bits/s (X. Chen et al. 2015; Nakanishi et al. 2018)) than the preliminary demonstration here (0.43 bits/s), their reliance on exogenous stimulation to detect visual attention requires sustained concentration by the user and makes their performance unlikely to translate in self-paced, real-world applications. Indeed, when similar implementations are tested in ecological settings much lower performances are reported (Sellers et al. 2010). In particular, the bitrate achieved in this study over a few training sessions is comparable to the 0.31 bits/s throughput measured over 2.5 years of at-home, all-day use of an EEG speller (Sellers et al. 2010). Promisingly, our learning analyses indicating that motor unit control improves with more training suggests that the developed NMI might not only be more flexible than existing non-invasive BMI implementations but also may approach the throughput of these best-in-class implementations over time. An NMI enabling independent motor unit control, as the one demonstrated in this study, thus possesses promising qualities that could lead to a new class of high-performance, non-invasive translational devices.

Limitations and future directions

Stable, online detection of motor unit activity using non-invasive recording technologies remains challenging in ecological settings. The waveform of recorded motor unit action potentials and consequently their detection in surface EMG recordings are known to be sensitive to movement artifacts and to relative positioning of skin to muscle, which can be especially deleterious in anisometric conditions (Farina 2006). In our study, we overcame these limitations through physical constraints imposed by the orthosis and by instructing participants to avoid performing overt movements when trying to control the selected motor units. We confirmed these relative static recording conditions through kinematics recordings. In more dynamic settings, improved algorithms for motor unit detection may be required to increase reliability. Notably, while global EMG features are often used as a proxy for motor unit activity in non-invasive recordings, their lower information content is likely to hinder performance (Kapelner et al. 2019; Farina et al. 2017), as also suggested by our results showing dimensionality increases that are greater in motor unit firing rates than in iEMG. Alternatively, minimally-invasive intramuscular electrodes could enable individual motor unit recordings during anisometric contractions (Harwood et al. 2012). These advances in motor unit detection both could afford for a more practical understanding of the inherent dimensionality of a given muscle and are likely necessary for real-world use of a translational device controlled by individual motor units.

Our results suggest that neurofeedback enables an increase in the control space of individual motor units within the biceps brachii beyond that observed during isometric contractions, providing new insights on the flexibility of motor unit recruitment orders. However, it remains unclear whether the observed increase in control dimensionality represents an absolute increase in motor output dimensionality. In particular, while we were able to exclude isometric elbow flexion and forearm supination contractions as an explicit strategy to generate independent motor unit control, it is possible that participants' strategies involved subtle co-contractions of synergistic muscles to modulate recruitment, and thus the observed increase in biceps motor unit activity dimensionality may have come at the expense of the ability to control other muscles. Further studies capable of testing more ecological movement scenarios while simultaneously recording the activity of synergistic muscles are necessary to further explore the potential of the proposed NMI for motor augmentation.

The population-level dynamics across motor units observed in neurofeedback tasks suggest each dimension of the system can be driven by sets of motor units, as opposed to a single motor unit. This can increase robustness to experimental instabilities and can facilitate a finer-grained measurement of a dimension's amplitude by incorporating multiple units' firing rates. Similarly, the decoder can periodically be tuned to optimize for performance or for similarity to previously learned decoders, leveraging the fact that neural activity resides in a persistent, low-dimensional manifold (Sadler et al. 2014).

Additionally, this current study did not explicitly identify nor target selection of motor units that had been selected in previous days of training. The presence of within-day learning in our study and the intracortical BMI literature (Ganguly and Carmena 2009; Orsborn et al. 2014) both suggest that retaining similar sets of motor units over days may increase overall performance. This can be addressed by longitudinally tracking individual motor units over training and prioritizing selection of those units (Martinez-Valdes et al. 2017). Alternatively, chronically implanted intramuscular electrodes could enable recordings that stably identify motor units across days, though such a system has yet to be shown.

Finally, this study solely tested participants with no history of motor impairments, and so future studies motivated by clinical translation should be performed to determine the efficacy of a motor-unit NMI in people with sensorimotor disabilities. Promisingly, however, recent studies in people with cervical spinal cord injury demonstrated that residual activity in functionally-paralyzed muscles (Ting, Del Vecchio, et al. 2021) and impaired movements (Pierella et al. 2021) can be successfully harnessed for powering peripheral human-machine interfaces. Additionally, initial neurofeedback studies reported selective motor unit control in a variety of muscles, suggesting this study's results can generalize outside of the biceps brachii (John V. Basmajian and De Luca 1985).

Conclusion

In conclusion, we have demonstrated that skilled independent control of individual motor units belonging to the same muscle can be enabled through an NMI. We uncovered properties of

individual motor unit control useful for translational and augmentative applications. Concurrently, we shed light on long-standing questions surrounding the applicability of recruitment order often measured in stereotyped movements to volitional control of individual motor units. Advances in both motor control theory and neurotechnologies are critical to push the field towards more widely-applicable devices. Our study provides advances in both, potentially leading to improved therapeutics for people with sensorimotor disabilities and to a new class of neurotechnologies for human augmentation.

2.4 Methods

Experimental Procedures

All experiments were approved by the Committee for Protection of Human Subjects (CPHS) of University California, Berkeley, and were performed in compliance with local COVID-19 regulations. The recruited participants were healthy individuals — with no history of cognitive or sensorimotor impairments — between 22 and 30 years old, of which 3 were female and 5 male. Experiments were carried out on 6 consecutive days, with each session lasting a maximum of approximately 1 hour and 50 minutes.

Setup and initial calibration | At the beginning of each session, participants were seated on a chair and fitted with a sensorized orthosis that constrained the elbow joint at 100 degrees and the wrist at its natural position (**Figure 1.1A**). After cleaning the skin with a mildly abrasive paste and isopropyl alcohol, a high-density 64-channel grid of surface EMG electrodes (GR10MM0808, OT-Bioelettronica, Torino, Italy) was placed over the short and long heads of the biceps brachii, with the proximal/distal edges of the grid positioned at approximately 60%/80% of the distance between the acromion and the distal insertion of the biceps brachii tendon (Barbero et al. 2012). Velcro straps were used to ensure a tight fit of the orthosis around each participant’s arm. Markings on the skin were used to ensure stable grid positioning across days.

We next calibrated the decomposition model used to extract individual motor unit activity from the measured neuromuscular signals. This initial calibration was performed offline on a recording of 60 seconds, during which the participants were instructed to perform subtle biceps contractions that would activate only a few motor units. To help participants in this task, we educated participants in recognizing individual motor unit action potentials from displayed raw neuromuscular signals, and encouraged them to use this simple form of neurofeedback to gauge their muscle activity. Participants were then introduced to the exploration procedure.

Exploration procedure | A computer screen and headphones were used to provide participants with real-time auditory and visual neurofeedback of the detected motor unit activity (**Figure 2.1B**). Visual neurofeedback consisted of color-coded LED-like indicators that flashed when an action potential was detected and plots of the corresponding multi-channel waveforms. Auditory neurofeedback mapped detected action potentials into pitch-coded 150-ms-long stimuli.

Neurofeedback signals were updated at 60 Hz. Detected activity and corresponding neurofeedback signals were divided into three categories: selected units, unselected units, and unsorted activity. Selected and unselected units represented motor unit activity successfully classified by the decomposition model, while unsorted activity represented residual threshold-crossing events that were not matched with previously recognized motor units. Selected units were assigned to unit-specific neurofeedback features (i.e., colors and pitches), while those for unselected units and unsorted activity had categorical features.

Participants were instructed to use the provided neurofeedback signals to explore covert strategies to selectively recruit different motor units — mimicking pioneering studies on individual motor unit control (Virginia F. Harrison and Mortensen 1962; J. V. Basmajian 1963; V. F. Harrison and Koch 1972; Illyés 1977) — and had approximately 30 minutes to select and sort in order of controllability the 3 motor units to use in the center-out task. To guide participants in their motor unit selection, we designed an algorithm that monitored motor unit activity in real-time and suggested units showing substantial evidence of independent control. Participants could rely on this algorithm to automatically define which units to be included in the selected units category but could also include, exclude, and reorder units at will.

Throughout the exploration period, the decomposition model was periodically updated until a maximum of 25 different motor units were detected. Participants could thus use the unsorted-activity neurofeedback to steadily recruit unsorted units of interest and assist the update algorithm in detecting these units.

Center-out task | Participants controlled a computer cursor using the 3 motor units selected during the exploration procedure to achieve targets displayed on a screen. The activity of the selected motor units was mapped into the 2D position of a computer cursor using a population-coding strategy (**Figure 2.1C**). Each motor unit was assigned to a unique direction by dividing the 2D space into three equal subspaces (i.e., with a 120 degrees angle between each other) and provided a vectorial contribution to the cursor position along this direction and proportional to its normalized firing rate. To provide intuitive feedback on this control strategy, the cursor position was indicated by an arrow — representing the population vector — originating at the center of the screen. Motor unit firing rates were computed over a rolling window of 50 bins of 16 ms (800 ms in total) using a half-Hamming window profile that gave larger weight to the most recent bins. This firing rate was then normalized between 0 and the 90th percentile of the firing rate displayed during the exploration procedure. In some cases, this normalization value was manually adjusted between 10 to 20 Hz.

A total of 13 active targets and 1 rest target were designed. Active targets included 12 peripheral targets and 1 center target. Peripheral targets were defined by polar rectangular regions with a $\Delta\theta$ of 45° and Δr of 0.39 population-vector magnitude and were divided into exclusive targets (T1, T2, and T3) and simultaneous targets (T4), depending on their center angle: exclusive targets were centered on the assigned motor unit directions and thus required exclusive recruitment of an individual motor unit; simultaneous targets laid between the assigned directions and thus

required simultaneous recruitment of two units. Peripheral targets were also divided by distance: close targets were centered at 0.395, while far targets were centered at 0.785 magnitude. To achieve peripheral targets participants had to hold the cursor position within the target for a minimum of 0.5 seconds. The center target (T5) was defined by a circular region located at the center of the screen and had a radius of 0.2 magnitude. To achieve this target, participants were required to recruit all selected motor units at a minimum normalized firing rate of 0.33, while also keeping the cursor within the target boundaries. In contrast to active targets, the rest target required participants to avoid motor unit recruitment by holding the cursor within a distance of 0.1 from the screen center for 2 seconds.

The task was divided into trials and inter-trial periods. At the beginning of each trial, an active target was randomly selected from a pool of available targets and participants had 60 seconds to achieve it (**Figure 2.1D**). The rest target was then displayed and participants could initiate the next trial by completing it. To promote learning, active targets were grouped into 3 difficulty levels, which were progressively made available depending on participants' performance. At the beginning of each session, only the center target (T5) and the motor unit #1 and #2 exclusive targets (T1 and T2) were available. An algorithm monitored the average trial success rate over a window of 5 min and if this surpassed a threshold of 3 trials per minute, targets belonging to the next difficulty level were made available: T3 targets were added first, T4 targets last.

To promote engagement and incentivize learning, task and trial performance metrics were displayed on the task monitor. Finally, in addition to the arrow indicating the cursor position, participants received neurofeedback of the selected unit action potentials via the same LED-like indicators and audio stimuli utilized in the exploration procedure. Participants trained on this task for approximately 60 minutes per day during the first 3 days, and for a minimum of 30 minutes per day on the last 3 days of experiments.

Force-control task | Participants were instructed to perform isometric elbow flexion and forearm supination contractions to match target force profiles displayed on a computer screen. The forces measured by the sensorized orthosis were displayed in real-time by a bar indicator (**Figure 2.5A**). Target forces followed a trapezoidal profile — with onset, hold, and offset periods each trial — and were displayed adjacently to the measured forces. To prepare participants for a change in force profile, the target force expected 1-second ahead was also displayed.

Three isometric contraction types were tested: elbow flexion, forearm supination, and simultaneous elbow flexion and forearm supination. Each contraction type was tested 5 times at 3 different loads, for a total of 45 trials. Loads of 500, 1000, and 1500 grams were default but in some cases decreased to avoid fatigue (lowest maximum load of 1000 g). Trials were separated by 2 seconds of rest period. Onset and offset ramps each occurred over 1, 2, and 3 seconds for the three total loads, respectively. Hold duration was 1 second in all trial types. Trials of different types were ordered randomly.

Speller task | The same 3 motor units from the center-out task were used to operate a cursor to navigate a virtual keyboard in a copy-typing speller task. The keyboard layout (OPTI-II) and target sentences mimicked those of previous BMI studies (Pandarinath et al. 2017; Rick 2010). The keyboard divided the screen in 30 square keys (6x5) and included all the alphabet letters, 2 space keys, and 2 delete keys; misselection of a character required participants to select the delete key.

To facilitate navigation, the keyboard featured wraparound borders and the cursor was controlled in velocity. In particular, the population vector used in the center-out task to compute the cursor position was here used to control the cursor velocity. These design features allowed full 2D space navigation even with control of a single motor unit, though this would result in extremely low performances. Letter selection was triggered by simultaneously recruiting the 3 selected motor units above a threshold normalized firing rate and for a threshold amount of time — similar to how center-out T5 targets were achieved. Firing rate and time thresholds were default to 0.5 Hz and 0.5 seconds, and sometimes slightly adjusted according to participants' preference.

Participants were tested in this task for approximately 30 minutes in the last 3 days of experiments, after training for a minimum of 30 minutes in the center-out task and reaching sufficient proficiency. 4 participants also tested this task prior to the 3rd day, but only after completing a minimum of 60 minutes of center-out task. 1 participant only completed 1 day of the speller task.

Motor unit NMI

EMG recordings | Biceps brachii EMG signals were acquired using a PZ5M neurodigitizer amplifier and an RZ2 bioamp processor from Tracker-Davis Technologies (TDT) at 12.2 kHz. The 64-channels grid of electrodes was connected via 32-channels ZIF-clip TDT headstages and Omnetics connectors. Signals were band-pass filtered between 10 and 900 Hz using a 6th order Butterworth filter. Notch filters at 60, 120, 180, and 240 Hz were also used to remove the powerline noise. Filtered signals were then used to extract 56 bipolar derivations parallel to the muscle fibers. A multichannel threshold crossing algorithm was then used to detect time windows with potential motor unit activity; thresholds were set to 6 times the signal's standard deviation and were calibrated at the beginning of each session using 10-second recordings during which participants were instructed to avoid biceps contractions and not move. A threshold crossing event at any of the bipolar channels triggered a shared dead-time of 20 ms that limited overall threshold crossing detection rate at 50Hz. Critically, since multiple motor units could be detected per threshold crossing, this capped detection rate only limited the maximum detectable firing rate of individual motor units and not that of the entire motor unit population (see **Decomposition model**). In addition, since firing rates of biceps brachii motor units during isometric contractions range between 5 and 30 Hz (Kukulka and Clamann 1981), this detection rate was largely sufficient to cover motor units' entire dynamic range. Threshold crossing events and filtered bipolar signals were downsampled to 2 kHz and streamed to the decomposition model. All these processing steps

were performed using custom software written for the RZ2 bioprocessor, which ensured a maximum of 0.5 ms delay between acquisition and streaming.

Decomposition model | Bipolar EMG signals were decomposed into motor unit activity using a convolutive blind source separation model. This model included a previously validated offline EMG decomposition model (Negro et al. 2016) and shared similar logic to recent techniques for online EMG decomposition (Barsakcioglu et al. 2020).

The offline decomposition model used convolutive blind source separation to define the motor units underlying the measured EMG signals (Negro et al. 2016). Briefly, the filtered bipolar EMG signals were extended and whitened. An extension factor of 16 was used (Negro et al. 2016). Next, a 2-step iterative algorithm was used to find sparse components that best reconstructed the whitened data. First, a fixed-point iteration algorithm was used to estimate the next component using the logarithm of the hyperbolic cosine as a contrast function to optimize sparseness and an orthogonal constraint to promote estimates of unique sources. The logarithm of the hyperbolic cosine was used because of its superior robustness to outliers compared with simpler contrast functions (Negro et al. 2016). Second, an iterative algorithm was used to minimize the variability of the inter-spike intervals of detected spike-trains. After projecting the data onto the candidate component, K-means++ ($k=2$) was used to estimate a threshold on the peaks in the squared projected data. The estimated component was then refined according to those peaks. This process repeated until the inter-spike interval converged. Since the coefficient of variation for spike-trains generated by multiple motor units are intrinsically more variable than those generated by a single motor unit, this second step was shown to ameliorate source estimation by exploiting the regularity of motor unit firings (Negro et al. 2016). The resulting component was then added to the matrix of estimated components if the signal to noise ratio (SNR) of the spikes detected along this component was greater than a fixed threshold; SNR was measured using the Silhouette coefficient and a threshold of 0.85 was used (Negro et al. 2016). This iterative algorithm, which is described in greater detail in Negro et al., 2016 (Negro et al. 2016), was repeated until a maximum of 25 sources were detected. A post-processing step was then introduced to further de-duplicate the number of components underlying the same motor units. Indeed, despite the orthogonal constraint used in the fixed-point algorithm to increase the number of unique estimated sources, this approach can lead to components capturing delayed versions of the same motor unit action potentials (Farina et al. 2017; Negro et al. 2016). Spike-trains were thus extracted from each estimated component and only components with less than 30% of coincident spikes — as measured by the rate-of-agreement (Negro et al. 2016, 20) across spike timings — were kept. Note that while a minor inconvenience in offline analyses, an excessive number of duplicated components would largely impact computational load required by our NMI.

The offline model was initialized on the 60-second dataset acquired at the beginning of each session. This calibration was used to compute the whitening matrix and to initialize the decomposition matrix with the first set of estimated sources. This whitening matrix was then fixed for the remainder of the session. A batch update algorithm was then used throughout the

exploration procedure to periodically update the decomposition matrix with potential new components. To optimize computational efficiency and allow for quick model updates (update time < 30s), instead of using the full EMG stream this algorithm only ran on the windows of EMG signals surrounding the detected threshold crossing events (10 ms before the peak multichannel amplitude and 20 ms after). The update algorithm was triggered every 750 threshold crossing events with no extracted motor units and ran until a maximum of 25 total motor units were detected or until the end of the exploration procedure.

Individual motor unit activity was continuously estimated in real-time through this decomposition model from the streamed EMG signals using the 30ms windows surrounding detected threshold-crossing events. For each threshold crossing, data windows were whitened, extended, and projected to the source space by multiplying each extended multichannel sample with the most current decomposition matrix. A motor unit was then considered detected if at least a sample of the squared projected data exceeded the decomposition model's threshold for a given source, determined with k-means during the offline decompositions. Using this algorithm, multiple units could be detected from one threshold crossing event. If the projected activity did not surpass any component's threshold, the event was then classified as unsorted activity. Notably, windows alignment to a particular peak had negligible impact on the detection of concurrent motor unit action potentials, even when considering that the threshold crossing detection rate was capped at 50 Hz (see **EMG recordings**). Indeed, motor unit classification was based on 8 ms data bins (extension factor = 16; sampling rate = 2000 Hz) and the overlap between windows at high threshold-crossing event rates (max 10 ms at 50 Hz) ensured detection of motor unit action potentials even when occurring between two consecutive events.

Online and offline decomposition models were implemented through custom-written GPU-accelerated Python programs. All data was streamed between multiple computers with minimal latency and high bandwidth through River (Botros 2021), an open-source C++ library based on Redis. Overall latency from data acquisition to motor unit activity detection was generally under 70ms.

Motor unit selection algorithm | This algorithm monitored the dimensionality of motor unit activity throughout the exploration procedure and suggested motor units with potential for independent control. A circular buffer (size of 2^{16} samples) was used to collect sorted motor unit activity. The firing rate of each motor unit was then computed over overlapping windows of 1 s with 500 ms overlap. Non-negative Matrix Factorization (NMF) was then used to detect motor units explaining most firing rates variance. First, components required to explain a minimum of 90% of the total firing rates variance were selected. Second, the motor unit with the largest weight for each of the selected components was chosen and used to update the subset of suggested motor units. Suggested motor units were updated every 20 seconds.

Force and kinematic recordings | The sensorized orthosis was custom-designed and 3D printed using a Form 2 (Formlabs, Somerville, MA) printer with standard resin. The orthosis embedded 2

load cell sensors (a CB6 from DACELL, Korea and a TAL220 from HT Sensor, China) to measure elbow-flexion and forearm-supination forces, respectively, and inertial measurement units (BNO055, Bosch Sensortec, Germany) to capture movements. Load and IMU signals were sampled at 50 Hz using a Raspberry Pi 4. A HX711 analog-to-digital converter (Avia Semiconductor, China) was used to acquire the load data. Data was streamed online to other modules using River.

Behavioral Analysis

Center out task day 1 | Center out performance at day 1 was evaluated using the percentage of successful trials for each target category (T1, T2, T3, T4, and T5). A trial was considered failed if the presented target was not achieved within the 60s of trial and successful otherwise. Participants that did not reach the second and third difficulty levels were excluded when analysing the corresponding target categories (T3 and T4 respectively). Hypothesis testing was performed using bootstrapping (n=10000 iterations) and Bonferroni correction for multiple comparisons (**Figure 2.2C-D**).

Trial performance metric | While the percentage of successful trials allows to evaluate whether independent motor unit control is possible, this metric fails to capture the quality of this control. A more holistic performance metric was thus computed to assess motor unit control quality and evaluate learning over time. This metric combined together 3 independent metrics using Principal Component Analysis (PCA). The normalized distance between the cursor position with respect to the presented target center was calculated for every time point within each recorded trial; normalization was performed with respect to the maximum target distance. The average and integral of this distance for each trial were then linearized using a log transform. These metrics were used to capture the cursor error and trial duration and were the first 2 independent metrics. The third metric was used to reward motor unit specificity. A specificity score was first computed for each trial's time point as a value between -1 and 1, where -1 corresponds to selective recruitment of motor units that are not required for achieving the considered target and 1 to selective recruitment of the target motor units. The mean specificity was then calculated for each trial and linearized using the logistic transform. A PCA model was then fit on all the collected trials to combine these 3 metrics; the single holistic metric was then the first component of this PCA model, standard scaled to improve interpretability of the results. **Figure 2.2A** and **B** show how this holistic metric relates to the 3 underlying metrics prior linearization, as well as to the trial duration — a feature commonly used for evaluating performances in trial-based tasks. Feature linearization was performed to conform with the assumptions of the statistical techniques used to analyze learning over time. These analyses excluded T5 targets.

Learning analyses | Collected center-out data are characterized by hierarchical and crossed dependencies: trials (at the first hierarchical level) are grouped in days (at the second level) and in

participants (at the third level), while target categories are crossed at all hierarchical levels. To account for these dependencies, learning analyses were performed using linear mixed-effects models (LMMs) — an extension of linear regression models that allow to separate the overall effects of a model term (i.e. the fixed effects) from the variability in the data generated by different sources of stochastic variations (i.e., the random effects) (Lazega and Snijders, n.d.).

When analyzing the overall within- and across-day learning (**Figure 2.3C-D**), trial performance was modeled by the following equations (2.1-2.10) representing our general LMM:

$$y_{jti\theta r} = \beta_{jt\theta r}^0 + \beta_{jt\theta}^1 \text{within}_{jti\theta}^{\text{cwc}} + \beta_{j\theta}^2 \text{across}_{jt\theta}^{\text{awc}} \quad (2.1)$$

$$+ \gamma^3 \text{within}_{jti\theta}^{\text{cwc}} \text{across}_{jt\theta}^{\text{awc}}$$

$$+ \gamma^4 \text{within}_{jt\theta}^{\text{mean}} + \gamma^5 \text{within}_{jt\theta}^{\text{mean}} \text{across}_{jt\theta}^{\text{awc}} + \epsilon_{jti\theta r}$$

$$\beta_{jt\theta r}^0 = \gamma^0 + \mu_j^{00} + \mu_{jt\theta}^{01} + \mu_{j\theta}^{02} + \mu_{\theta}^{03} + \mu_{\theta r}^{04}, \quad (2.2)$$

$$\beta_{jt\theta}^1 = \gamma^1 + \mu_{jt\theta}^{10}, \quad (2.3)$$

$$\beta_{j\theta}^2 = \gamma^2 + \mu_{j\theta}^{20}, \quad (2.4)$$

$$\epsilon_{jti\theta r} \sim N(0, \sigma_e^2), \quad (2.5)$$

$$\mu_j^{00} \sim N(0, \sigma_{\mu_j^{00}}^2), \quad (2.6)$$

$$\mu_{\theta}^{03} \sim N(0, \sigma_{\mu_{\theta}^{03}}^2), \quad (2.7)$$

$$\mu_{\theta r}^{04} \sim N(0, \sigma_{\mu_{\theta r}^{04}}^2), \quad (2.8)$$

$$\begin{pmatrix} \mu_{jt\theta}^{01} \\ \mu_{jt\theta}^{10} \end{pmatrix} \sim N \left(0, \begin{bmatrix} \sigma_{\mu^{01}}^2 & \rho_2 \sigma_{\mu^{01}} \sigma_{\mu^{10}} \\ \rho_2 \sigma_{\mu^{01}} \sigma_{\mu^{10}} & \sigma_{\mu^{10}}^2 \end{bmatrix} \right), \quad (2.9)$$

$$\begin{pmatrix} \mu_{j\theta}^{02} \\ \mu_{j\theta}^{20} \end{pmatrix} \sim N \left(0, \begin{bmatrix} \sigma_{\mu^{02}}^2 & \rho_1 \sigma_{\mu^{02}} \sigma_{\mu^{20}} \\ \rho_1 \sigma_{\mu^{02}} \sigma_{\mu^{20}} & \sigma_{\mu^{20}}^2 \end{bmatrix} \right) \quad (2.10)$$

where j , t , i , θ , and r refer to the participant, day, trial, target angle, and target distance indexes, respectively; γ^n refers to the fixed effect estimated for the n^{th} independent variable; μ_{xyz}^{nl} refers to the l^{th} random effect for the n^{th} independent variable caused by the random factor xyz ; β_{xyz}^n refers to the combined random and fixed effects; and $\epsilon_{jti\theta r}$ refers to the model residuals. This model describes trial performance $y_{jti\theta r}$ as a function of within- ($\text{within}_{jti\theta}^{\text{cwc}}$) and between-day ($\text{across}_{jt\theta}^{\text{awc}}$) time variables, an interaction term between these 2 variables ($\text{within}_{jti\theta}^{\text{cwc}} \text{across}_{jt\theta}^{\text{awc}}$), and two additional variables used to control for potential across-day effects caused by differences in number of performed trials ($\text{within}_{jt\theta}^{\text{mean}}$ and $\text{within}_{jt\theta}^{\text{mean}} \text{across}_{jt\theta}^{\text{awc}}$). The within-day time variable $\text{within}_{jti\theta}^{\text{cwc}}$ was calculated as the centered, normalized trial index i . For each day t , subject j , and target direction θ , trials were centered with respect to half of the performed trials. Such centering within-cluster (CWC) was used to segregate within-day effects from higher order effects. A normalization factor of 100 trials was used. The subtracted means from $\text{within}_{jti\theta}^{\text{cwc}}$ were in turn CWC centered and included in the model through the $\text{within}_{jt\theta}^{\text{mean}}$ term, which was used to account

for possible changes in performances caused by the different number of performed trials for each recording. The across-day variable $across_{jt\theta}^{awc}$ consisted of the aligned and normalized day index t . Alignment was performed within-cluster (AWC) with respect to the first day t for which participant j performed θ targets. While for targets belonging to the first difficulty level (i.e., T1 and T2 targets) AWC had no effect, this alignment strategy allowed to take into account participants' across-day heterogeneity in reaching T3 and T4 targets, effectively comparing across-day performances with respect to the number of days of practice instead of those of experiment. This variable was normalized with respect to 6 days. Maximal random effects were used to minimize Type I errors during hypothesis testing (Barr et al. 2013). Random effects included: random intercepts for each participant (μ_j^{00}), target direction (μ_θ^{03}), combination of participant and target direction ($\mu_{j\theta}^{02}$), combination of participant, target direction, and day ($\mu_{jt\theta}^{01}$), and combination of target direction and distance ($\mu_{\theta r}^{04}$); and random slopes for both the within- and across-day time variables ($\mu_{jt\theta}^{10}$ and $\mu_{j\theta}^{20}$, respectively). Random effects were modeled as 0-centered Normal distributions with estimated standard deviations σ and optional correlation parameter ρ . The centering and alignment choices used for $within_{jti\theta}^{cwc}$ and $across_{jt\theta}^{awc}$ made the fixed-effect of the model intercept γ^0 to capture the overall performance of a general participant on the center-out task at day 1. The modeled fixed effects for the within- and across-day time variables represented the overall improvement in performance a general participant would obtain in the center-out task by training over 100 trials and 6 days, respectively.

Learning analyses performed for each of the selected motor units separately (**Figure 2.3E**) were carried out using a similar LMM, which included the same fixed-effect terms but reduced random-effects:

$$y_{jtir} = \beta_{jtr}^0 + \beta_{jtr}^1 \text{within}_{jtir}^{cwc} + \beta_{jtr}^2 \text{across}_{jtr}^{awc} + \gamma^3 \text{within}_{jtir}^{cwc} \text{across}_{jtr}^{awc} + \gamma^4 \text{within}_{jtr}^{\text{mean}} + \gamma^5 \text{within}_{jtr}^{\text{mean}} \text{across}_{jtr}^{awc} + \epsilon_{jtir} \quad (2.11)$$

$$\beta_{jtr}^0 = \gamma^0 + \mu_j^{00} + \mu_{jtr}^{01} + \mu_{jr}^{02} \quad (2.12)$$

$$\beta_{jtr}^1 = \gamma^1 + \mu_{jtr}^{10} \quad (2.13)$$

$$\beta_{jtr}^2 = \gamma^2 + \mu_{jr}^{20} \quad (2.14)$$

$$\epsilon_{jtir} \sim N(0, \sigma_e^2) \quad (2.15)$$

$$\mu_j^{00} \sim N(0, \sigma_{\mu_j^{00}}^2) \quad (2.16)$$

$$\begin{pmatrix} \mu_{jtr}^{01} \\ \mu_{jtr}^{10} \end{pmatrix} \sim N \left(0, \begin{bmatrix} \sigma_{\mu^{01}}^2 & \rho_0 \sigma_{\mu^{01}} \sigma_{\mu^{10}} \\ \rho_0 \sigma_{\mu^{01}} \sigma_{\mu^{10}} & \sigma_{\mu^{10}}^2 \end{bmatrix} \right) \quad (2.17)$$

$$\begin{pmatrix} \mu_{jtr}^{02} \\ \mu_{jtr}^{20} \end{pmatrix} \sim N \left(0, \begin{bmatrix} \sigma_{\mu^{02}}^2 & \rho_1 \sigma_{\mu^{02}} \sigma_{\mu^{20}} \\ \rho_1 \sigma_{\mu^{02}} \sigma_{\mu^{20}} & \sigma_{\mu^{20}}^2 \end{bmatrix} \right) \quad (2.18)$$

$$(2.19)$$

where terms follow the same conventions as in the previous model. In particular, since different models were used to evaluate learning over T1, T2, and T3 targets, random effects that were used to account for variations caused by different target directions were removed. Random slopes for the within-day term were computed for each combination of participant, day, and target distance, while random slopes for the across-day term were computed for each combination of participant and target distance.

Analyses of participants' performance on the T4 targets were conducted using a generalized linear mixed-effects model (GLMM) with a Poisson link function. Specifically, the rate of successful T4 trials over time was modelled as:

$$\log \left(\frac{\text{successful_trials}_{jt}}{\text{task_duration}_{jt}} \right) = \beta_j^0 + \beta_j^1 \text{day}_{jt} + \epsilon_{jt} \quad (2.20)$$

$$\beta_j^0 = \gamma^0 + \mu_j^{00} \quad (2.21)$$

$$\beta_j^1 = \gamma^1 + \mu_j^{10} \quad (2.22)$$

$$\epsilon_{jt} \sim N(0, \sigma_e^2) \quad (2.23)$$

$$\begin{pmatrix} \mu_j^{01} \\ \mu_j^{10} \end{pmatrix} \sim N \left(0, \begin{bmatrix} \sigma_{\mu^{00}}^2 & \rho_0 \sigma_{\mu^{00}} \sigma_{\mu^{10}} \\ \rho_0 \sigma_{\mu^{00}} \sigma_{\mu^{10}} & \sigma_{\mu^{10}}^2 \end{bmatrix} \right) \quad (2.24)$$

where terms follow the same convention as above, $\text{task_duration}_{jt}$ indicates the duration in hour of the center-out task performed by participant j at day t , and day_{jt} indicates the t^{th} experiment day of participant j .

All models above parameters were fitted using the restricted maximum likelihood (REML) approach. Confidence intervals used for hypothesis testing were computed using the profile method. Model assumptions were tested using the White's Lagrange Multiplier test, for testing heteroskedasticity of the residuals, and the D'Agostino and Pearson's test, for testing residuals Normality. All models (general, T1, T2, T3, and T4 models) displayed homoscedastic residuals ($p=0.08$, $p=0.9$, $p=0.3$, $p=0.9$, and $p=0.07$, respectively), but only the residuals for the GLMM resulted normally distributed ($p=0.65$ for the T4 model, $p<0.001$ for the others). However, LMMs have been shown to be highly robust to violations of distributional assumptions and the kurtosis ([1.2, 0.77, 0.5, 2.5]) and skewness ([0.23, -0.19, 0.15, 0.7]) of our models with non-normal residuals' fell largely within acceptable ranges, shown to have minimal impact on the validity of LMMs estimates (Schielzeth et al. 2020).

Kinematic analyses | Measured IMU Euler angles were preprocessed using an artifact removal algorithm and a 6th order Butterworth low-pass filter at 6 Hz. Artefact removal was used to ignore samples with prominence superior to 10° , which accounted for less than 0.1% of all samples. Principal Component Analysis (PCA) was then used to align the rotational axis of the IMU sensor to the axis of largest variation. Kinematic analyses during the center-out task (**Figure 2.7B**) used

the 1st principal component to compute the mean absolute velocity (MAV) for trial and inter-trial periods. The median MAV was then computed for each day, each participant, and trial category and used to evaluate target-specific movement strategies. Statistics of active targets were compared with respect to those of rest targets; hypothesis testing was performed by bootstrapping ($n = 10000$ iterations) the distribution of the paired differences for each recording and using Bonferroni correction of the estimated confidence intervals for multiple comparisons. Decoding analyses were performed using a multilayer perceptron (MLP) with 1 hidden layer of 100 artificial neurons having hyperbolic tangent as activation function. For each session, a MLP was trained to predict the cartesian coordinates (x, y) of the center-out task cursor from the measured Euler angles and their derivatives. Time bins during rest trials were excluded. Data were then shuffled and split in training (70%) and testing (30%) sets. Decoder performances were evaluated on the testing sets using the average R^2 of the 2 output variables (**Figure 2.7C**).

Speller | A common metric for assessing information throughput in self-paced BMIs is the achieved bitrate, which combines the number of possible symbols to select (i.e. the number of characters on a keyboard) with the net number of correct symbols selected per second (P. Nuyujukian et al. 2015). This metric is typically considered an underestimate of the true information throughput of a device, as it penalizes errors relatively harshly compared to other information throughput metrics (P. Nuyujukian et al. 2015). It is defined as:

$$B = \frac{\log_2(N) \max(S_c - S_i, 0)}{t} \quad (2.25)$$

where S_c is the number of correct symbols transmitted, S_i the number of incorrect symbols transmitted, and N the number of symbols. In our case, $N = 27$, due to the 26 letters and the “space” character on the keyboard (excluding the delete key). Smoothed bitrates (**Figure 2.8B**) were computed from 5-minute sliding windows taken every 30 seconds; peak bitrate was the maximum smoothed bitrate value during a given session. Average bitrate was the achieved bitrate B computed over the entire spelling session. Correct characters per minute were computed similarly as the net number (correct symbols minus incorrect symbols) of correct characters spelled. Changes in average bitrate over days of training (**Figure 2.8D**) were modelled with a linear mixed-effects model where the number of days of training were centered within-subject to account for differences in amount of training between subjects. This model fit a fixed-effect slope and intercept for days of training and was fit using the restricted maximum likelihood (REML) approach. Model assumptions were tested as described in the above learning analyses.

Motor unit activity analysis

Pooled motor unit decomposition | A separate offline motor unit decomposition was run for the EMG collected during the force-control task with the same parameters as the decompositions run online. Then, for each day, the motor units identified across both the online and offline decompositions were pooled together, and all of the EMG data for that day was then re-decomposed with these motor units, yielding a superset of motor unit action potential timings

relative to those detected online. Motor units exhibiting more than 30% of coincident spikes, according to the rate-of-agreement between action potential timings, were considered duplicates, and only one of the duplicate units was retained. No duplicates were found within selected motor units in any session. All analysis that used firing rates (**Figures 4-6, 8**) uses these pooled motor units. This methodology allowed for motor units to be identified for analysis purposes even when they had not been identified during the online sessions.

Integrated EMG and motor unit firing rates | The integrated EMG (iEMG, **Figure 2.5**) for channel i at time t was computed as the sliding window sum of rectified EMG:

$$\text{iEMG}_i(t) = \sum_{j=t-N}^t |\text{EMG}_i(j)| \quad (2.26)$$

where N was fixed as the number of samples corresponding to a 200 millisecond window. The data was then downsampled by a factor of 25 to approximately 81 Hz. Smoothed motor unit firing rates were computed from the pooled motor unit firings and were computed in the same manner as computed online for the center-out task.

Firing rates for pooled motor units were computed as in the center-out task. For analysis based on firing rates during the center-out task (**Figures 5-6**), any time bins occurring during T5 or rest trials were excluded. For analysis during the speller task (**Figure 2.8**), time bins used for letter selection were explicitly excluded as well. When necessary, both firing rate and iEMG were linearly interpolated in time in order to align with other streams of data (e.g. aligning with load sensor data).

Exploration Period Analysis | In order to identify groups of units that were often mutually active during the exploration period, motor unit activity was decomposed into 3 separate components via non-negative matrix factorization (NMF). NMF aims to find two low-rank matrices, W and H , from a non-negative data matrix X such that

$$\frac{1}{2} \|X - WH\|_2^2 \quad (2.27)$$

is minimized and such that W, H are also nonnegative. NMF was performed via a coordinate descent solver with NNDSVD initialization. Since the relative scales of the projections (W) and its components (H) are typically arbitrary, we resolved ambiguity by scaling each component to unit L2-norm and scaling its corresponding transformation by the appropriate reciprocal factor. We then computed the CIFT for each of the 3 components relative to one another, as described in the following section.

CIFT metric | For analysis of the speller and exploration period, a simple time-based metric, the cumulative independent firing time (CIFT), was devised (**Figure 2.4**). CIFT is defined as the fraction of total time a motor unit was independently active relative to the total time the motor unit

was active, and thus takes values between zero and one. A motor unit was considered “active” if its smoothed firing rate exceeded 5 Hz, and was considered “independently active” if both it was active and no other motor units had firing rates simultaneously exceeding 5 Hz. This 5 Hz threshold corresponds to the approximate physiological minimum motor unit firing rate (Duchateau and Enoka 2011). Throughout this analysis, we utilize the CIFT as a general measure of relative independence of motor units and use it across various contexts (**Figures 4, 8**). Note that our use of CIFT in the exploration period extends its use from comparing motor units to comparing NMF components.

Dimensionality Computation | The participation ratio (PR) was computed to quantify the dimensionality of the iEMG and firing rate data (Recanatesi et al. 2019; Gao et al. 2017b). The PR is a metric computed on the covariance matrix of a feature and represents the approximate dimensionality of the manifold spanned by that feature; a higher participation ratio means more principal components are needed to explain a given proportion of the feature’s variance. Participation ratio is defined as:

$$\text{PR} = \frac{(\sum_i \lambda_C(i))^2}{\sum_i (\lambda_C(i)^2)} \quad (2.28)$$

where C_i is the i -th eigenvalue of the covariance matrix C of the corresponding feature (iEMG or firing rates). Participation ratio was computed across the periods spanning the force-control tasks, center-out tasks, and speller tasks (**Figures 5, 8**). In our data, the participation ratio corresponded to the number of principal components needed to explain approximately 83% of the total feature variance.

The relationship between selected and unselected motor units during the force-control and center-out tasks was characterized using linear regression (**Figure 2.5F-G**). Linear regression was used to predict the unselected motor units’ activity from the activity of the selected ones. The quality of this prediction was characterized by the coefficient of determination (R^2).

Recruitment Thresholds | Recruitment thresholds for each motor unit were computed for both elbow flexion and wrist supination from force-control task data (**Figure 2.6**). First, force data from load sensors was smoothed with a median filter and normalized within each session to values between zero and one. Then, for force-control task trials in which elbow flexion (forearm supination) was the sole movement indicated, the recruitment threshold for a particular motor unit for elbow flexion (forearm supination) was identified as the average across trials of the measured load at the beginning of the first occurrence of 3 consecutive firings with inter-spike interval (ISI) less than 200ms.

Control Spaces | To compare possible joint firing rates observed across the 3 motor units selected for the center-out, a control space was computed for firing rates spanning the center-out and force-

control tasks. First, each unit's firing rate was normalized to 1.5 times the normalization value used for that particular unit in the center-out task, in order to focus analysis on the same, task-relevant firing rates between tasks, firing rates above this normalization value were discarded. Normalized firing rates were then binned into 1 of 10 equally-spaced bins between 0 and 1, with any time points exceeding a normalized value of 1 ignored. Each time point was thus converted into a 3-dimensional point representing the firing rate bins of each of the 3 selected motor units. From this, control spaces for the center-out and force-control tasks were then computed as the set of joint firing rate bins that occurred for at least 160ms (i.e., 10 bins) in a given task. The maximum number of populated bins, i.e. the largest possible control space, was thus 10^3 .

To facilitate 2-D visualization of the 3-D control space in **Figure 2.6D**, non-negative matrix factorization, as used in the exploration period analysis, was performed for each session on the force-control's firing rates to compute 2 non-negative components that best factorized the data. Components were ordered such that component #1 captured more variance in the data than component #2. These components were then normalized for the day such that the maximum projected value across either component was 1. Finally, firing rates for the force-control and center-out tasks were then projected onto these 2 components. After this normalization, the control space was then computed as above, i.e. component projections were binned between 0 and 1 and counted.

Ordinary least-squares was used to determine the best-fit line between mean center-out task performance and the relative fraction of control space bins that were novel in the center-out (**Figure 2.6F**). For performance analysis focusing only on T1 and T2 targets, the above control space computation was repeated with only the first 2 selected motor units as opposed to all 3.

Statistics

Statistical tests, their significance values, and the relevant number of samples are reported in the appropriate figure legends and/or relevant method section. Error bars used in point-plots represent 95% confidence intervals. No data were excluded from the analyses, unless specifically reported.

2.5 Supplementary Material

Supplementary Discussion 1 | Subjective strategies reported by participants

Strategies used to activate motor units independently widely varied between participants. To qualitatively record their experience and strategies, a free-form text form was filled out by each participant at the conclusion of each of their training days. The majority of subjects incorporated forearm supination and elbow flexion into their strategies, though not often did these movements in isolation correspond to isolated activation of a particular selected motor unit. Instead, participants reported performing subtle combinations of flexion, extension, supination, and/or pronation to achieve independent motor unit control across the 3 motor units, often combined with some level of generating "tension" in their biceps. Imagining naturalistic movements also seemed to prove a useful strategy; some participants reported imagining "bringing [their] hand towards [their] chest", "reaching for a cup", or "raising [their] arm while contracting the biceps" to activate

particular motor units in isolation. One participant reported their strategy as “abstract contractions” of their biceps, reporting that they “didn’t know what [they] were doing” but “would try to just keep doing the same thing”. Participants reported that thorough exploration and careful selection of motor units during the exploration procedure was critical for center-out and speller task performance, though time spent in the exploration procedure should be balanced against limits on attention and muscle fatigue.

Strategies within a day remained mostly stable for participants; however, on some occasions, participants reported the amount of force required to activate a unit increased over time or that the motor units activated by particular strategies would switch or no longer activate. The subjective difficulty of activating a motor unit also occasionally modulated throughout a session, with some reports of “momentum” where activating a motor unit in subsequent trials was easier if it was just activated in the previous trial. Across days, consistency in strategies was variable: participants reported qualitatively different strategies across most days. Interestingly, the two participants with the highest center-out task performance had high consistency in their strategies across their last 3 days of training.

Supplementary Video 1 | Center-out task demonstration. Available at <https://iopscience.iop.org/article/10.1088/1741-2552/ac35ac/data>. Video demonstrating one participant performing 7 trials of the center-out task that spanned all possible target categories (T1-5, close/far, and rest targets). All videos and data seen within this video are synced in time. **Left:** top-down video of the participant performing the task; the sensorized orthosis on her right arm is visible, as well as the EMG grid on her biceps underneath it. **Right:** the user interface that the participant saw when performing the center-out task. Motor unit action potential indicators (blue, red, yellow) are visible at the top of the interface, in addition to the indicators for unselected units (“∞”) and unidentified threshold crossings (“-1”). Each of the three selected units have a corresponding auditory pitch that is audible when an action potential is detected. The middle displays the center-out task, where the tip of the black arrow corresponds to the cursor’s position according to a population-coding scheme and where trial targets are highlighted in blue. **Bottom:** real-time EMG and motor unit data, not visible to the participant. A representation of the 56-channel bipolar derivations of the surface EMG is presented in the bottom left, where hues represent the smoothed, total energy in a particular channel in recently detected action potentials. The top row of the 2D grid represents the row of channels most proximal on the biceps, while the left column represents the most lateral, most external (i.e. towards the biceps long head) column of channels. Three bipolar EMG channels are selected for representation in the middle (faded gray, with rows highlighted with the appropriate motor unit colors). Overlaid on the raw EMG voltages in this middle plot are the timings of detected motor unit action potentials for the three selected units for the center-out task, with these units’ normalized firing rates displayed in the bottom right. These normalized firing rates are summed up along their three vectorial axes to yield the black arrow’s position in the center-out task.

Supplementary Video 2 | Speller task demonstration. Available at <https://iopscience.iop.org/article/10.1088/1741-2552/ac35ac/data>. Video demonstrating one participant performing the speller task, correctly typing 9 characters in 1 minute. All videos and data seen within this video are synced in time. **Left:** same as in Supplementary Video 1. **Right:** the user interface that the participant saw when performing the speller task. Motor unit indicators are the same as in the center-out task, described in Supplementary Video 1. The OPTI-II keyboard is visible in the middle of the interface, with the target sentence and pending letters (gray or blinking letters) visible above the keyboard. The black dot is the cursor whose velocity is controlled by the normalized firing rates of the 3 selected motor units. Cursor clicks are performed similar to the center-out's T5 target, through simultaneous co-activation of all three motor units. **Bottom:** same as the bottom pane of Supplementary Video 1.

Chapter 3

A neural model of proximity to reward

Throughout learning, refinement of neural activity in cortex, a process termed “credit assignment”, is thought to underlie the refinement of behavioral actions leading to reward. While the striatum relates behavior to reward, it remains unclear how the striatum might relate cortical activity to reward to facilitate credit assignment. By rewarding specific motor cortical patterns while recording from cortex and striatum, we demonstrate that neural activity in the striatum forms an internal, continuous model of the proximity of behaviorally-relevant cortical activity to reward. Dorsomedial and ventral striatum played complementary yet distinct roles in this model, emerging as cortical activity consolidated over learning. Striatal activity thus constitutes a neural model of cortical progress towards reward, suggesting one mechanism by which the brain implements credit assignment to refine behavior.

3.1 Main

When we first begin learning a new skill, initial behavior is highly variable as we explore which actions lead to reward (R. M. Costa 2011). As behavioral variability decreases over training, neural activity patterns in cortex that relate to those behaviors simultaneously reduce in their variability (Ganguly and Carmena 2009; Lemke et al. 2019). The computational theory underlying this process of neural consolidation during skill learning is frequently termed the “credit assignment problem”: to learn, the brain must discover and subsequently bias task-relevant cortical neurons to execute intended behavior more consistently (Richards et al. 2019), despite reinforcement signals being both spatially non-specific (Hamid et al. 2021) and temporally delayed (Schultz et al., n.d.; Sutton and Barto 2018). While particular subcortical areas have been identified as necessary or sufficient for credit assignment (Koralek et al. 2012; Neely et al. 2018; Athalye et al. 2018), a full elucidation of the neural basis of skill learning remains elusive.

As animals attempt to maximize reward over time throughout learning, a solution to the credit assignment problem inevitably involves the relationship between ongoing behavior, relevant cortical activity, and proximity to reward. To this end, myriad studies suggest the striatum as a strong candidate for a home for this relationship: the striatum is necessary for skill learning (Neely et al. 2018; Atallah et al. 2007; Yin et al. 2009), encodes reward-related quantities (Shin et al. 2021), and is anatomically well-suited to bridge cortex and reward-relevant subcortical areas (Dudman and Krakauer 2016). Existing studies suggest different roles during learning across the dorsal/ventral axis of the striatum (Atallah et al. 2007; O’Doherty 2004) and across cell types (Atallah et al. 2014), predicting heterogeneity in striatal encoding. However, it remains

unclear how striatum might facilitate credit assignment to task-relevant cortical neurons and how such a role might differ between striatal regions.

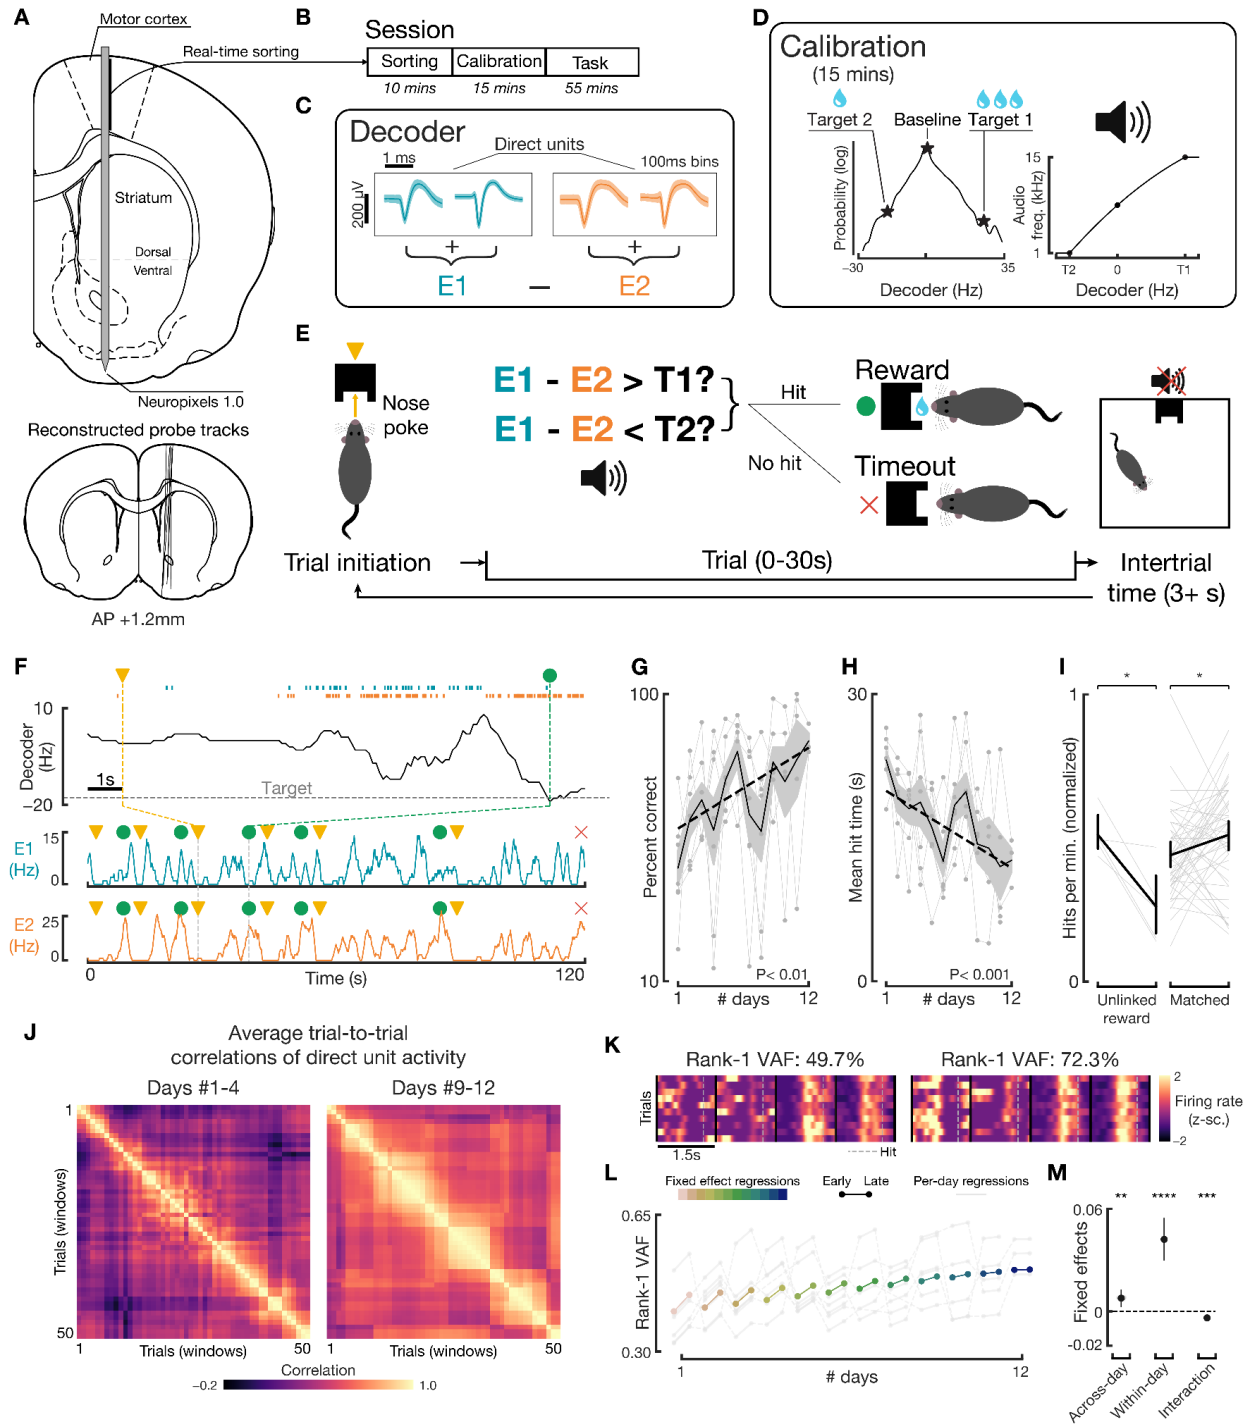


Figure 3.1. Rats learn a neuroprosthetic task driven by motor cortex units via an increasingly consistent neural strategy. (A) Targeted chronic implant trajectory, Neuropixels 1.0 probes in rats (n=7). Bottom: probe tracks via histological imaging confirming desired positioning. (B) Phases of daily training. (C) Example mean waveforms (shading ± 1 SD) of direct units and

decoder calculation. **(D)** Decoder targets and its mapping to auditory tone according to calibration period. **(E)** Task structure. **(F)** Example task data, displaying E1 (middle, blue) and E2 firing rates (bottom, orange) and trial initiation (yellow triangle), hit (green circle), and failure (red X). Top: each dot represents a spike. **(G)** Percentage of correct trials increased ($p=0.003$). Black line & shading is mean ± 1 SEM; dotted line regression line. **(H)** Mean successful trial length decreased ($p=0.0003$). **(I)** Hits per minute (max-normalized per-session) decreased after unlinking of reward from the task (left; $p=0.02$, $n=5$ sessions), whereas performance-matched sessions increased (right; $p=0.04$, $n=48$ sessions). **(J)** Pairwise correlations of trial direct unit activity, averaged across all sessions in the days #1-4 ($n=7$ rats) and days #9-12 ($n=6$ rats). **(K)** Rank-1 variance-accounted-for (VAF) computation. **(L)** Regressions from linear mixed model (fixed effects: colored lines; per-day regressions incorporating random effects: gray lines) displaying changes in Rank-1 VAF both within-day (early vs late trial windows, indicated by endpoints of each line) and across-day (x-axis). **(M)** Fixed effects (slopes) from linear mixed model. Across-day slope (Rank-1 VAF per day) 0.008, $p=0.006$; within-day slope 0.042, $p<0.0001$; interaction term -0.0038, $p=0.0004$. For all, bars indicate 95% confidence interval.

Taking advantage of a neuroprosthetic task paradigm (Koralek et al. 2012; Neely et al. 2018; Athalye et al. 2018) in which observed neural activity causally determines behavior, we studied how the striatum models the relationship between task-relevant cortical activity and reward and how this encoding emerges over learning. We simultaneously recorded tens to hundreds of single units from the motor cortex, dorsomedial striatum (DS), and ventral striatum (VS) of 7 rats over 8-15 days of training ($n=80$ days total, Figure 3.1A-B, Supplementary Figure 3.1) using chronically implanted Neuropixels 1.0 electrodes (Jun et al. 2017). Each day, we divided four well-isolated units located in the motor cortex, henceforth designated as “direct units”, into two ensembles, *E1* and *E2*, and defined a simple decoder that computed the difference between the summed firing rates within each ensemble (“*E1* - *E2*”, Figure 3.1C). The rat’s goal was to modulate cortical activity such that the decoder output exceeded either a positive (T1) or negative (T2) target (Figure 3.1D-F). This “hit” yielded either smaller or larger amounts of sucrose water reward, depending on target (Figure 3.1D-F). All cortical units other than direct units were designated as “indirect units”. During trials, an auditory tone with its frequency derived from the decoder provided external feedback to rats of task progress.

Over learning, rats improved in both the proportion of successful trials ($p=0.003$; Figure 3.1G) and the average length for successful trials ($p=0.0003$; Figure 3.1H). Unlinking reward from the task in a subset of high-performing, late-learning sessions significantly decreased performance, whereas normal sessions with matched initial performance increased in a similar time frame, confirming goal-directed behavior. (Figure 3.1I). Additionally, while differing reward amounts were provided at the two decoder targets, rats did not prefer targets based on reward amount, direction, or calibration bias, but instead preferred whichever target they initially achieved, a strategy correlated with higher performance and with learning (Supplementary Figure 3.2). To this end, further analysis in this study considers only the successful trials for the rat’s preferred target on that day (average of 86.8% successful trials per session achieved the same target), unless otherwise noted.

We then investigated rats’ neural strategies by analyzing activity of the direct units surrounding target achievement. Strong trial-to-trial correlations in direct unit activity emerged by late training (Figure 3.1J), suggesting rats converged to a single consistent strategy each day. Quantifying consistency via principal component analysis (Figure 3.1K), direct units’ activity increased in consistency both across- and within-day, plateauing in late training (Figure 3.1K-M).

Thus, rats progressively refined rewarded cortical activity patterns, likening such cortical activity to that of skilled movements (Lemke et al. 2019; Peters et al. 2014), and echoing findings from previous neuroprosthetic studies (Athalye et al. 2018; 2017; Ganguly and Carmena 2009). Hence, this neuroprosthetic task established a suitable framework to study credit assignment to cortical units.

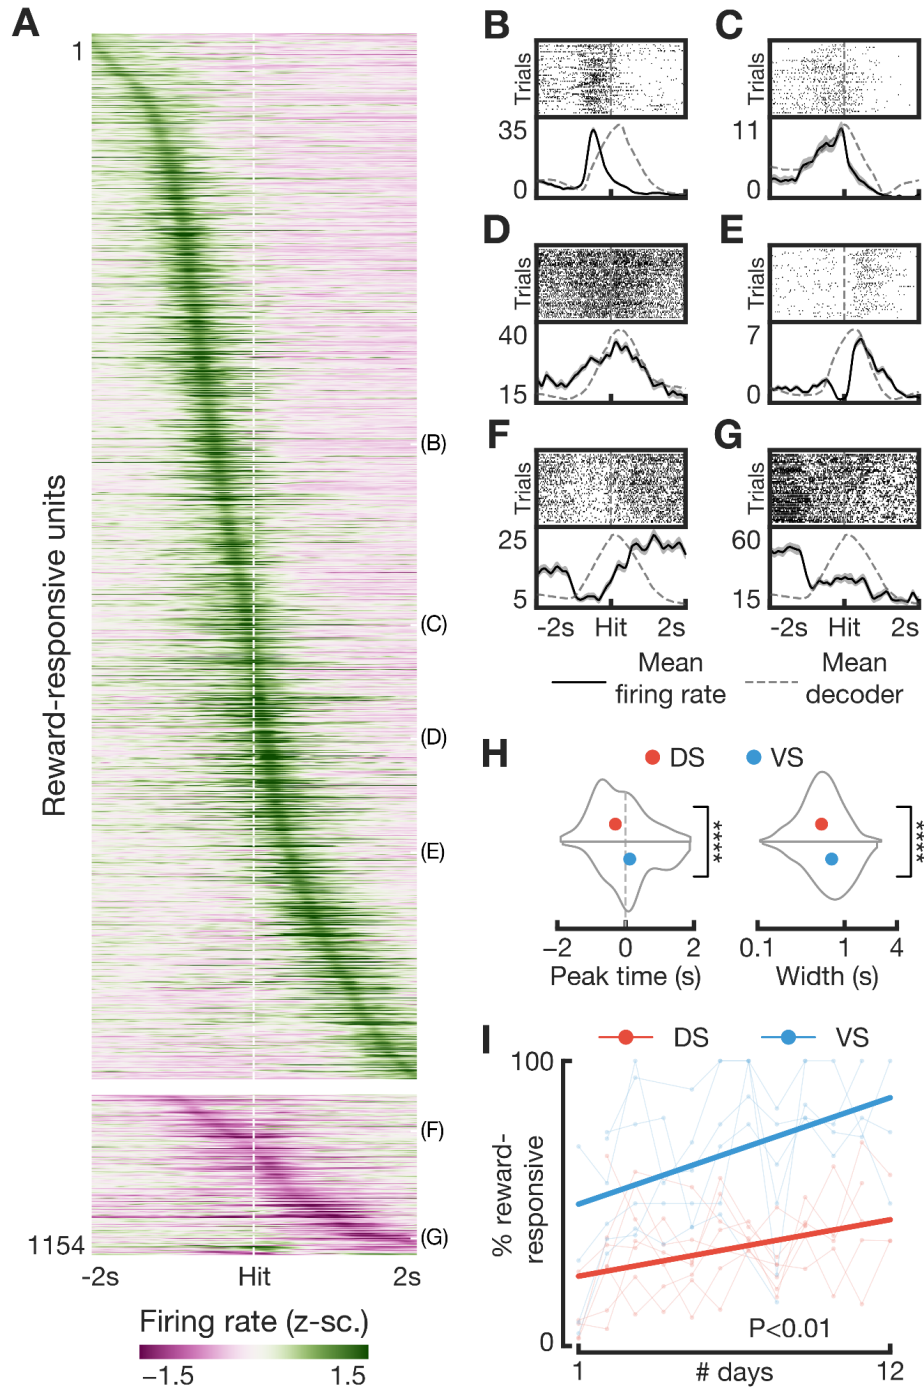


Figure 3.2. Striatal units develop varied and stereotyped behavior surrounding task hit. (A) Mean firing rates (z-scored) of reward-responsive striatal units during the 2 seconds preceding and following a task hit (white dotted line). Top block: up-modulating units; bottom block: down-modulating units. (B)-(G) Example reward-responsive striatal units with narrow (B-C-E) or broad (D-F-G) modulations. Top: raster spike plot, time-locked to hit (gray dotted line in center). Bottom: mean firing rates of the same unit (blue; shading 1 SEM) overlaid on mean decoder value (gray dotted line). (H) Distribution of peaks (left) and widths (right) of the mean firing rates for all significantly modulating dorsal (red) and ventral (blue) striatal units. Peaks ($p < 0.0001$) and widths ($p < 0.0001$) differ significantly. (I) Change in relative proportion of significantly modulating units across days (increase, $p = 0.004$).

We then asked how individual units in striatum responded throughout this task. Striatal units ($n = 4633$) exhibited widely varied, yet stereotyped, response profiles according to task events. Many striatal units were “reward-responsive” ($n = 1193/3943$ in dorsal striatum, $n = 361/690$ in ventral striatum), characterized by a significant positive or negative modulation in the unit’s mean firing rates in the seconds before and after hit (Figure 3.2A-G). These units’ firing rates modulated across a wide temporal range. Such modulation was not visible in the same units during simulated hits within the preceding calibration period (Supplementary Figure 3.4). DS reward-responsive units, on average, modulated both earlier and quicker than those in VS (Figure 3.2H). Additionally, the relative proportion of reward-responsive units increased over days, with roughly twice the proportion in ventral versus dorsal striatum (Figure 3.2I). These results indicate a broad heterogeneity among striatal units suggesting different, emergent encodings during the task.

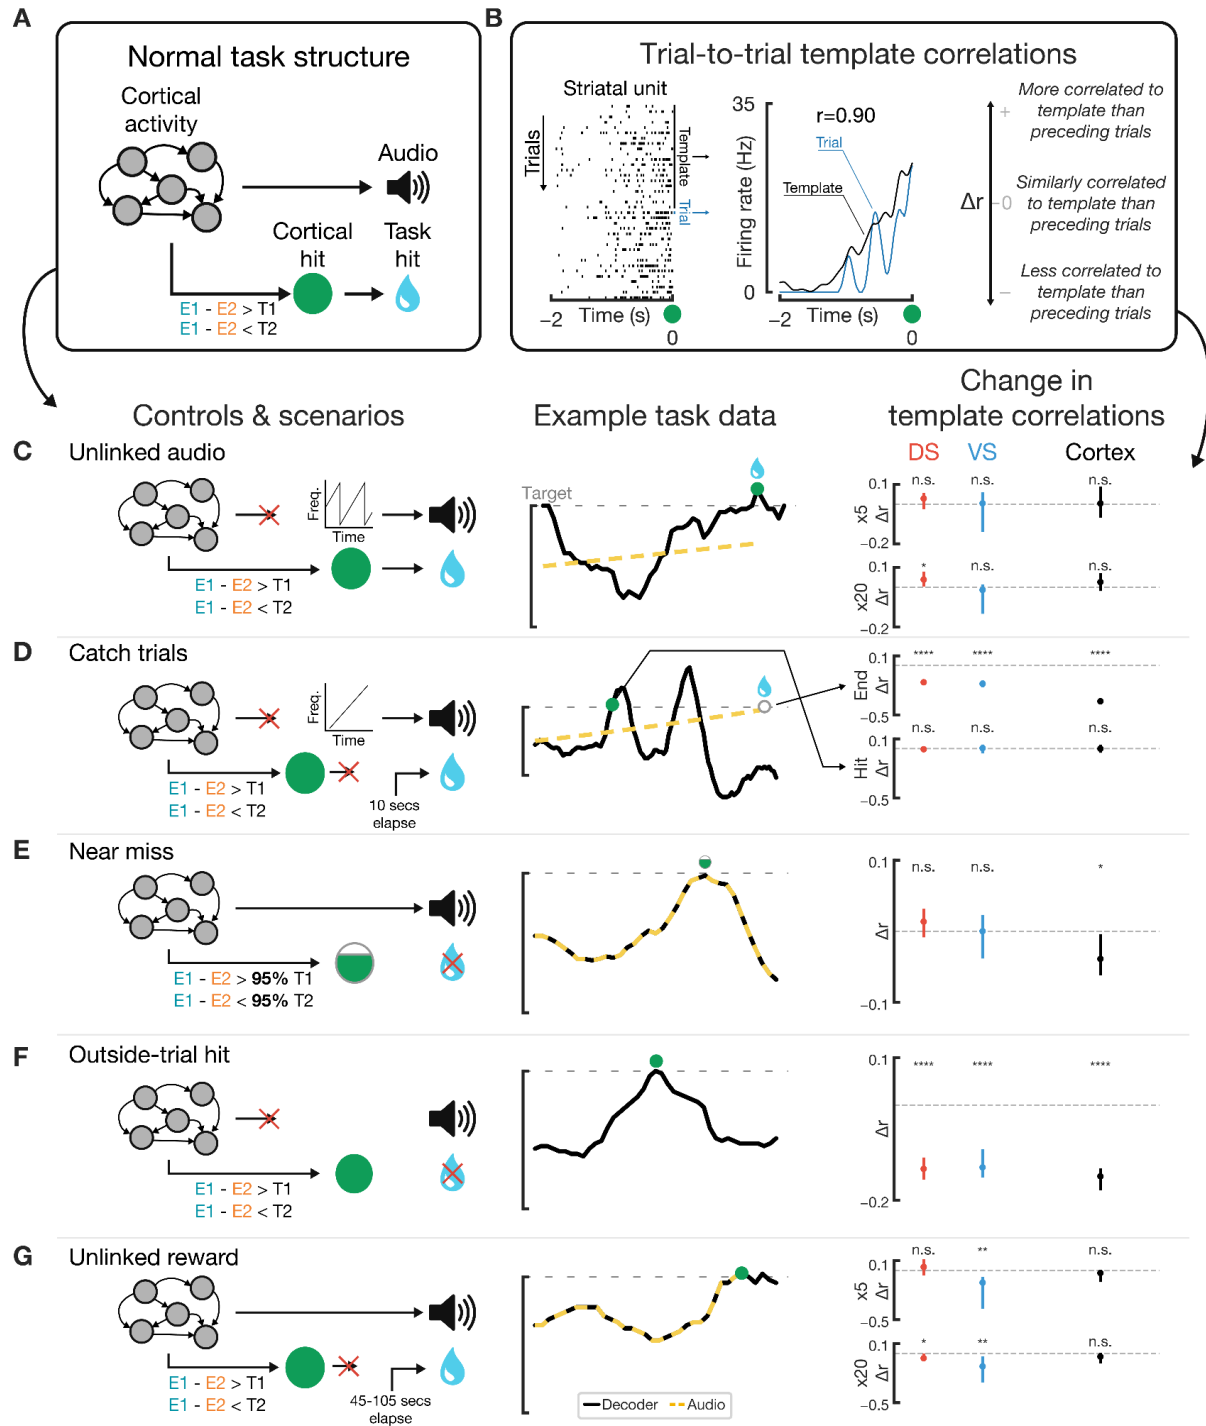


Figure 3.3. Striatal task-locked activity is independent of external feedback and sensitive to reward availability. (A) Normal task structure illustration. (B) Trial-wise template correlation (“ Δr ”) illustration. (C) In particular sessions, audio was unlinked from cortical activity. Right: Δr in the 5 (top) and 20 (bottom) trials after unlinking audio. (D) Catch trials unlinked both audio and task output. Right, top: Δr for ends of catch trials. Right, bottom: Δr for cortical hits performed during catch trials, despite no reward given. (E) Δr for instances at $>95\%$ target but no hit. (F) Δr for hits outside of a trial, despite no reward. (G) In particular sessions, reward was unlinked. Right: Δr in the 5 (top) and 20 (bottom) trials after unlinking reward. For all, n.s. $p > 0.05$; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$.

In order to ascertain what this stereotyped activity within striatum might encode, we probed the relationship of these reward-responsive striatal units to the task's two primary means of external feedback: the auditory tone mapped to the decoder output and reward. To quantify changes in striatal representation during particular scenarios deviating from normal task structure (Figure 3.3A), we compared each unit's firing rates to the scenario's preceding trials' mean firing rates via correlation and subsequent normalization (" Δr ", Figure 3.3B; Methods).

In a subset of late, high-performing days ($n=7$ days), we unlinked the auditory tone from the decoder output in all remaining trials; instead, auditory tone frequencies repeatedly swept a frequency range (Figure 3.3C). However, reward remained dependent on the decoder output reaching either target. Strikingly, after unlinking the audio, we observed no significant changes in performance (Supplementary Figure 3.5). Concomitantly, we observed no changes in striatal activity nor in direct units' activity, with representation even growing more consistent in dorsal striatum after 20 trials (Figure 3.3C). In addition, we analyzed striatal activity during catch trials, in which the rat's behavior was entirely decoupled from the task: the auditory tone linearly ramped across the full frequency spectrum over 10 seconds, always ending in a reward regardless of the rat's neural activity (Figure 3.3D). At the end of these catch trials when reward was given, striatal activity significantly differed relative to the prior successful trials (subplot "end" in Figure 3.3D), even for catch trials occurring during earlier or more poorly-performing periods (Supplementary Figure 3.6). Additionally, despite not resulting in a reward, rats still occasionally performed a cortical hit (i.e., when the decoder output reached a target) during catch trials; striatal activity during these catch-trial hits did not significantly differ (Figure 3.3D). Taken together, these results demonstrate that striatal activity is not dependent on any external measure of task progress within a trial and does not encode reward anticipation or reward retrieval.

To assess striatal activity's relationship to reward, we then probed two additional scenarios in which cortical pattern execution did not lead to reward achievement. "Near-miss" cortical activity, i.e. activity that led to decoder output close to, but not sufficient for, achieving a reward, led to no significant changes in striatal representation (Figure 3.3E). Conversely, when a cortical hit that would typically grant reward was executed outside of a trial, we observed significant differences in striatal representation relative to preceding successful trials (Figure 3.3F), even for those hits where the activity of direct units was more similar to preceding trials (Supplementary Figure 3.7). Finally, to probe the sensitivity of this striatal representation to devaluation, in a subset of late, high-performing days ($n=5$ days), we stopped delivering the water reward as a result of cortical hits (Figure 3.3G). As mentioned above, task performance significantly decreased in the minutes following this unlinking (Figure 3.1I). Within twenty trials following the unlinking of reward, both dorsal and ventral striatal representations significantly changed. These results show that this striatal representation is dependent on the link between cortical activity and reward.

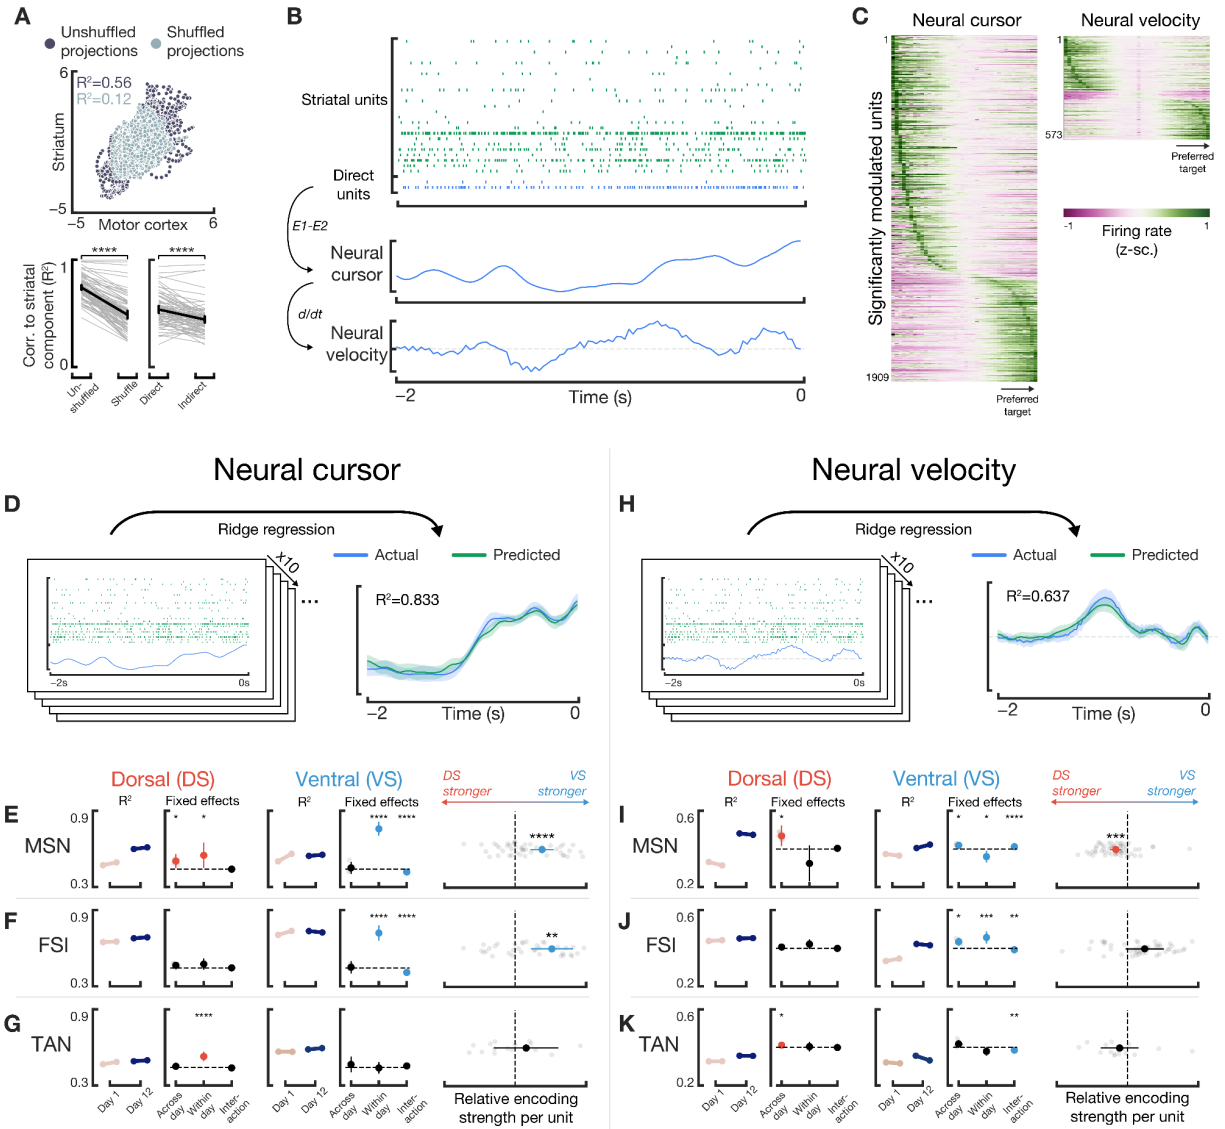


Figure 3.4. Striatum heterogeneously models both the proximity of cortical activity to reward and the changes of that proximity. (A) Top: Cross-area activity between striatum and motor cortex for unshuffled data (light gray) or shuffled data (dark gray). Bottom: coefficient of determination for cross-area activity for each session (light thin lines) or in average (black) when cross-decomposing unshuffled and shuffled motor cortical data (left, $p=1e-32$ different) with striatal data and when cross-decomposing striatal data with direct units' data vs indirect units' data (right; $p=2e-11$ different). (B) Example of the “neural cursor” and its time derivative, “neural velocity” within the 2 seconds preceding cortical hit. (C) Tuning to neural cursor (left) or neural velocity (right). Only units deemed significantly modulating are shown. Units are ordered by max bin. (D) Population model utilized to analyze encoding strength of the neural cursor. Actual (blue) and predicted (green) cursor values (shading SEM). (E-G) For each section: left: regressions from linear mixed model of cross-validated, unit-normalized R^2 for various subgroups of units on day #1 and day #12. Right: fixed effect slopes (in log-transformed R^2 , see Methods) for neural cursor encoding; positive slopes indicate increases in encoding strength. Faded gray dots indicate random effects, when applicable. Third section: difference between DS vs VS encoding strength per unit recorded. Values farther right indicate a stronger relative encoding for VS than DS; dotted line indicates no difference. Putative cell types for each row are: medium spiny neurons (MSN), fast-spiking interneuron (FSI), and tonically active neurons (TAN). (H-K) Same as (D-G) except with neural velocity, instead of neural cursor. For all, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$, otherwise not indicated $p > 0.05$; bars represent 95% confidence intervals.

After interpreting these findings concerning reward and independence from external feedback in the context of a cortically-driven task, we then focused on the reward-related relationship between cortical activity and the striatum. Suggesting a specific focus on the cortical activity driving behavior, we found statistically significant correlations on the shared latent variable between motor cortex and striatum, with direct units' activity correlated more strongly to striatum than indirect units' activity at both single-unit (Supplementary Figure 3.8) and population (Figure 3.4A) levels.

We then considered two particular reward-related quantities derived from direct units' activity as candidates for striatal encoding. First, by utilizing a neuroprosthetic task, we *a priori* define the true relationship between cortical activity and reward via the decoder. Thus, the projection of firing rates of direct units onto this decoder subspace, termed the "neural cursor" here, directly represents the proximity of cortical activity to reward (Figure 3.4B). Secondly, seeking a continuous analog of action-related measures from more typical motor tasks (Jin et al. 2014; Klaus et al. 2017; Sales-Carbonell et al. 2018), we computed the time derivative of the neural cursor, termed "neural velocity" here, to assess encoding of the change in proximity to reward of task-relevant cortical activity (Figure 3.4B). Many individual striatal units displayed consistent, significant tuning to the neural cursor and neural velocity spectrum (Figure 3.4C). Seeking to interrogate population-level effects, we then analyzed encoding of neural cursor and velocity within populations of striatal units, taking the cross-validated coefficient of determination (R^2) from a ridge regression model as a measure of encoding strength (Figure 3.4D).

Considering all striatal units, we found a strong encoding of the neural cursor within the striatum across all days of training, with a significant increase in encoding strength within each day of training ($p=0.0006$; Supplementary Figure 3.9). To investigate specific effects of striatal regions and putative cell-types (Supplementary Figure 3.10) on the encoding of the neural cursor we repeated population decoding analysis across different subsets of striatal units. We observed heterogeneous trends dependent on both cell-type and region (Figure 3.4E-G). In particular, both putative medium spiny neurons (MSNs, Figure 3.4E) and putative fast-spiking interneurons (FSIs, Figure 3.4F) in the VS exhibited significant within-day increases in encoding strength ($p<1e-10$) and significantly stronger encoding than their DS counterparts ($p<0.002$). Regarding neural velocity, we observed moderate encoding strengths (Figure 3.4H; Supplementary Figure 3.9). However, in contrast, we observed strong across-day increases in encoding strength ($p<1e-4$), with no significant changes within-day ($p=0.38$; Supplementary Figure 3.9). Repeating analysis across cell-types and regions (Figure 3.4I-K), DS MSNs exhibited across-day increases in neural velocity encoding ($p=0.025$) and significantly stronger encoding than that of their VS counterparts ($p=0.0002$; Figure 3.4I). Finally, neural cursor and velocity encodings were independent of auditory feedback but dependent on reward (Supplementary Figure 3.11) and were stronger than time-related alternative hypotheses (Supplementary Figure 3.12). All taken together, these data provide evidence that the striatum heterogeneously and simultaneously encodes both the proximity of task-relevant cortical activity to reward and its change over time.

Ventral striatum plays an essential role in enabling stimuli to tune associations between behavior and reward, as VS lesions impair particular forms of learning in instrumental conditioning tasks (Atallah et al. 2007; Nicola 2010; Rothenhoefer et al. 2017; V. D. Costa et al. 2016). However, while prior studies (van der Meer et al. 2010; McGinty et al. 2013; Atallah et al. 2014) congruently show VS firing rates can modulate as animals approach reward, they do not differentiate whether such VS encoding represents stimuli or the rewarded behavior that stimuli might predict. By conditioning a neural pattern, rather than behavior, and subsequently unlinking external stimuli, we provide strong evidence that VS forms an internal representation of the proximity of task-relevant cortical activity to reward, not of stimuli. These results thus shed light on the precise role of VS during learning. Such an internal model may arise from integration across VS's highly-convergent inputs (Yoshida et al. 2020) from hippocampus (Aronov et al. 2017), amygdala (V. D. Costa et al. 2016), and prefrontal cortex (Aoki et al. 2019), each likely containing distinct, and perhaps even competing, information (Floresco 2015).

Accurate internal encoding of proximity to reward engenders accurate reward-prediction errors critical to learning (Schultz 2019). We observe remarkably strong neural cursor encoding in VS, a rapid change in VS representation immediately following task devaluation, and significant within-day increases in neural cursor encoding strength. An accurate, updateable encoding of reward proximity underlies temporal credit assignment in computational models of reinforcement learning (Sutton and Barto 2018), and thus VS may manage temporally delayed reinforcement in the brain. In particular, via strong bidirectional connections with the ventral tegmental area (VTA) (Athalye et al. 2018), this encoding may facilitate appropriate dopamine release at the time of reward across both striatum and cortex, known to be critical for skill learning (Athalye et al. 2018; Hamid et al. 2021).

Shifting focus to dorsal striatum, our stereotyped patterns of dorsomedial striatal units in response to target achievement echo task-locked firing rate modulations observed in prior neuroprosthetic (Neely et al. 2018), navigation (Thorn et al. 2010), and motor (Kupferschmidt et al. 2017) tasks. In conjunction with dorsolateral striatum and dopamine, task-locked dorsomedial activity is hypothesized to guide selection and shaping of actions by coincident activation with cortical activity that leads to the desired action (Kupferschmidt et al. 2017; Klaus et al. 2019). The assembly of shorter, behavioral “syllables” into coarser behavioral actions is mediated by the striatum (Markowitz et al. 2018), implying this striatal influence in goal-directed action selection may be continuous as an action unfolds (Jin and Costa 2010; Sales-Carbonell et al. 2018).

Our results show DS forms an internal model of the cortical activity responsible for changing the animal's proximity to reward. In congruence with this theorized role in action selection and shaping, DS may leverage such a model to continuously provide immediate feedback to the motor cortex via downstream thalamic circuits (Aoki et al. 2019) to quickly positively or negatively bias unfolding cortical activity. Critically, a mapping of the neural velocity would provide more immediate feedback relative to that of the neural cursor, reminiscent of the responsiveness benefits of derivative-based controllers widely utilized in control theory (Nise 2015). Such a mechanism can then underlie structural credit assignment: by dynamically filtering

cortical activity into desired and undesired activity – i.e. activity bringing the animal closer and farther from reward, respectively – dorsomedial striatum can participate in a positive-feedback recurrent loop biasing those cortical neurons leading towards reward (Hunnicutt et al. 2016). These findings of context and reward sensitivity within dorsal striatum seemingly contrast with those of a recent study, which showed invariance to task context in the dorsal-striatum-cortex relationship (Peters et al. 2021). However, these studies’ neural recordings vastly differ in spatial scales, suggesting task context selectively affects the transfer function between particular cortical and dorsal striatal neurons without affecting broader functional connectivity.

Our results elucidate the neural circuits underlying goal-directed learning by demonstrating an internal, continuous model of the proximity of cortical activity to reward within the striatum, with differential representation across dorsal/ventral region and cell types. Thus, we provide evidence of a dynamic and heterogeneous coupling between motor cortex and striatum that may participate in the brain’s solution to the credit assignment problem.

3.2 Materials and Methods

Animals

All rat experiments were performed in compliance with the regulations of the Animal Care and Use Committee at the University of California, Berkeley. Singly housed, male Long-Evans rats on a 12h light/dark cycle weighing 200-300 g were used for the experiments. All rats that had at least 4 well-isolated units in the motor cortex performed the neuroprosthetic task for as many days as possible, until no more motor cortex units remained or until sufficient days of training elapsed. Only rats (n=7) that performed at least 8 days of the task were included in analysis. In 1 rat, 2 days in late training (days #14 and #17) were omitted from analysis due to a different neuroprosthetic task structure attempted. Otherwise, no rats, trials, or sessions were excluded from analysis.

Implant Construction

All rats were chronically implanted with Neuropixels 1.0 probes, mounted within a custom-designed, lightweight 3D-printed enclosure containing both the probe and the Neuropixels headstage PCB. This enclosure added onto a prior study’s design (Luo et al. 2020, 202) and enabled quick and stress-free probe connections by chronically mounting the headstage and utilizing snap-fit joints to connect the headstage PCB. A transport-friendly container utilized for safekeeping and for sterilization was designed to secure the probe mounted within its enclosure. Explantation of probes was possible with this implant. Construction was largely similar to that of Luo et. al (Luo et al. 2020); detailed construction steps can be found alongside the 3D models of this study’s head-mounted connector.

Rat Surgery

Targeted stereotactic coordinates were 1.4mm anterior to bregma and 2mm lateral from midline, though exact implant coordinates varied based on cortical vasculature and removal of

dura mater. Probes were lowered as deep as possible, up until approximately 8.5mm deep from the cortical surface.

Fully assembled probes within their enclosures were sterilized with ethylene oxide gas sterilization. Immediately before surgery, probes were maximally lowered into an excess quantity of lipophilic dye for histological probe tracking (Vybrant DiI Cell-Labeling Solution, ThermoFisher Scientific, Waltham, MA) repeatedly for several seconds.

After approximately 1 week of manual handling, rats of age 8-9 weeks old, 200-300g were anesthetized with isoflurane (0.5-3%), set in a stereotactic frame (Kopf, Tujunga, CA, USA). Artificial tears were applied, and a closed-loop heating pad utilized to maintain 35°C (RightTemp, Kent Scientific Corporation, Torrington, CT). Dexamethasone (0.75 mg/kg) and buprenorphine (0.05 mg/kg) were administered subcutaneously for anti-inflammation and analgesic properties. Saline (10 ml/kg) was administered subcutaneously for hydration. Bupivacaine (1.5 mg/kg, diluted further by 50% with saline) was administered locally to the surgical site to provide secondary local anesthesia. Hair was cleared with an electric shaver and fully removed with brief application of hair-removal product. After verifying surgical levels of anesthesia, rats were secured in blunted ear bars, and the surgical site was sterilized with isopropyl alcohol and chlorhexidine. A midline incision was then made with a scalpel, and skin retracted with Alm retractors (Fine Scientific Tools, Foster City, CA). A spatula and forceps were used to clean the skull of all overlying tissue while keeping the skull moist with sterile saline kept on ice. Additionally, hydrogen peroxide was used to aid removal of remaining tissue and to aid in drying the skull. A #15 blade was used to gently yet firmly scrape the surface of the skull in a cross-hatch pattern to increase roughness of the skull for improved eventual adhesion of dental cement. The skull was fully dried and all bleeding contained.

6-8 1mm diameter holes were drilled around the perimeter of the skull utilizing a custom-made impedance-sensing automated drill system which stopped precisely at the detection of dura mater and/or CSF, based loosely on the system of another study (Pak et al. 2015). M1x2mm screws (McMaster Carr, Robbinsville, NJ) were then firmly screwed into each hole until dura was just touched. A rectangular craniotomy of 2mm x 2.5mm centered around the target implant site was then drilled using the same automated drill system, which enabled manual specifications of craniotomy depths and shapes. The perimeter of the craniotomy was repeatedly drilled until the piece of skull could be excised gently with minimal pressure. During excision of skull, constant irrigation with cold saline minimized sticking of dura to excised skull pieces. Once the dura mater was exposed, a 30G needle with its tip bent at 45 degrees was utilized in conjunction with angled forceps (Dumont #5/45 forceps, Fine Scientific Tools) to very gently cut and remove patches of dura mater in areas clear of large vasculature. Bleeding frequently occurred despite utmost caution; in these cases, gelfoam soaked with cold saline was gently moved over bleeding sites until bleeding stopped. Only the minimal dura was removed to enable implantation of the probe clear of surface vasculature. A piece of saline-soaked gelfoam was then placed in the craniotomy space, and saline periodically applied as necessary to ensure the exposed cortex was hydrated.

The skull was then redried completely. The probe enclosure was then mounted stereotactically (Model 1766-AP Cannula Holder, Kopf) and straightened appropriately. The enclosure was moved near to its desired implant site, and the ground wire was connected by repeated wrapping around 2-3 skull screws. Ground wire/screw connectivity was confirmed with impedance tests between the screw and the rat's paw. After re-ensuring the skull was completely dry, C&B Metabond (Parkell, Edgewood, NY) was used to completely cover the exposed skull, including all skull screws and the ground wire. After the cement dried, gelfoam within the craniotomy was removed, and the probe tip was lowered to the cortical surface using a 25x zoom surgical microscope. Utilizing custom-built software controlling a vertically-mounted linear translation stage (ThorLabs NRT100/M, Newton, NJ) mounted on top of passive vibration isolation legs to minimize vibration (ThorLabs PWA075), the probe was then carefully and slowly lowered to minimize damage to brain tissue and minimize eventual glial encapsulation of the probe. The first 1mm of insertion was lowered at 1-2 microns per second, the next 5mm of insertion was lowered at 4 microns per second, and the remaining insertion was performed at 1-2 microns per second; saline was periodically applied to keep cortex hydrated, and the surgical table was minimally touched to reduce induced vibrations. When the probe was within 1mm of its final depth, a small quantity (~1-2 microliters) of soft silicone elastomer (DOWSIL 3-4680, Dow) was carefully injected with a micropipette to seal the craniotomy while minimizing vibrations to the surgical table; too much elastomer could lead to overflow onto the skull and prevent proper sealing with cement in subsequent steps. The probe was then lowered to its final, maximal depth (≤ 8.5 mm or just before the implant touched the skull). We then applied a viscous dental composite (Absolute Dentin, Parkell) to seal the craniotomy and fix the implant in place to the skull. Once the viscous cement completely dried (~4 minutes), less viscous dental cement (Ortho-Jet, Lang Dental, Wheeling, IL) was applied liberally around the implant and skull; the loose skin surrounding the implant was used to gently mold dental cement in order to ensure smoothness and improve surgical time. Rats were given meloxicam (2 mg/kg), another supplement of dexamethasone (0.2 mg/kg), and saline (10 mL/kg) subcutaneously, and antibiotics were liberally applied. Rats received a taper schedule of dexamethasone (1.0 mg/kg 2 days, 0.5 mg/kg 1 day post-surgery), and were allowed five days post-surgery to recover before experiments began. Antibiotics were applied daily; rats that did not recover weight by five days post-surgery were allowed extra time to recover as needed.

Electrophysiology

Unit activity from Neuropixels 1.0 probes (Jun et al. 2017) was recorded and visualized via OpenEphys GUI (Siegle et al. 2017) (Open Ephys), and all neural data was streamed to disk. Custom software written in Python and C++, spanning both standalone services and OpenEphys plugins, enabled real-time readout of spikes of selected units. Services were networked with a high-throughput, low-latency streaming framework with C++ and Python bindings, River (Botros 2021).

Each day, 2 2-minute baseline sessions were first recorded in order to determine the enabled channels for the Neuropixels 1.0 probe; one session captured the deepest 384 channels, and another session captured the next-dorsal block of 384 channels. We then utilized Kilosort2 (www.github.com/MouseLand/Kilosort2) to sort each session individually, from which the 384 channels capturing the most units across the probe length were selected via an optimization algorithm (Choi et al. 2020). A third 2-minute baseline session was then performed with this final channel configuration and was sorted via Kilosort2.

Utilizing a custom-built GUI, 4 well-isolated units within the motor cortex were selected from these Kilosort2 results as “direct units”, with preference given to selecting units with similar waveforms, depths, and/or firing properties as previously-selected direct units for that rat. In cases where similar units were not obviously available, direct units were chosen from similar depths within the motor cortex as preceding days. All rats exhibited a gap in well-isolated units beginning between 2.0mm and 2.5mm from the cortical surface, putatively corresponding to the dorsal edge of white matter tracts of the corpus callosum; direct units were chosen strictly dorsal to the start of this gap. Post-mortem histology (described below) confirmed the location of direct units within the motor cortex.

We then desired to adopt Kilosort2-identified spikes for our 4 selected direct units into a methodology commonly used for online spike sorting: classification via ellipsoids within a PCA-determined space of waveforms of threshold crossings. To do this, for each selected direct unit, we identified an appropriate threshold for each relevant channel based on that channel’s RMS noise in order to perform threshold-crossing extraction. From these threshold crossings, we then identified and subsequently projected these voltage waveforms into a 3-dimensional space via PCA. Particular threshold crossings corresponding to Kilosort2-identified spikes then formed the putative “ground truth” for online sorting, and a best-estimate ellipsoid in PC-space was constructed via stochastic gradient descent to maximize accuracy (Rossant et al. 2016). The parameters for the PCA projection and subsequent ellipsoid classification were then transferred to OpenEphys for online spike identification. For subsequent sessions that required real-time online sorting, an OpenEphys program was used that processed data online in similar fashion to Kilosort2: common median referencing across all channels, followed by a 150 Hz, 3rd order highpass filter. Online-identified spikes for each of the 4 direct units were then streamed out from OpenEphys in real-time via River. Firing rates were then continuously computed via these streamed spikes in 100ms bins and smoothed using a running average of 10 bins (1 second). Frequencies for the auditory tone were simultaneously computed as well, according to the thresholds set during the calibration period, described below.

In a subset of rats (n=4), improvements were enacted to increase the agreement between online-identified spikes and Kilosort2-identified spikes for the (putatively) same unit within a session, in cases where neural recordings were not fully stationary within a day. Though not common, some direct units tended to appear to slowly drift along the dorsal/ventral axis, i.e. along the length of the probe. To combat this, instead of a single-channel’s 3 dimensions in PC space, each unit was identified using up to a 15-dimension ellipsoid, with 3 PC dimensions originating

from each of 5 channels stacked along the dorsal/ventral axis. Then, every 6 minutes with a 3-minute overlapping window, PC scores of all threshold crossings on a unit's channels were fit with a Gaussian mixture model (`sklearn.mixture.GaussianMixture`) with a number of components chosen to minimize the Akaike Information Criterion (AIC). The cluster containing the highest proportion of online-identified spikes became the new "ground truth" for that unit's online sorting, and a new best-estimate ellipsoid was fit to this cluster and formed the new basis for identifying spikes for that given unit online. This updated clustering was only accepted if it had at least 70% agreement with the prior sorting. This methodology thus enabled tracking of slow changes in neural recordings across channels.

Behavioral Task

Upon recovery from surgery, rats were deprived of water for 24 hours before initiation of experiments. During training, rats only received access to water during the behavioral task, unless supplemental water was needed to maintain 90% body weight. After initial sorting and selection of direct units were performed as described above, rats had a 15-minute calibration period, in which they freely moved around the cage and were passively given water rewards paired with a reward tone every 45-105 seconds (uniformly random). After 15 minutes, the calibration period data was used to determine assignments of the 4 direct units into ensemble #1 (E1) and #2 (E2), the positive (T1) and negative (T2) neuroprosthetic targets, a baseline value, and the mapping of decoder output to the frequency of the auditory tone. The goal of calibration was to find the set of parameters leading to a success rate during calibration of 35-40% of trials.

In order to set these parameters, we simulated the task performance of the rat during this 15-minute calibration period for all 6 combinations of E1 and E2 and across a sweep of thresholds for T1 and T2. In particular, T1 and T2 values were selected based on exhaustive search that minimized the difference between simulated and goal rates of success (35-40%). E1 and E2 assignments were then selected as the assignments that yielded a success rate close to the goal, had subjectively fair thresholds, and were balanced between T1 and T2 rewards. The baseline was set as the mean E1 - E2 value during calibration. The auditory tone's frequency was set according to a 2nd-order polynomial fit interpolating decoder values of T2, 0, and T1 to frequencies 1, 8, and 15 kHz, respectively. These frequencies fall within the auditory spectrum of the rat. The task structure used during simulations was identical to the true task structure, and was as follows (and as illustrated at a high-level in Figure 3.1).

First, firing rates were first computed using spike counts in 100ms bins and smoothed with a 1s rolling average filter. Then, for each time bin, decoder output was computed as

$$\begin{aligned} \text{decoder} = & (\text{E1 \#1 firing rate}) + (\text{E1 \#2 firing rate}) \\ & - (\text{E2 \#1 firing rate}) - (\text{E2 \#2 firing rate}) - \text{baseline} \end{aligned}$$

Note that, elsewhere in this study, we might refer to the decoder as *E1 - E2* as shorthand for the above full formula for the decoder. Trials were initiated once the decoder output crossed zero at

least once between trials and once at least 1 nose-poke was detected between trials. During a trial, if the decoder output exceeded either T1 or T2 for 1 bin, the trial was declared successful; an extended tone with frequency matching that target's frequency was played, and reward was delivered simultaneously. Reward was 30% sucrose water. Each rat was randomly assigned a target to be more highly rewarded, where rewards for this target yielded three times as much water as the other target; this assignment was held constant for the duration of experiments for that rat. If 30 seconds elapsed after trial initiation without decoder output exceeding either T1 or T2, the trial was declared failed, and a white noise sound played. After either reward delivery after a success or a white noise tone after a failure, an inter-trial period of at least 3 seconds was enforced, after which the animal was free to initiate another trial. The auditory tone played continuously during a trial and was muted when no trial was ongoing. Every 20 trials (regardless of success or failure in those 20 trials), catch trials were performed, in which the auditory tone's frequency either ramped from 1 to 15 kHz or 15 to 1 kHz over 10 seconds regardless of the rat's direct unit activity, after which the animal would hear a reward tone and receive a reward (Figure 3.3D). Each day's session would end after approximately 55 minutes of task time or if the animal reached satiety and stopped initiating trials.

Additionally, in a subset of high-performing ($\geq 75\%$ trial success rate in a trailing 10 minute window and ≤ 6 mL water consumed), late-learning (day ≥ 8) sessions, we unlinked either the auditory tone (Figure 3.3C) or the reward (Figure 3.3G) from cortical activity for the remainder of that day's trials. In the case of the former, linear sweeps of frequencies from the lowest (1 kHz) to the highest (15 kHz) ends of the spectrum were performed over a period of 10 seconds and repeated; otherwise, the task structure remained identical. In the case of unlinking reward, we instead delivered a reward every 45-105 seconds, approximately matching the rat's calibration success rate ($\sim 35\text{-}40\%$); otherwise, the task structure remained identical.

Explanation & Histology

After the completion of experiments, rats were first anesthetized with isoflurane in preparation for explantation. After surgical anesthetic levels were confirmed, rats were fixed in a stereotactic frame. Screws securing pieces of the 3D-printed implant were loosened, and the headstage PCB was removed and disconnected. After ensuring vertical alignment to minimize damage to the probe, the same motorized system used to implant the probe was used to gently remove the probe from the implant and brain. Upon successful explantation, the probe was immediately soaked in 10% Tergazyme solution for 24 hours. After this, the probe was soaked in deionized water for 5 minutes, and then soaked in isopropyl alcohol for 1 minute. In cases where residual tissue or elastomer remained on the probe, the probe was cleaned via an ultrasonic cleaner (CREWORKS) filled with isopropyl alcohol, as residual matter would interfere with subsequent recordings. For stuck-on elastomer, the probes were soaked in elastomer solvent (DS-2025, Dow) for 18-24 hours, followed by a 10-minute soak in DI water and then another minute in an ultrasonic cleaner with isopropyl alcohol.

Immediately after explantation of the probe, rats were injected with sodium pentobarbital and transcardially perfused with PBS followed by 4% paraformaldehyde. Brains were removed and post-fixed in 4% paraformaldehyde overnight at 4°C, after which they were stored in PBS at 4°C. When ready, brains were then transferred to 15% sucrose/PBS solution until they sunk, and then a 30% sucrose/PBS solution until they sunk. Brains were then mounted and sliced via a cryostat into coronal slices. Brain slices were then serially imaged with DAPI and DsRed. Images were preprocessed to adjust for orientation and contrast via manual scripts. Images were then aligned in three dimensions to the Waxholm Space Rat Atlas (Papp et al. 2014) via the QuickNII tool (Puchades et al. 2019). Probe tracks within each slice were programmatically segmented via ILastik (Berg et al. 2019), 3D coordinates of probe tracks spanning each brain's slices were extracted via Nutil (Groeneboom et al. 2020), and linear probe tracks were finally reconstructed from these 3D point clouds via custom software. 1 rat's brain was unable to be imaged successfully due to technical issues during the slicing process.

Analysis & Statistics

Analyses were performed in Python 3.8 (<https://www.python.org/>) with custom-written scripts utilizing publicly available software packages, including numpy (Harris et al. 2020), scipy (Virtanen et al. 2020), pandas (Reback et al. 2021), and scikit-learn (Pedregosa et al. 2018). Data pipelines were constructed utilizing Apache Airflow (<https://airflow.apache.org/>), and the majority of data files were stored in the Parquet file format (<https://parquet.apache.org/>), a cross-platform, high-performance columnar data format.

Unless explicitly noted otherwise, in all figures, n.s. $p > 0.05$; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$.

Additionally, where indicated for use below, linear mixed models (LMMs) were implemented via the pymer4 package (Jolly 2018, 4). Each model was inspected for homoscedasticity in residuals via the Het-White test or visually. Normality of residuals was not strictly enforced, as LMMs are relatively robust to non-normality in residuals (Schielzeth et al. 2020); instead, kurtosis and skew of residuals of all models were both confined to be less than 3.0 and are stated as needed per model below. Confidence intervals were computed via Wald estimates.

Behavioral Analysis

Behavioral analysis across sessions focused on percent of successfully completed trials (“% correct”, Figure 3.1) and the average length of successful trials (“Mean hit time”, Figure 3.1), considering trials of all targets. Analysis was done with linear regression (Figure 3.1). To study within-session behavioral changes after the unlinking of reward (Figure 3.1I) and unlinking of audio (Supplementary Figure 3.7), we utilized the number of hits per minute, a measure that incorporated both the length of trials and the time the rat took to initiate trials. We took the mean hits/min in the 10 minutes preceding and the 10 minutes following the start of the first unlinked trial; statistical significance was tested via a paired t-test.

Within each day, the rat preferred one target more than the other (Supplementary Figure 3.2). The preferred target for that day was considered to be the target that had more successful trials for that day. Unless explicitly noted otherwise, all analysis described below utilizes successful trials to the rat's preferred target for that day.

Data Analysis: Preprocessing into Firing Rates

Each day's session was first sorted via Kilosort2. All unit waveforms were manually curated via Phy; units subjectively determined to be noise were excluded. Of these non-noise units, only high-quality units were then utilized for the remainder of this study's analysis, as defined by: having a mean firing rate of greater than 0.2 Hz both globally and within the 4 seconds centered on the end of each successful trial, having less than 50% ISI violations, less than 50% missing spikes due to low SNR, and a trough in the unit's mean waveform occurring temporally before its peak to exclude axonal spikes (Deligkaris et al. 2016). Additionally, in 1 rat, during surgery it was noted the probe was angled relative to the brain during insertion, and histological analysis confirmed the ventral portion of the probe did not intersect with the ventral striatum; all units in this ventral portion of this particular probe were excluded from analysis. Finally, the units and spikes identified online as direct unit activity were merged with these Kilosort2-determined units and spikes. For the remainder of the analysis, all "direct unit activity" corresponds to the units/spikes determined online.

Spike counts for each unit were binned at 20ms resolution and divided by bin size to yield (unsmoothed) firing rates. These firing rates were then extracted time-locked to the end of each trial. Owing to the broader binning/smoothing used online (100ms bins with 1s rolling average) compared to that of the offline analysis done here – as well as to non-zero jitter in the experimental setup itself – neural data varied in its degree of alignment with the online-determined "trial end" across trials. In order to correct for this, time shifting on each trial was performed to optimize the trial-to-trial correlations of direct unit activity with the successful trials of each target (Williams et al. 2020) (Supplementary Figure 3.3). The computed optimal shift for the direct units was then applied to the firing rates of all units to globally align all units to the time of cortical hit. Notably, we only shifted units in time but did not warp time, and the same shift was applied to all units for each trial. Additionally, the shifts per session were adjusted such that the correlations of the unaligned and aligned decoder outputs were maximal, thereby minimizing deviations between the two while increasing overall trial-to-trial alignment. Finally, these aligned firing rates were smoothed with a centered Gaussian kernel with standard deviation 60ms to yield firing rates for each unit. For all subsequent neural data analysis, unless otherwise noted, we utilized these aligned, Gaussian-smoothed firing rates.

Units were then segmented into approximate brain regions according to probe insertion depth. As stated above, all rats exhibited a gap in units beginning between 2.0mm and 2.5mm relative to the cortical surface. Thus, a given unit was considered to be in the dorsal striatum if the channel containing the largest amplitude in the unit's mean waveform was between 2.5mm and

5.5mm below the approximated cortical surface; all units more ventral than 5.5mm were considered ventral striatal units, and those dorsal to the gap were considered motor cortex units.

Furthermore, the direct units and spikes identified online were duplicated in the motor cortex units and spikes identified by offline Kilosort2 sorting. To identify these offline-sorted units corresponding to the online-sorted direct units, we computed the fraction of spikes between each offline and each online-sorted direct unit that occurred within 1ms of one another. Those offline units that had at least 50% agreement with a direct unit and were less than 120 μ m away from that direct unit were identified as a duplicated direct unit; these relatively relaxed criteria aimed to establish an upper bound on duplicated direct units. These offline-identified, duplicated direct units were then excluded from Granger causality and cross-correlation analysis of “indirect” motor cortical units in Supplementary Figure 3.8 and Figure 3.4, respectively, described below.

Data Analysis: Direct units

To analyze direct unit activity (Figure 3.1J-M), firing rates of each direct unit were first z-scored, based on the mean and standard deviations computed from all successful trials towards the preferred target between 1.2 seconds before the hit and 0.3 seconds after (i.e. 75 time bins each trial). Z-scored direct unit activity was then concatenated within-trial across the four units, yielding 300 time bins per trial. Then, we constructed overlapping, sliding windows of 10 trials each, and used 1-component PCA across trials to yield a 300-element projection representing the normalized “average” of the direct unit activity within that window. In Figure 3.1J, the projections between sliding windows were compared pairwise with Pearson correlation. In Figures 1K-M, we took the variance-accounted-for by this 1-component PCA (Rank-1 VAF) as a measure of consistency of direct unit activity within each window. Note that for this measure and window size of 10 trials, a standard normal Gaussian noise process would expect a Rank-1 VAF of 0.133 (simulations not shown). Sessions with fewer than 10 successful trials were ignored. Additionally, since across-day changes were a focus, only sessions within the first 12 days were considered to ensure at least 3 rats’ sessions were included for each day.

We then utilized a linear mixed model to analyze the changes in Rank-1 VAF over within- and across-day timescales. First, we designated “early trials” as the first 15 windows of 10 trials (i.e. spanning the first 25 trials), and “late trials” as all other windows of trials. Then, we constructed a model with: fixed effects of number of days, indicator of early or late trials (-1 or 1), and an interaction term, a random intercept per session, and a random across-day slope per rat. The model converged and met appropriate statistical assumptions (heteroscedasticity test $p=0.28$, kurtosis/skew of residuals 0.74/0.40).

Data Analysis: Reward-responsive modulation

To analyze striatal unit activity, each unit in the striatum was first z-scored, based on the mean and standard deviations of the time periods between 3 and 6 seconds both before the hit and after the hit. Note these windows were intentionally selected outside of the rewarded period to be analyzed in order to capture each unit’s relative baseline. Mean z-scored firing rates were then

computed for the period between 2 seconds before and 2 seconds after the hit (i.e. 200 time bins total), and modulation depth for each unit was computed as the difference between maximum and minimum z-scores.

Reward-responsive striatal units (Figure 3.2) were then determined using a sliding z-score threshold on modulation depth in order to account for varying numbers of successful trials each day. In particular, we computed the z-score threshold corresponding to a 95% confidence that the modulation depth in the unit's mean firing rates was not generated by a stationary, independent Gaussian noise process with 200 samples per trial. This threshold was the minimum z-score threshold T such that, for a given number of trials N , the following probability was below 5%:

$$X \sim \mathcal{N}\left(0, \sqrt{\frac{2}{N}}\right)$$

$$P(\text{any MD exceeds } T) = 1 - (P(X < T) \cdot P(X > -T))^{200C_2}$$

We then ensured all z-score thresholds on modulation depth were non-trivial by setting a minimum threshold of 0.75. This process yielded z-score thresholds such as 1.5 for sessions with 20 trials or 0.75 for sessions with 79 trials or more.

After filtering units according to this z-score threshold, we then identified units as positively modulating or negatively modulating depending on whether the magnitude of the max z-score or min z-score were greater. The time at which this peak occurred was denoted as the peak time (Figure 3.2H). Then, to further reduce the effects of noise or trial size on determination of reward responsiveness, only those units also exhibiting a mean firing rate at the peak time significantly different than zero were finally deemed reward-responsive (Wilcoxon signed-rank test, $p\text{-value} < 0.05$). The width of the modulation (Figure 3.2I) was determined according to 50% of the peak's prominence in the mean firing rate (`scipy.signal.find_peaks` with `rel_height=0.5`). Statistical tests on the distributions of peak time and modulation width used medians (testing different than zero) or the nonparametric Mann–Whitney U test (comparing distributions). Changes in relative proportions of DS and VS reward-responsive units were modeled with a linear regression simultaneously including both DS-specific and VS-specific slopes and intercepts.

Data Analysis: Manipulations

We then sought to analyze changes in activity in reward-responsive striatal units during particular deviations from the normal task structure (Figure 3.3). In particular, we examined: the ends of each trial within 5 or 20 successful trials after unlinking the auditory tone (Figure 3.3C) and reward (Figure 3.3G) from cortical activity; the first cortical hit occurring during each catch trial (if any), as well as the ends of all catch trials (Figure 3.3D); instances where the local maximum of the decoder output reached 95% of either target T1 or T2 (Figure 3.3E); and instances where a cortical hit occurred at least 10 seconds outside of any trial (Figure 3.3F).

For each of these scenarios, we extracted the 2 seconds leading up to a particular scenario (e.g. the 2 seconds preceding a cortical hit). For near misses, outside-trial cortical hits, and catch trial cortical hits, firing rates were first re-aligned according to direct unit activity to include these

new instances of cortical hits, in addition to including all successful trials to the preferred target as done initially (Supplementary Figure 3.3). Then, for each scenario, the mean firing rate for each unit for all preceding successful trials was computed and labeled the “template” for that particular scenario and unit. The firing rate for that scenario was then compared to the template via Pearson correlation. This process was then repeated for the 10 successful trials preceding this scenario, effectively computing a “baseline” correlation of successful trials to their templates relative to this unit and scenario. Δr was then computed for each unit and scenario as the difference between the scenario’s Pearson correlation coefficient and the mean of these 10 “baseline” correlation coefficients. Comparing Δr values across units, instead of raw Pearson correlation coefficients, thus accounted for local differences in stability of the running mean firing rate for each unit. Finally, to analyze these trial-to-trial template correlations, the median Δr value for each unit was taken across scenarios, and then BCa-corrected confidence intervals and p-values were computed using bootstrapping to test if median Δr values significantly differed from zero.

Using this methodology, significant positive values in Δr indicated scenario firing rates closer to the template than preceding trials, i.e. indicating a strengthening representation; Δr values not significantly different from zero indicated a representation not significantly different from preceding trials; Δr values significantly negative indicate scenario firing rates that are farther from the template than preceding trials, indicating a changing representation (Figure 3.3B).

Trial-to-trial template correlations were computed similarly for cortical activity as for striatal activity, except: the time period extracted was between 1.2 seconds before the scenario and 0.3 seconds after the scenario, firing rates for each direct unit were first z-scored and then concatenated within-trial across the direct units before the computation of correlation coefficients, and statistical tests were computed directly on the Δr values without taking the median across scenarios.

Data Analysis: CCA

To assess the relationship between cortical and striatal populations, we utilized canonical correlation analysis (CCA, `sklearn.cross_decomposition.CCA`) to compute the projections that were maximally correlated between the two populations. Similar to principal component analysis (PCA), CCA also aims to reduce the dimensionality of data by finding projections that explain the originating dataset well; however, unlike PCA, CCA finds two components that simultaneously reduce dimensionality across two populations and is thus well-suited to describe cross-area dynamics in neural data (Veuthey et al. 2020).

To this end, CCA was performed on the firing rates within the trailing 2 seconds of all successful, preferred-target trials, concatenated together. For all below analysis, 10-fold, shuffled cross-validation (`sklearn.model_selection.KFold`) was used to compute the cross-validated squared correlation coefficient (R^2) as the mean test R^2 over the 10 folds.

First, to determine whether there was significant motor cortex / striatum cross-area activity, for each session, we shuffled cortical data in time 100 times, and took the mean cross-validated R^2

as a baseline. Comparisons were made between this and the unshuffled data via a paired t-test (Figure 3.4A, left).

Secondly, we sought to compare the specificity of striatal encoding by comparing the population components of striatum and direct units versus the population components of striatum and indirect units (Clancy et al. 2014; Ganguly et al. 2011). CCA was performed in two ways: first, cross-decomposing the activity of all striatal units and the 4 direct units for each session, and then cross-decomposing the activity of all striatal units and indirect cortical units. To correct for cortical neuron sample size in the latter, 4 random indirect units were chosen for analysis and the cross-decomposition performed, and this process was repeated 100 times and results averaged. Again, comparisons between striatum/direct and striatum/indirect cross-area correlations were performed via a paired t-test (Figure 3.4A, right).

Data Analysis: Granger

Similar to the CCA analysis, we desired to investigate the level of functional connectivity between motor cortex and striatum, but now at the level of individual units. To this end, we computed the directed functional connectivity between all pairs of motor cortex and striatal units via Granger causality analysis (up to 100ms lags allowed, i.e. 5 bins). To establish whether a given value was significant, each unit's data was shuffled in time 100 times, and a given unit-pair was considered connected if its (unshuffled) F-test value exceeded the 95th percentile of the shuffled F-test values. Un-smoothed neuronal activity was used for this analysis. The same window of firing rate data was used as in the above CCA analysis (i.e. trailing 2 seconds of all successful, preferred-target trials).

Data Analysis: Neural cursor and velocity computation

Seeking concrete, task-relevant, cortically-derived quantities that could be encoded in the striatal population, we next defined the neural cursor and neural velocity. The neural cursor was defined in the same manner as the decoder utilized for the task: the sum of the firing rates of E1 minus the sum of the firing rates of E2 minus the baseline value computed during calibration. Notably, this differed from the decoder output only in the smoothing parameters: while the decoder was computed with 100ms bins and smoothed causally over 1 second, the neural cursor was computed using the binning (20ms bins) and smoothing (centered Gaussian filter with 60ms standard deviation) used for all analysis. Additionally, since it was computed offline, the neural cursor exactly aligned with binning utilized for all striatal activity. This neural cursor value was finally normalized to values between -1 and 1, where 1 represented the neural cursor value sufficient to hit the preferred target and -1 the other target. The neural velocity was then computed as the time derivative of the (unsmoothed) neural cursor followed by a more broad smoothing (centered Gaussian with 120ms standard deviation) than the cursor, since time differentiation significantly amplifies noise.

Data Analysis: Single-unit tuning curves

To construct mean firing rates of individual units at particular points in the neural cursor and velocity spectrum (Figure 3.4C), we considered all task trial data, across the full length of all trials (though excluding any catch trials and those where reward or audio was unlinked). We first divided the neural cursor spectrum into 40 bins between -1 and 1, and the neural velocity spectrum into 40 bins between -3 and 3. Then, each unit's firing rates were z-scored and the mean z-scored firing rate for each bin was computed. For determination of significantly-modulated units (Figure 3.4C), similar to the minimum threshold utilized in Figure 3.3, we utilized a threshold of 0.75 for the difference between the maximum and minimum mean z-scored firing rates across bins.

Data Analysis: Cell type determination

For use in the population decoding analysis described below, putative cell types were determined via similar means as previously published studies (Peters et al. 2021; Schmitzer-Torbert and Redish 2008) (Supplementary Figure 3.10). First, the template utilized by Kilosort2 to identify each particular unit was extracted and was classified as “narrow” if the width of the template was less than or equal to 400 μ s. Post-spike suppression indirectly measures the refractory period of a unit and was computed as the time needed for the firing rate to exceed the average firing rate in the 100-400ms period following a spike; this was computed via the auto-correlogram smoothed with a 25ms centered Hamming window. Finally, the phasic ratio was computed as the relative fraction of time a given unit spent in an interspike interval longer than 2 seconds. For units that were not considered narrow, units were deemed medium spiny neurons (MSNs) if their post-spike suppression was greater than 40ms and tonically active neurons (TANs) otherwise. For narrow units, units were deemed fast-spiking interneurons (FSIs) if their post-spike suppression was less than 40ms and their phasic ratio less than 10%. All other units were classified as unidentified interneurons and excluded from cell-type-specific analysis in Figure 3.4.

Data Analysis: Population Decoding

We then utilized a population decoding model to analyze the encoding within striatal units of two particular quantities, the neural cursor and the neural velocity (Figure 3.4). Similar to analysis seen in Figure 3.1J-M, sliding windows of 10 successful trials were first computed across all sessions. Again, sessions with fewer than 10 successful trials were ignored, and only sessions within the first 12 days were considered. Firing rates spanning the 2 seconds preceding the cortical hit of each trial were extracted, and lags of -100ms and +100ms were added for each unit. Then, within each sliding window, encoding strength was assessed via cross-validated ridge regression (sklearn.linear_model.RidgeCV) fit to the trial-concatenated firing rates. In particular, 5-fold shuffled cross-validation was utilized to determine the optimal regularization constant, and all coefficients of determinations (R^2) used in figures and models were the average test R^2 across the 5 folds within that window. For both the neural cursor and neural velocity, a separate ridge regression model was computed within each sliding window for all striatal units and for the 6 combinations of region (dorsal and ventral striatum) and putative cell type (MSN, FSI, and TAN).

To analyze changes in encoding strength across timescales, we utilized a linear mixed model. First, as encoding strengths were highly skewed towards 1, we transformed the cross-validated R^2 values using a log-transform to make it more amenable for linear modeling:

$$\hat{R}^2 = -\log_{10}(1 - R^2)$$

This transformation thus resulted in values between 0 and positive infinity, where more positive values represented encoding strengths closer to 1. This became our dependent variable for all population decoding models. This transformation significantly improved homoscedasticity and normality of residuals, as described below.

As done above, we first designated “early trials” as the first 15 windows of trials (i.e. spanning the first 25 trials), and “late trials” as all other windows of trials. Then, we constructed a model with fixed effects of number of days, indicator of early or late trials (-1 or 1), an interaction term between number of days and early/late indicator, and the number of units recorded that day; a random intercept per session, and a random across-day slope per rat. Critically, the inclusion of the number of units in the model accounted for inevitable increases in encoding strengths by including more predictors, and this slope was always highly significant. The models converged; while heteroscedasticity tests had p-values less than 0.05, homoscedasticity of residuals was confirmed visually for each model, and kurtosis/skew of residuals were within reasonable values (<1 for both). In some cases, the variance explained by the per-rat across-day slope was near zero, resulting in a singular random effects matrix; since this indicates there exists near-zero variation in this parameter across rats, in these cases, the per-rat across-day random slope was removed and the model re-fit.

To directly compare encoding strengths between dorsal and ventral striatum within a given cell type, we took advantage of simultaneous dorsal and ventral striatal recordings by comparing encoding strengths across all sliding windows. In particular, to account for differences in number of units across regions, we first computed the \widehat{R}^2 per unit recorded for each region and sliding window, and computed difference in this relative encoding strength between dorsal and ventral striatum across all sliding windows. A linear mixed model was then used to test whether this difference was significantly different than zero through a fixed effect intercept and a random intercept per session, yielding a p-value for the fixed effect intercept shown in Figure 3.4 (right columns for each cell type). A simpler model testing the mean difference across sliding windows per session for significant differences from zero via a 1-sample t-test yielded the same trends (data not shown).

Additionally, to analyze changes in encoding strength after the unlinking of audio and reward, we compared the encoding strengths of the five sliding windows occurring entirely and immediately before the unlinking to the five sliding windows occurring entirely and immediately after (Supplementary Figure 3.11). A linear mixed model with fixed effect slope and intercept and a random intercept per session was used to analyze differences in encoding strength before and after these two manipulations. All linear mixed models converged, had homoscedastic residuals, and had kurtosis/skew within reasonable bounds (< 2). Relatedly, to assess striatal encoding of the

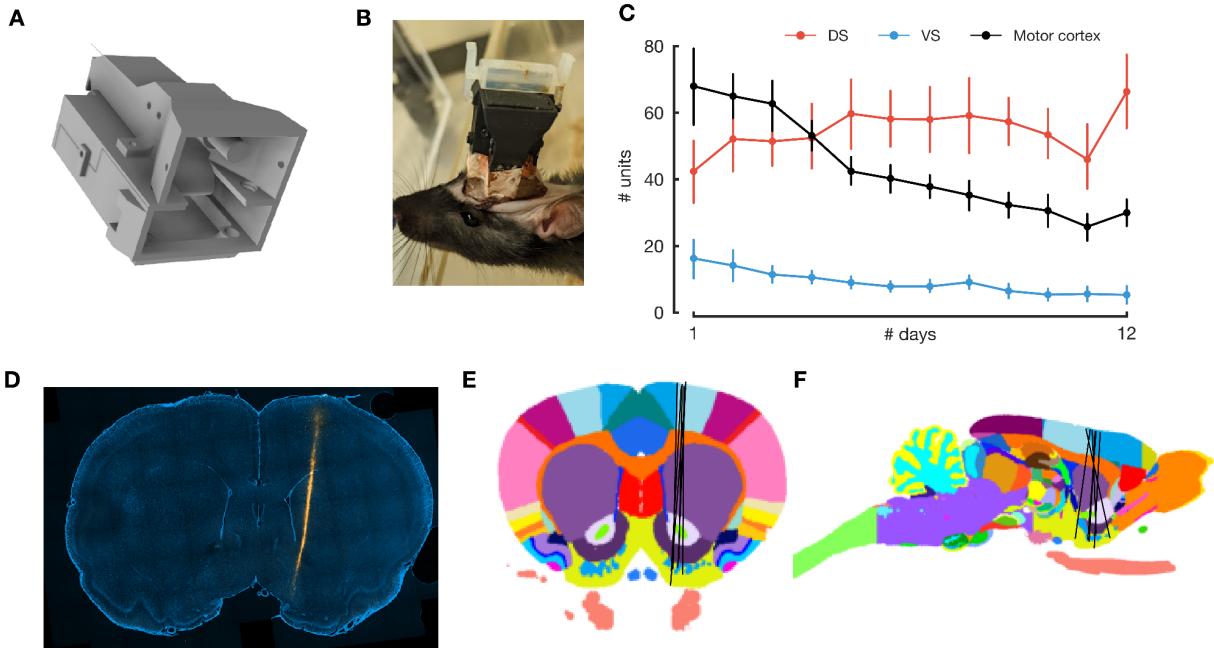
neural cursor and auditory tone frequency during catch trials, we utilized a population decoding model across the full length (~10 seconds) of all catch trials within a given session, and compared encoding strengths with a paired t-test.

To analyze alternative hypotheses of striatal encodings, we compared the encoding strength of the neural cursor with that of a linear representation of time, as striatum has been reported to be correlated with measures of time (Mello et al. 2015) (Supplementary Figure 3.12). We considered the same time periods as previous population decoding analyses (2 seconds before hit), except models were fit to an entire session's trials, instead of windows of trials, to simplify analysis. As a measure of time, a line was constructed that ramped from 0 to the max value of the mean neural cursor, starting at an onset swept across a range of values (-2s, -1.5s, -1.4s, -1.0s, -0.6s, -0.5s, -0.4s, -0.3s, and -0.2s relative to hit). Then, in order to enable a fair comparison of encodings strengths with the neural cursor, each representation of time was variance-matched to the neural cursor: per-time-bin variance was computed for the neural cursor across all successful trials, and Gaussian noise with equal variance was injected into the line. A population decoding model was then fit using units within the respective subgroups of striatal region and cell types. We then only retained the onset for the time representation that had the highest encoding strength. Lastly, similar to Supplementary Figure 3.11, a linear mixed model with fixed effect slope and intercept and a random intercept per session was used to assess differences in encoding strengths.

Similarly, as an alternative hypothesis to neural velocity encodings, striatum could have encoded a discrete "go" signal beginning at some point before hit, corresponding to a step function (Supplementary Figure 3.12). We performed a similar analysis as investigating time, except a step function was used, instead of a line, and amplitudes/variances were matched to the neural velocity signal instead of neural cursor.


3.3 Supplementary Figures and Data

Supplementary Figure 3.1




Supplementary Figure 3.1. Custom-made implant yields tens to hundreds of units each day, and histology confirms placement of probe within targeted areas of striatum. **A:** 3D CAD model of assembled custom-made 3D-printed implant, heavily based on a prior validated design (Luo et al. 2020) (40). View is from the end of the implant that will be facing dorsally. The Neuropixels 1.0 headstage PCB is secured via M1 screws in the bottom half of the implant (as viewed from this angle); the cylinder visible in the top half is used to attach the implant to the stereotaxic frame via a cannula holder. **B:** entire implanted probe assembly on a ~250g rat after a few days of recovery post-surgery. Black and transparent white material is the 3D-printed implant; orange wrapping is copper foil used for shielding. **C:** number of well-isolated units recorded per day, divided into regions as defined in Methods; bars represent SEM. **D:** example stained and imaged histological brain slice. Slice was stained with DAPI (blue); probe track (orange) fluoresces from Vybrant DiI dye applied to the probe shank before implantation. **E, F:** reconstructed probe tracks (n=6) overlaid on coronal (**E**) and sagittal (**F**) planes of the Waxholm Space Atlas. In 1 rat VS units were excluded due to posterior location of probe (see methods). Different colors represent the default colors used to differentiate atlas-defined brain regions.

Supplementary Figure 3.2




Supplementary Figure 3.2. Rats preferred the target they initially achieved. **A:** display of all included sessions ($n=80$) and which target (T1 or T2) was more highly preferred that day, i.e. the target that had more successful trials. Each row represents a single rat, and each box a day. In total, 86.8% of trials went to the preferred target for that day. **B:** measures of possible factors contributing to a rat choosing a particular target on a particular day. Rats did not prefer the target leading to more reward (binomial test for $37/80$, $p=0.58$), did not consistently prefer a particular target overall (T1 chosen arbitrarily; binomial test for $35/80$, $p=0.31$), and did not prefer the target that might have been biased due to calibration (bias determined as the target with more rewards in simulated calibration; each (neuroprosthetic) session was then determined as not biased or biased according to a binomial test p -value of less than 0.05; final predictor significance determined via binomial test for $34/80$ with 3 choices, $p=0.10$). However, rats did significantly prefer whichever target they achieved more frequently within the first 5 trials (binomial test for $63/80$, $p<1e-6$). **C:** The proportion of trials per session going to a preferred target correlates significantly with the proportion of correct trials (left, $p=1.9e-5$) and with time (right, days #1-12 only, $p=0.003$). Individual dots represent 1 session; black line and gray shading indicate regression line and lowess, respectively.

Supplementary Figure 3.3




Supplementary Figure 3.3. Demonstration of time-shifting utilized to align trials, correcting for trial-by-trial variations in precise timing. Optimal time shifts were computed via previously demonstrated algorithms (Williams et al. 2020) that optimized the inter-trial correlations of direct unit activity for successful trials of each target. Left: unaligned, simultaneous firing rates for a unit in ensemble 1 (top) and a unit in ensemble 2 (bottom). White ticks indicate the model’s prediction of the alignment point and will be the points that become “time = zero” after alignment. Right: aligned, simultaneous firing rates for the same units and trials. All other units’ activities (not shown) were then shifted by the same amount as the direct units to maintain alignment.

Supplementary Figure 3.4




Supplementary Figure 3.4. Mean firing rates for reward-responsive units time-locked to target hit during the task (**A**) and to simulated target hit during each task's preceding calibration (**B**). Units are the same across rows. Firing rates were z-scored independently between task and calibration.

Supplementary Figure 3.5




Supplementary Figure 3.5. Performance does not change when the auditory tone is unlinked from cortical activity. In a subset of high-performing, late-learning sessions, the frequency of the auditory tone was unlinked from cortical activity and instead linearly ramped across the frequency spectrum, while the link between cortical activity and reward remained intact. There was no significant change in the hits per minute between the 10 minutes preceding and the 10 minutes following unlinking of audio ($n=7$ sessions, $p=0.66$). In performance-matched, normal sessions, there was a slight but significant increase in performance in a similar time period ($n=48$, $p=0.04$; same as pictured in Figure 3.11).

Supplementary Figure 3.6



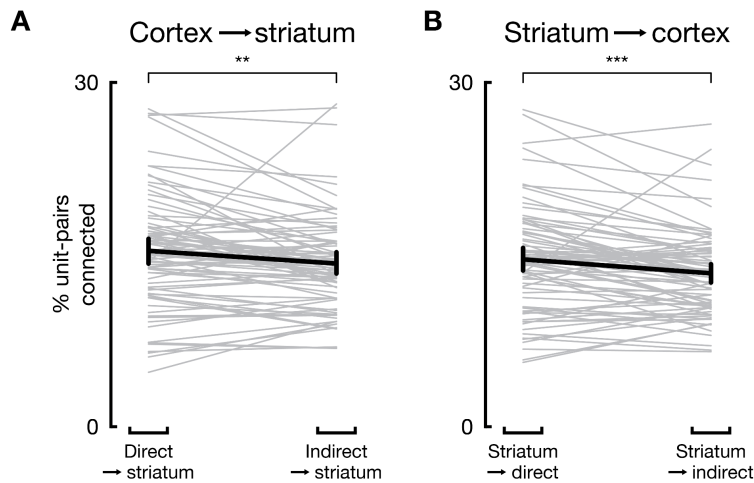
Supplementary Figure 3.6. Further analysis of striatal representations at the end of catch trials. Display and analysis is the same as that of Figure 3.3, with analysis performed on different subsets of catch trials. **A:** trial-to-trial template correlations computed at the end of all catch trials only on the first 2 days. **B:** correlations analysis performed on catch trials following a period of 20 trials that were in the lowest 25th percentile of mean hit times, representing catch trials following poor performance. **C:** correlations analysis for the first catch trials of each day. For all, n.s. $p > 0.05$; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$.

Supplementary Figure 3.7



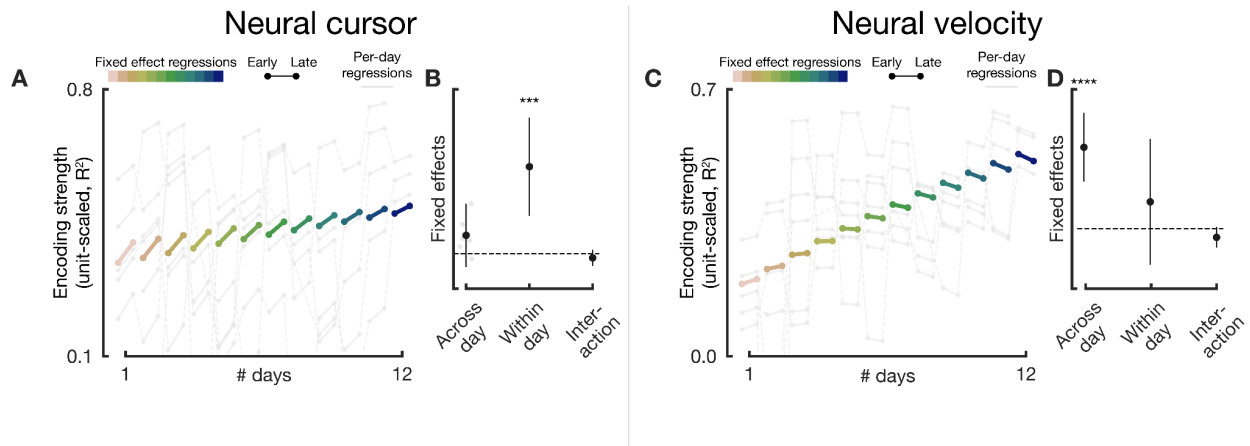
Supplementary Figure 3.7. Further analysis of striatal representations during outside-trial hits. For each outside-trial hit analyzed, the median Δr value for all dorsal (left) / ventral (right) striatal units (x-axis) was plotted against the Δr value for the direct units (y-axis), and a linear regression model was fit (black line: regression line, shading lowess). The intercepts for both dorsal and ventral striatum were significantly negative (DS -0.09, $p < 0.0001$; VS -0.07, $p < 0.001$), indicating that even when cortical activity in extra-trial hits is similar to preceding successful trials, striatal activity still significantly deviates from its current representation.

Supplementary Figure 3.8




Supplementary Figure 3.8. Granger causality analysis between striatal and motor cortical units shows functional connectivity is higher between striatum and direct units than striatum and indirect units. **(A)** Percent of unit-pairs connected with direct (left) or indirect (right) cortical units as the source, and striatal units as the targets ($p = 0.002$, paired t-test). **(B)** Same as A but with source and target reversed ($p = 0.0003$, paired t-test). For all, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$, otherwise not indicated $p > 0.05$; bars represent 95% confidence intervals.

Supplementary Figure 3.9



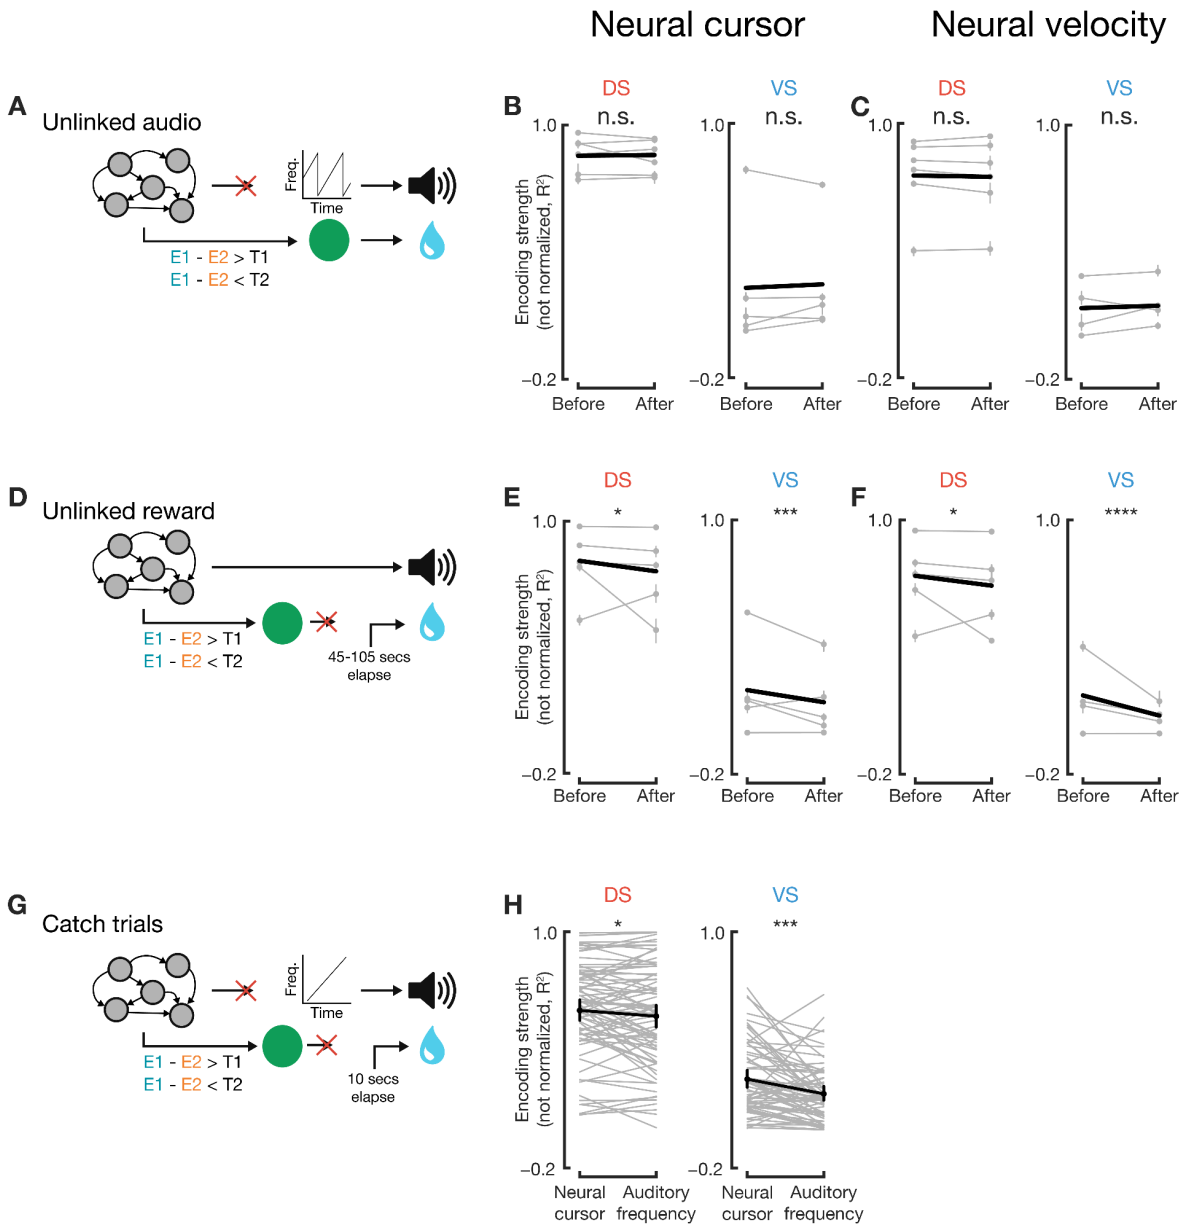
Supplementary Figure 3.9. Results of population decoding model when all striatal units are included for neural cursor encoding (left) and neural velocity encoding (right). **(A)** Regressions from linear mixed model (fixed effects: colored lines; per-day regressions incorporating random effects: gray lines) displaying changes in neural cursor encoding strength both within-day (early vs late trial windows, indicated by endpoints of each line) and across-day changes (x-axis). Encoding strengths shown are scaled to 20 units (“unit-scaled”, Methods). Model intercept for unit-scaled R^2 0.37 ($p < 1e-4$). **(B)** Fixed effect slopes for neural cursor encoding; positive slopes indicate increases in encoding strength. Faded gray dots indicate random effects, when applicable. Y-axis for fixed/random effects is log-transformed R^2 (Methods). **(C-D)** Same as (A) and (B), except with neural velocity encoding. Model intercept for unit-scaled R^2 0.21 ($p < 1e-6$). For all, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$, otherwise not indicated $p > 0.05$; bars represent 95% confidence intervals.

Supplementary Figure 3.10



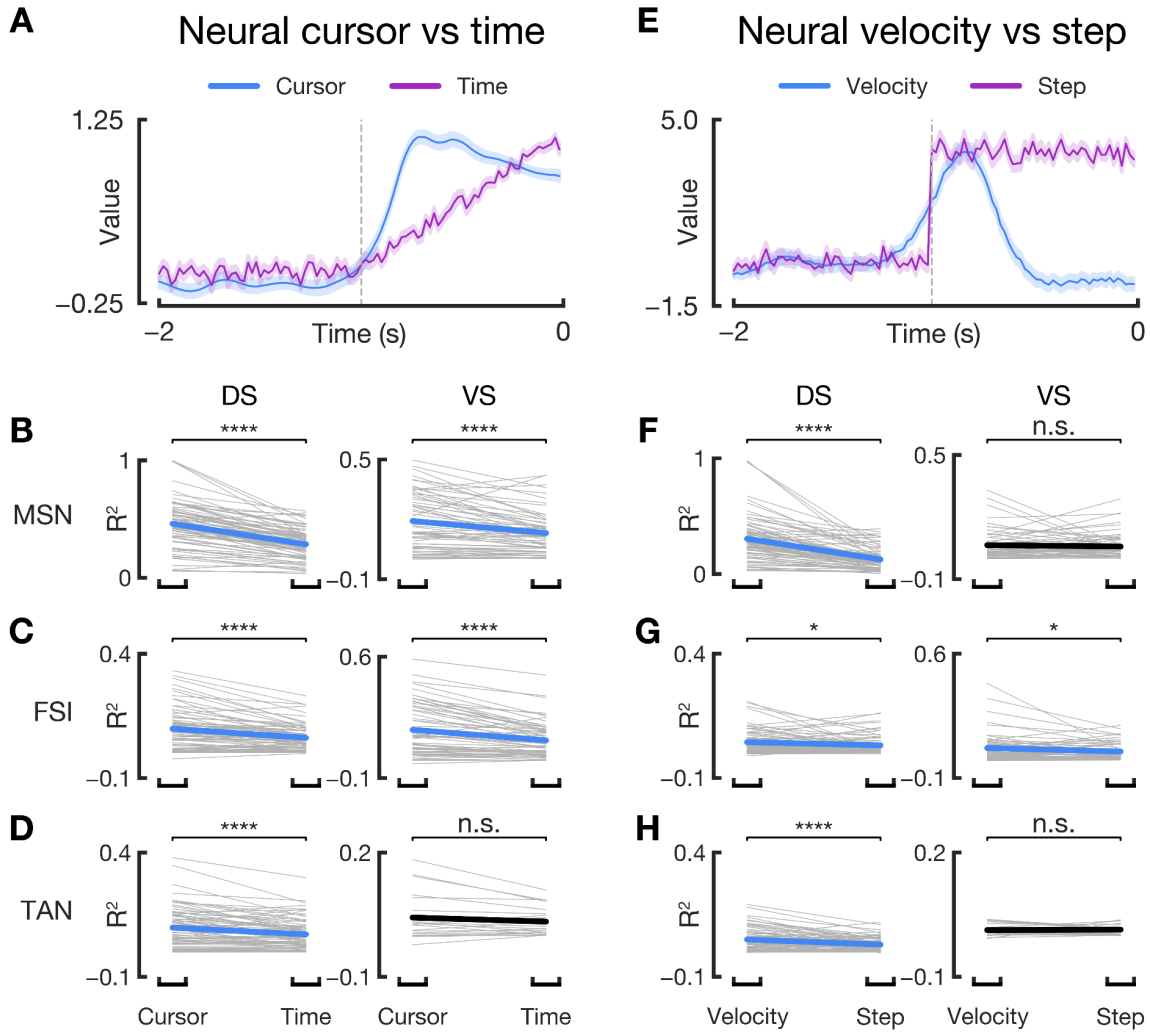
Supplementary Figure 3.10. Classification of units into putative cell types via electrophysiological characteristics. **A:** left: mean template extracted by Kilosort2 for all units classified as a particular cell type (MSN: medium spiny neuron, FSI: fast-spiking interneuron, TAN: tonically active neuron). Shading is 1 standard deviation. Right: mean auto-correlograms for each cell type, displayed from -150ms to 150ms (x-axis). Shading is 1 standard deviation. **B:** distribution of firing rate properties for all well-isolated units within the dorsal striatum identified as particular cell types. Colors are the same as in (A); ISI interspike interval. **C:** same as (B) but for units in ventral striatum (VS). **D:** relative proportion of each cell type within the population of well-isolated dorsal and ventral striatal units. 3943 dorsal striatal units were identified as 3174 MSNs, 228 FSIs, 411 TANs, and 130 unidentified interneurons; 690 ventral striatal units were identified as 387 MSNs, 231 FSIs, 35 TANs, and 37 unidentified interneurons.

Supplementary Figure 3.11



Supplementary Figure 3.11: Population decoding model applied to instances where audio and reward were unlinked and during catch trials, as also depicted at the single-unit level in Figure 3.3.. **A:** when the auditory tone was unlinked from cortical activity in a subset of sessions, there were no significant changes in the encoding strength within dorsal or ventral striatum for the neural cursor or velocity ($p=0.67$, $p=0.20$, $p=0.55$, $p=0.39$ for dorsal-cursor (B, left), ventral-cursor (B, right), dorsal-velocity (C, left), ventral-velocity (C, right) encodings, respectively). **D:** when reward was unlinked from cortical activity, encoding strength for both neural cursor and velocity significantly decreased in both dorsal and ventral striatum ($p=0.03$, $p=0.0003$, $p=0.02$, $p=0.0001$ for dorsal-cursor (E, left), ventral-cursor (E, right), dorsal-velocity (F, left), ventral-velocity encodings (F, right), respectively). **G:** during catch trials, the auditory tone and the neural cursor were decoupled. Both dorsal (H, left) and ventral striatum (H, right) encoded the neural cursor more strongly than the auditory tone's frequency (DS $p=0.03$, $n=80$ sessions, VS $p=0.0005$, $n=67$ sessions; paired t-test). In all, thin gray lines represent individual session's data; thick black lines indicate regressions. Bars indicate 95% confidence interval.

Supplementary Figure 3.12



Supplementary Figure 3.12: Population decoding model applied to alternative encoding hypotheses, where striatal activity may encode a linear representation of time (Mello et al. 2015) instead of the neural cursor (left column), or may encode the time derivative of this representation of time, a step function, instead of the neural velocity (right column). **(A)** Example neural cursor and time data. Time (purple) functions were constructed based on a linear ramp from 0 to the max value of the mean neural cursor, where the onset time (gray dotted line) was swept across a range of times. Additionally, Gaussian noise was injected into the time function to match the variance of the neural cursor to enable a fair comparison for encoding. Shading SEM. **(B-D)** Comparison of encoding strengths (R^2 , not unit-scaled) across an entire session for the encoding of the neural cursor (left datapoint of each column) and time (right datapoint of each column) for all 6 combinations of regions & putative cell types. For each session, the best onset for the time function was taken. Significance of the slope (i.e. the difference between cursor and time encoding strengths) was computed via a linear mixed model. Bold line indicates fixed effect; black bold line means slope is not significant, otherwise the color represents which quantity is more strongly encoded. Gray lines represent each session. As in other figures: n.s. $p > 0.05$; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$. **(E)** Similar to (A), a step function was constructed, with scale equal to the max mean neural velocity value and onset time swept across a range of times. Variances were matched as well. Data originates from the same session as (A). **(F-H)** Same as (B-D), but comparing neural velocity (left datapoint) and step function (right datapoint) encodings.

Chapter 4

Scalable Software for Neuroscience

Shifting gears, we now delve into two open-source software frameworks built as a part of this thesis’s work. We will only briefly describe each tool by providing context and an implementation overview, but otherwise delegate to documentation linked within each software’s repository for details on installation, usage, and development.

4.1 River: a cross-platform, high-throughput, structured streaming framework

Repository: <https://github.com/pbotros/river>

Modern experiments in neuroscience demand increasingly complex and performant systems. Technological advances in the past decades have led to an explosion in neural data volume, with current technologies capable of simultaneous recordings of thousands of individual neurons (Steinmetz et al. 2021; Ota et al. 2021). Concomitantly, to improve scientific efficiency (Z. S. Chen and Pesaran 2021) or modulate neural or physical systems (Shanechi 2019), closed-loop experimental paradigms that multiplex neural and non-neural (e.g. kinematic) data have come to the forefront of research, where real-time computations on various data streams directly influence the experiment. Considering these ever-increasing demands on both latency and throughput of data, there is a dire need in the field to stream high-throughput data between devices, balancing requirements for performance with simplicity for non-software experts.

To tackle these problems, we developed River, a structured streaming solution to streamline data management for modern neuroscience experiments. Producers of data – “writers” – first define a structure for each data sample for a particular stream. Then, utilizing function calls similar to those of file I/O, writers can then write data to that stream, while an unlimited number of “readers” can consume this data in real-time by specifying the appropriate stream name. Each reader can read at its own pace, with support for blocking for a given period of time or number of samples. Readers and writers only require network connectivity to a centralized server, enabling a horizontal scaling paradigm useful for computationally-intensive or cloud-reliant experiments. Finally, leveraging the structured nature of each stream, a separate “ingester” process consumes each stream’s data and writes data to disk in a cross-platform tabular format, Parquet (<https://github.com/apache/parquet-format>), for post-hoc analysis. Once persisted by the ingester, sufficiently stale data is then deleted from the River stream, enabling streams to be indefinitely long. River is written in C++ with Python and MATLAB bindings and works under Linux, Mac OSX, and Windows. An example deployment, highlighting River architecture, can be found in the below figure.




Figure 4.1. Example deployment of 1 writer and 2 readers utilizing River, where the writer reads data from a probe, one reader visualizes the data, and another reader utilizes the data to actuate an external device. Readers and writers may exist on different devices, with network connectivity to the in-memory database, Redis, the only requirement. Dotted line demarcates the boundary of River APIs, with monospaced-font commands referring to River APIs invoked to write, read, and ingest data. Note all interactions with the underlying database, Redis, and nuances of ingestion are abstracted behind River APIs, enabling users to both stream online and analyze offline data with simpler commands.

Eschewing a bespoke, research-specific solution to reduce maintenance burden, we utilized an industry-standard database, Redis (<https://redis.io>), to underlie River, enabling sub-millisecond latencies, low jitter, and 20+ MB/s throughput. Critical to researchers, River guarantees that no written data can be dropped by any readers, even in the case of long pauses in reading, and that all persisted data exactly mirrors the data streamed online. These guarantees come in contrast to the best-effort paradigm of another popular research-oriented streaming framework (<https://github.com/scn/labstreaminglayer>). Additionally, timestamp management is centralized and thus stream synchronization is simplified. Other streaming solutions widely used in the software industry, such as Kafka (<https://kafka.apache.org/>) and RabbitMQ (<https://www.rabbitmq.com/>), have more strict guarantees on data availability and delivery but subsequently suffer in raw streaming performance and can be difficult to set up robustly.

Finally, River has been successfully used in two closed-loop neuroscience studies by the author (Formento et al. 2021) (Botros et. al in preparation). While River was designed with neuroscience in mind, we believe it can have applications in other fields that might have similar data management needs.

4.2 DuraPy: a service container for service-oriented neuroscience experiments

Repository: <https://github.com/pbotros/durapy>

Modern experiments in neuroscience demand increasingly complex and performant systems, with exploding neural data volumes (Steinmetz et al. 2021) and rising interest in closed-loop systems from both scientific (Z. S. Chen and Pesaran 2021) and neuromodulation (Shanechi 2019) perspectives. Current approaches to neuroscience software adopt a monolithic architecture, where a single parent process running on a single computer manages the entirety of the experiment (e.g., OpenEphys (Siegle et al. 2017), RigBox (J. Bhagat et al. 2020), and OpenBCI (https://github.com/OpenBCI/OpenBCI_GUI)) (Figure 4.1). A monolithic architecture enables simple deployment, as it is simply a single process to run, and can be performance-optimal since all data and code resides in the same computer. However, especially in neuroscience, a monolithic architecture hinders scalability: multiple devices (e.g. neural probes, cameras, and position trackers) are likely queried simultaneously at high sampling rates, and computationally-intensive computations may need to be performed in real-time (Barsakcioglu et al. 2020; Hu et al. 2018). These functions compete for a limited set of hardware resources – in the form of PCI slots, USB slots, and network bandwidth – and computer resources such as CPU and memory, as monolithic scaling is entirely vertical. Additionally, a monolithic architecture does not encourage functional separation of components via data-based application programming interfaces (APIs), which could increase the complexity and thus hinder the maintainability of the software. Finally, a monolithic application constrains functionality to a single operating system, which may not be ideal if desirable tools only support different operating systems.

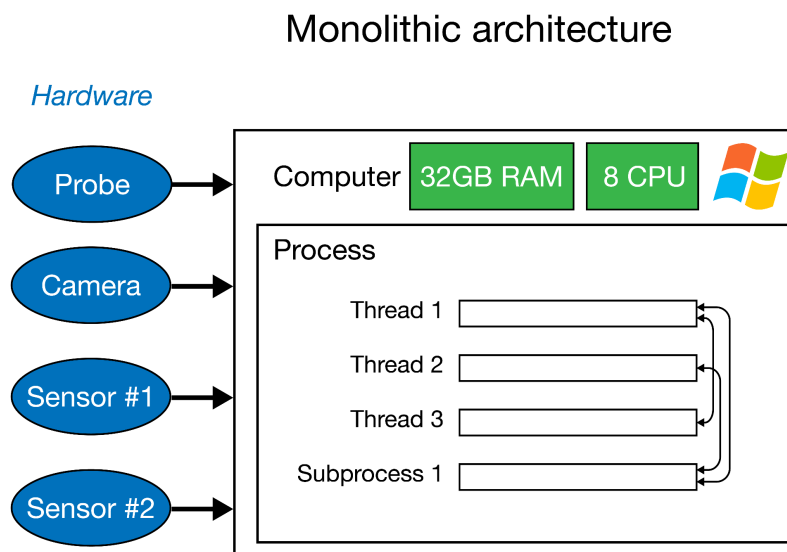


Figure 4.1. An example configuration of a monolithic application running a neuroscience experiment. A single computer runs a single process encapsulating all experimental data and logic. Within this process, many threads and subprocesses run,

communicating with one another via file, shared-memory, or across-memory mechanisms. All hardware (e.g. a probe, camera, and sensors) connects to this single computer.

The pitfalls of monolithic architecture have long been noted in the software engineering industry. The dominant solution in industry is to adopt a service-oriented architecture, which extends well-known software practices of abstraction and compartmentalization of code to the level of networks and processes. In a microservice-oriented architecture, a functional unit that performs some subset of functionality is run in its own process, termed a “microservice”. When microservices need to consume or produce data to other microservices, they communicate with one another over a network. This distributed architecture enables microservices to reside on different devices and different operating systems, powering horizontal scaling in addition to vertical scaling. In addition, the constraint that service-to-service communication occurs over a network forces the format of that data to be defined explicitly in code, which, if properly designed, abstracts away implementation-specific details of each functional unit and can thus increase the maintainability of the overall system. However, these design decisions of a service-oriented architecture create a more complex deployment, requiring additional tooling to utilize its benefits over a monolithic architecture.

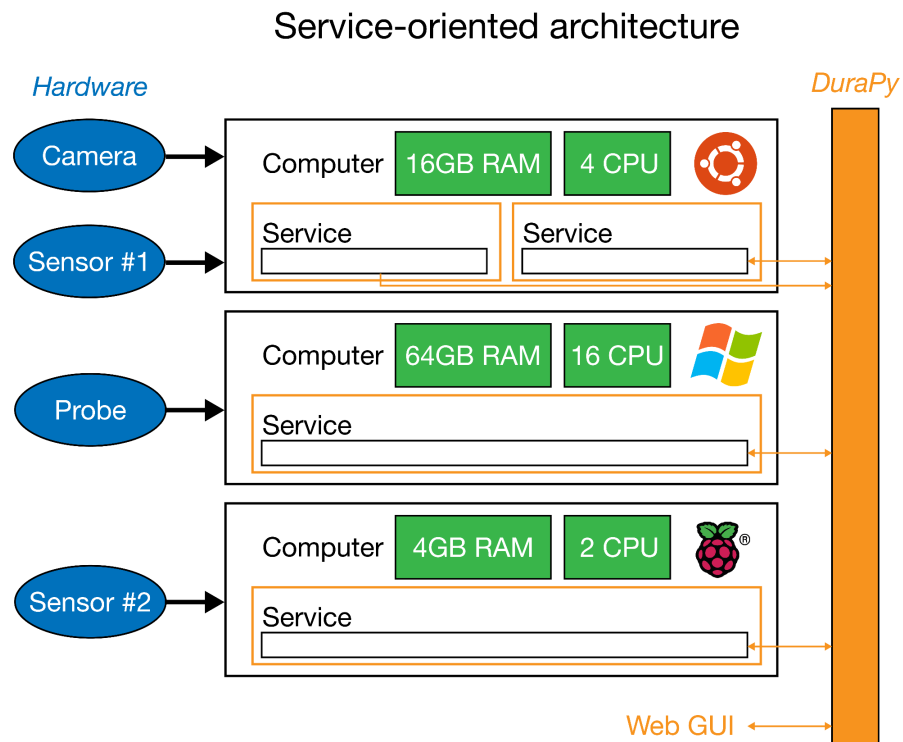


Figure 4.2. An example configuration of a service-oriented architecture running a neuroscience experiment, similar to the example functionality depicted in the monolithic architecture of Figure 4.1. Individual services, encapsulated in orange boxes, represent compartmentalized functional units, running within the DuraPy framework. Communication between microservices occurs over the network, via commands implemented through the DuraPy framework. Hardware connects to whichever computers require it, and DuraPy services can run on heterogeneous computers (as depicted by differing amounts of RAM and CPU) and operating systems (e.g. Raspberry Pi, Ubuntu, and Windows). Finally, a web graphical user interface (GUI) connects to the DuraPy framework, enabling experimenters to build monitoring tools or experiment GUIs within a web browser.

DuraPy is an open-source Python service container for neuroscience experiments utilizing a service-oriented architecture, aiming to provide the communication and process backends for each microservice (Figure 4.2). DuraPy, at its core, utilizes an event-driven paradigm: “commands” can be sent from any service and are broadcast to all services. Each command can have its own structure, and per-command handlers can be installed in each service that calls a specified function when receiving a command, enabling each service to perform custom functionality per command. As part of its service architecture, each DuraPy service runs in its own long-running process, with commands likely sent by the experimenter to initiate and terminate experiments. Logging, Git-based deployment, and POSIX-compatible run scripts are provided as well.

In addition, a web-based graphical user interface (GUI), running itself as a DuraPy microservice, is provided to enable experimenter control of the experiment (Figure 4.2). In addition to supporting the sending and displaying of commands, managing service deployment, and viewing logging, the web GUI implementation, built on an open-source web framework Tornado (<https://github.com/tornadoweb/tornado>), is explicitly designed to enable extensions, where experimenters can more easily build custom web pages for experiment monitoring or review. Extensions of this web GUI can additionally play a role itself in the experiment, such as displaying an interface to subjects (Formento et al. 2021). Finally, DuraPy is intentionally light on any imposed structure of each service, as more rigid models, such as those that utilize a domain-specific language (DSL) or finite-state-machine (FSM) to specify functionality (e.g., LiCoRICE (<https://github.com/bil/licorice>), RigBox (J. Bhagat et al. 2020)), can hinder flexibility. The default command database requires a running instance of Redis, an open-source in-memory database, but future changes could enable DuraPy to utilize other existing networked communication frameworks, such as Robot Operating System 2 (ROS2; <https://github.com/ros2>).

As currently designed, DuraPy’s communication backend is not explicitly designed for high-throughput, low-latency data streaming, as the commands broadcast to each microservice are not envisioned to contain large amounts of data (e.g. neural data, camera frames, raw sensor data, etc.). If desired, existing data streaming solutions can be leveraged to fill this need, such as River (<https://github.com/pbotros/river>), ZeroMQ, or RabbitMQ, where stream metadata can be shared via DuraPy commands but raw stream data remains within the purview of these frameworks. For performance-intensive processes running on the same computer, shared-memory stores such as Apache Plasma (<https://arrow.apache.org/docs/python/plasma.html>) and Cthulhu (<https://github.com/facebookresearch/labgraph>) can also be utilized for optimal performance.

DuraPy has been successfully used in two closed-loop neuroscience studies by the author (Formento et al. 2021) (Botros et. al in preparation) and enabled the bridging of multiple operating systems (Windows, Linux), multiple device types (cloud-based and local desktop computers, and Raspberry Pi computers), and multiple languages (C++ via a thin HTTP wrapper and Python). As a particular example highlighting the benefits of the microservices framework espoused by DuraPy, requirements changed over the course of these research studies calling for additional real-time machine learning functionality to be added, both of which benefited from new

hardware (i.e. high-powered GPUs). These new functionalities were implemented as separate DuraPy microservices and run on newly-introduced desktop computers (Botros et. al in preparation) or nodes running in Amazon's cloud (Formento et al. 2021) without requiring significant changes to other microservices.

While DuraPy was designed with neuroscience in mind, we believe it can have applications in other fields that might have similar architectural needs.

Chapter 5

Conclusion

By utilizing neural-machine interfaces to interrogate the functionality of the nervous system, the results described in this thesis contribute to our understanding of skill learning, demonstrating novel levels of flexibility in the peripheral nervous system with applications in clinical translation and suggesting a critical role of the striatum in skill learning as a part of credit assignment, a core neural mechanism of learning. In addition, the software frameworks implemented and open-sourced as a result of this thesis aim to accelerate future neuroscientists' endeavors.

5.1 Future Work

Motor-unit neuromuscular-machine interface in alternative muscles

The motor-unit neuromuscular-machine interface (NMI) described in the first part of this thesis explored the dimensionality of the biceps brachii muscle, which is a known multifunctional muscle classically governing elbow flexion and forearm supination functions. While the performed analysis showed categorical differences between neurofeedback-enabled motor unit activity and these two stereotyped movements and a resulting dimensionality beyond 2, exploring the motor unit dimensionality in other muscles could provide useful insight into the applicability of an NMI as an assistive device. In particular, investigation into muscles typically retained in disabled patients, such as those in the neck, ears, and shoulders, could shed light into the usefulness of such a device to restore functionality in the population most often targeted by brain-machine interfaces, i.e. tetraplegics.

Motor-unit neuromuscular-machine interface in disabled patients

As mentioned in Chapter 2, recent work demonstrated residual motor unit activity under voluntary control in the forearm muscles of a tetraplegic patient (Ting, Del Vecchio, et al. 2021). Stemming from this promising result, future research into a motor-unit NMI should investigate if this residual motor unit activity in functionally disabled limbs can be exploited for high-performance control. Relatedly, a recent BMI study enabled near-native rates of digital communication by decoding thoughts of handwriting from patients (Willett et al. 2021). In this study, a supplementary video displays the patient miming handwriting with their hands while the BMI is being used. This observation begs the question if there might be sufficient residual motor unit activity remaining in arm or hand muscles to similarly decode handwriting in real-time, with

such a device possibly retaining performance while obviating the need for invasive brain implant surgery.

Real-time detection of motor unit basis vectors

The NMI paradigm utilized in Chapter 2 relied on the participant to explicitly decide on three motor units that they deemed most independently controllable, and thus was a subjective and often challenging decision. Additionally, participants were only able to select a single motor unit per dimension to be controlled, resulting in a fairly straightforward but crude decoder. Future research could address this by devising algorithms to robustly detect controllable dimensions from motor unit activity. Such an identification algorithm is not trivial, since the dimensionality is not known *a priori* and the “basis vectors” of motor unit activity will undoubtedly change as the participant trains with the NMI and putatively increases dimensionality. A two-learner system has been discussed in the brain-machine interface world (Orsborn and Pesaran 2017) where both the decoder and the participant learn to optimize control, so principles from that line of research could be re-used. Possible methodologies to explore are linear decomposition methods such as principal component analysis (PCA), linear constrained decomposition methods such as non-negative matrix factorization (NMF), non-linear machine learning techniques that optimize a cost function assessing the independence of given sets of motor units.

Minimally-invasive, chronic motor unit recording modalities

While surface EMG recordings are dominant in research and industry since they are non-invasive and robust, surface EMG signals change as the skin moves relative to the muscle, a problem that often necessitates isometric, fixed setups in research in order to keep motor unit recordings stable over the length of an experiment. Furthermore, removal and replacement of the surface EMG grid across different days can confound the shape and identity of motor units, though some research studies use markers on the skin and particular algorithms to enable putative identification of the same motor units across days (Martinez-Valdes et al. 2017).

In order to improve the applicability of surface EMG recordings to the clinic, stable recordings of individual motor units should be possible in dynamic settings. One way this can be enabled is via minimally-invasive, chronically-implanted recording grids, which, like chronic brain implants, could enable tracking of individual motor units across days and during dynamic movements. However, such a system has yet to be demonstrated.

Perturbations to dorsal/ventral striatal circuitry to further interrogate learning

Chapter 3 investigates the role of the striatum in relation to cortical credit assignment, based on the use of a neuroprosthetic task and the correlations stemming from simultaneous recordings of cortex and striatum. Future research could manipulate the neural circuits at play to provide causal evidence of dorsal/ventral striatum’s role in skill learning. In particular, our study demonstrates stereotyped responses of individual striatal neurons constituting a model of the

proximity of cortical activity to reward. Optogenetic manipulations could theoretically enable both the disruption of such a model in trained animals and the artificial creation of such a model in untrained animals. Of additional interest would be the effect of the amount of training on the effect of these manipulations, as the striatum is known to change in roles as the behavior shifts from goal-directed to habitual (Graybiel and Grafton 2015). Finally, if the conclusions implied in the second part of this thesis hold true, then precise stimulation within the striatum during stereotyped cortical activity, immediately followed with stimulation to dopaminergic neurons in the VTA (Athalye et al. 2018), could form a completely internal system that alone drives cortical credit assignment and thus effectively demonstrates a standalone micro-circuit of learning. That said, while brain manipulation via optogenetics at the cellular level is increasingly possible, robust manipulation in deep structures and/or with freely behaving animals remains challenging (Pégaré et al. 2017), and so changes to the experimental structure, animal model, and neural recording modalities would likely be required.

Investigation of specific pathways involving the striatum for skill learning

The striatum has a multitude of cortical and subcortical inputs; in particular, the dorsomedial and ventral striatum receive inputs from prefrontal cortex, with both direct and re-entrant projections into motor cortex (Aoki et al. 2019). The ventral striatum also has a multitude of connections with reward-relevant subcortical structures, such as the amygdala, hippocampus, and VTA. Future research could tease apart the differential role of these various input and output pathways by pathway-selective optogenetic tracing and manipulation. In particular, different optogenetic fluorescent markers and anterograde/retrograde tracers could be used to separate recordings in the striatum according to their input pathways, and optogenetic manipulations of these tagged pathways could provide casual evidence of their role. One theory of the ventral striatum is that the role of the various subcortical inputs to the ventral striatum differs depending on the structure of the task; for example, navigational or mapping-intensive tasks might utilize the hippocampus more strongly (Aronov et al. 2017).

Open-source software dissemination

Finally, while the open-source software discussed in Chapter 4 has been released publicly and documented, effective dissemination of software remains challenging. Future work could partner with existing open-source software companies or organizations, such as OpenEphys, the Chan-Zuckerberg Initiative, or Schmidt Futures' Virtual Institute for Scientific Software, to help maintain and/or disseminate the software built here to accelerate the neuroscience community.

Chapter 6

Bibliography

- Aoki, Sho, Jared B Smith, Hao Li, Xunyi Yan, Masakazu Igarashi, Patrice Coulon, Jeffery R Wickens, Tom JH Ruigrok, and Xin Jin. 2019. “An Open Cortico-Basal Ganglia Loop Allows Limbic Control over Motor Output via the Nigrothalamic Pathway.” *ELife* 8 (September): e49995. <https://doi.org/10.7554/eLife.49995>.
- Aronov, Dmitriy, Rhino Nevers, and David W. Tank. 2017. “Mapping of a Non-Spatial Dimension by the Hippocampal–Entorhinal Circuit.” *Nature* 543 (7647): 719–22. <https://doi.org/10.1038/nature21692>.
- Atallah, Hisham E, Dan Lopez-Paniagua, Jerry W Rudy, and Randall C O’Reilly. 2007. “Separate Neural Substrates for Skill Learning and Performance in the Ventral and Dorsal Striatum.” *Nature Neuroscience* 10 (1): 126–31. <https://doi.org/10.1038/nn1817>.
- Atallah, Hisham E., Andrew D. McCool, Mark W. Howe, and Ann M. Graybiel. 2014. “Neurons in the Ventral Striatum Exhibit Cell-Type-Specific Representations of Outcome during Learning.” *Neuron* 82 (5): 1145–56. <https://doi.org/10.1016/j.neuron.2014.04.021>.
- Athalye, Vivek R, Jose M Carmena, and Rui M Costa. 2020. “Neural Reinforcement: Re-Entering and Refining Neural Dynamics Leading to Desirable Outcomes.” *Current Opinion in Neurobiology* 60 (February): 145–54. <https://doi.org/10.1016/j.conb.2019.11.023>.
- Athalye, Vivek R., Karunesh Ganguly, Rui M. Costa, and Jose M. Carmena. 2017. “Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control.” *Neuron* 93 (4): 955–970.e5. <https://doi.org/10.1016/j.neuron.2017.01.016>.
- Athalye, Vivek R., Fernando J. Santos, Jose M. Carmena, and Rui M. Costa. 2018. “Evidence for a Neural Law of Effect.” *Science* 359 (6379): 1024–29. <https://doi.org/10.1126/science.aao6058>.
- Averbeck, Bruno B., and Elisabeth A. Murray. 2020. “Hypothalamic Interactions with Large-Scale Neural Circuits Underlying Reinforcement Learning and Motivated Behavior.” *Trends in Neurosciences* 43 (9): 681–94. <https://doi.org/10.1016/j.tins.2020.06.006>.
- Barbero, Marco, Roberto Merletti, and Alberto Rainoldi. 2012. *Atlas of Muscle Innervation Zones: Understanding Surface Electromyography and Its Applications*. Milan ; New York: Springer.
- Barr, Dale J., Roger Levy, Christoph Scheepers, and Harry J. Tily. 2013. “Random Effects Structure for Confirmatory Hypothesis Testing: Keep It Maximal.” *Journal of Memory and Language* 68 (3): 255–78. <https://doi.org/10.1016/j.jml.2012.11.001>.
- Barsakcioglu, Deren Y., Mario Bracklein, Ales Holobar, and Dario Farina. 2020. “Control of Spinal Motoneurons by Feedback from a Non-Invasive Real-Time Interface.” *IEEE Transactions on Biomedical Engineering*, 1–1. <https://doi.org/10.1109/TBME.2020.3001942>.
- Basmajian, J. V. 1963. “Control and Training of Individual Motor Units.” *Science* 141 (3579): 440–41. <https://doi.org/10.1126/science.141.3579.440>.
- Basmajian, John V., and Carlo J. De Luca. 1985. *Muscles Alive: Their Functions Revealed by*

- Electromyography*. 5th ed. Baltimore: Williams & Wilkins.
- Bawa, Parveen N. S., Kelvin E. Jones, and Richard B. Stein. 2014. "Assessment of Size Ordered Recruitment." *Frontiers in Human Neuroscience* 8. <https://doi.org/10.3389/fnhum.2014.00532>.
- Benabid, Alim Louis, Thomas Costecalde, Andrey Eliseyev, Guillaume Charvet, Alexandre Verney, Serpil Karakas, Michael Foerster, et al. 2019. "An Exoskeleton Controlled by an Epidural Wireless Brain–Machine Interface in a Tetraplegic Patient: A Proof-of-Concept Demonstration." *The Lancet Neurology* 18 (12): 1112–22. [https://doi.org/10.1016/S1474-4422\(19\)30321-7](https://doi.org/10.1016/S1474-4422(19)30321-7).
- Berg, Stuart, Dominik Kutra, Thorben Kroeger, Christoph N. Straehle, Bernhard X. Kausler, Carsten Haubold, Martin Schiegg, et al. 2019. "Ilastik: Interactive Machine Learning for (Bio)Image Analysis." *Nature Methods* 16 (12): 1226–32. <https://doi.org/10.1038/s41592-019-0582-9>.
- Bhagat, Jai, Miles J. Wells, Kenneth D. Harris, Matteo Carandini, and Christopher P. Burgess. 2020. "Rigbox: An Open-Source Toolbox for Probing Neurons and Behavior." *ENEURO* 7 (4): ENEURO.0406-19.2020. <https://doi.org/10.1523/ENEURO.0406-19.2020>.
- Bhagat, Nikunj A., Anusha Venkatakrishnan, Berdakh Abibullaev, Edward J. Artz, Nuray Yozbatiran, Amy A. Blank, James French, et al. 2016. "Design and Optimization of an EEG-Based Brain Machine Interface (BMI) to an Upper-Limb Exoskeleton for Stroke Survivors." *Frontiers in Neuroscience* 10 (March). <https://doi.org/10.3389/fnins.2016.00122>.
- Blabe, Christine H., Vikash Gilja, Cindy A. Chestek, Krishna V. Shenoy, Kim D. Anderson, and Jaimie M. Henderson. 2015. "Assessment of Brain-Machine Interfaces from the Perspective of People with Paralysis." *Journal of Neural Engineering* 12 (4): 43002. <https://doi.org/10.1088/1741-2560/12/4/043002>.
- Borzelli, Daniele, Marco Gazzoni, Alberto Botter, Laura Gastaldi, Andrea d'Avella, and Taian M. Vieira. 2020. "Contraction Level, but Not Force Direction or Wrist Position, Affects the Spatial Distribution of Motor Unit Recruitment in the Biceps Brachii Muscle." *European Journal of Applied Physiology* 120 (4): 853–60. <https://doi.org/10.1007/s00421-020-04324-6>.
- Botros, Paul. 2021. *Pbotros/River: Alpha* (version v0.1). Zenodo. <https://doi.org/10.5281/ZENODO.4624271>.
- Bouton, Chad E., Ammar Shaikhouni, Nicholas V. Annetta, Marcia A. Bockbrader, David A. Friedenberg, Dylan M. Nielson, Gaurav Sharma, et al. 2016. "Restoring Cortical Control of Functional Movement in a Human with Quadriplegia." *Nature* 533 (October 2017): 247–50. <https://doi.org/10.1038/nature17435>.
- Buchanan, Thomas S., and David G. Lloyd. 1995. "Muscle Activity Is Different for Humans Performing Static Tasks Which Require Force Control and Position Control." *Neuroscience Letters* 194 (1–2): 61–64. [https://doi.org/10.1016/0304-3940\(95\)11727-E](https://doi.org/10.1016/0304-3940(95)11727-E).
- Burguière, Eric, Patricia Monteiro, Luc Mallet, Guoping Feng, and Ann M Graybiel. 2015. "Striatal Circuits, Habits, and Implications for Obsessive–Compulsive Disorder." *Current Opinion in Neurobiology* 30 (February): 59–65. <https://doi.org/10.1016/j.conb.2014.08.008>.
- Chen, Xiaogang, Yijun Wang, Masaki Nakanishi, Xiaorong Gao, Tzyy-Ping Jung, and Shangkai Gao. 2015. "High-Speed Spelling with a Noninvasive Brain–Computer Interface." *Proceedings of the National Academy of Sciences* 112 (44): E6058–67.

- <https://doi.org/10.1073/pnas.1508080112>.
- Chen, Zhe Sage, and Bijan Pesaran. 2021. “Improving Scalability in Systems Neuroscience.” *Neuron* 109 (11): 1776–90. <https://doi.org/10.1016/j.neuron.2021.03.025>.
- Choi, John, Katie Wingel, Adam Charles, Krishan Kumar, Mahdi Choudhury, Bijan Pesaran, and Mohammad Khazali. 2020. “Optimal Adaptive Electrode Selection to Maximize Simultaneously Recorded Neuron Yield.” Preprint. Neuroscience. <https://doi.org/10.1101/2020.10.06.328526>.
- Clancy, Kelly B, Aaron C Koralek, Rui M Costa, Daniel E Feldman, and Jose M Carmena. 2014. “Volitional Modulation of Optically Recorded Calcium Signals during Neuroprosthetic Learning.” *Nature Neuroscience* 17 (6): 807–9. <https://doi.org/10.1038/nn.3712>.
- Collinger, Jennifer L., Michael L. Boninger, Tim M. Bruns, Kenneth Curley, Wei Wang, and Douglas J. Weber. 2013. “Functional Priorities, Assistive Technology, and Brain-Computer Interfaces after Spinal Cord Injury.” *The Journal of Rehabilitation Research and Development* 50 (2): 145. <https://doi.org/10.1682/JRRD.2011.11.0213>.
- Collinger, Jennifer L, Brian Wodlinger, John E Downey, Wei Wang, Elizabeth C. Tyler-Kabara, Douglas J Weber, Angus J.C. McMorland, Meel Velliste, Michael L Boninger, and Andrew B Schwartz. 2013. “High-Performance Neuroprosthetic Control by an Individual with Tetraplegia.” *The Lancet* 381 (9866): 557–64. [https://doi.org/10.1016/S0140-6736\(12\)61816-9](https://doi.org/10.1016/S0140-6736(12)61816-9).
- Costa, Rui M. 2011. “A Selectionist Account of de Novo Action Learning.” *Current Opinion in Neurobiology* 21 (4): 579–86. <https://doi.org/10.1016/j.conb.2011.05.004>.
- Costa, Vincent D., Olga Dal Monte, Daniel R. Lucas, Elisabeth A. Murray, and Bruno B. Averbeck. 2016. “Amygdala and Ventral Striatum Make Distinct Contributions to Reinforcement Learning.” *Neuron* 92 (2): 505–17. <https://doi.org/10.1016/j.neuron.2016.09.025>.
- Dabney, Will, Zeb Kurth-Nelson, Naoshige Uchida, Clara Kwon Starkweather, Demis Hassabis, Rémi Munos, and Matthew Botvinick. 2020. “A Distributional Code for Value in Dopamine-Based Reinforcement Learning.” *Nature* 577 (7792): 671–75. <https://doi.org/10.1038/s41586-019-1924-6>.
- De Luca, C. J., and Z Erim. 1994. “Common Drive of Motor Units in Regulation of Muscle Force.” *Trends in Neurosciences* 17 (7): 299–305. [https://doi.org/10.1016/0166-2236\(94\)90064-7](https://doi.org/10.1016/0166-2236(94)90064-7).
- De Luca, C. J., and B. Mambrito. 1987. “Voluntary Control of Motor Units in Human Antagonist Muscles: Coactivation and Reciprocal Activation.” *Journal of Neurophysiology* 58 (3): 525–42. <https://doi.org/10.1152/jn.1987.58.3.525>.
- Deligkaris, Kosmas, Torsten Bullmann, and Urs Frey. 2016. “Extracellularly Recorded Somatic and Neuritic Signal Shapes and Classification Algorithms for High-Density Microelectrode Array Electrophysiology.” *Frontiers in Neuroscience* 10 (September). <https://doi.org/10.3389/fnins.2016.00421>.
- Desmedt, J. E., and E. Godaux. 1981. “Spinal Motoneuron Recruitment in Man: Rank Deordering with Direction but Not with Speed of Voluntary Movement.” *Science* 214 (4523): 933–36. <https://doi.org/10.1126/science.7302570>.
- Desmedt, John E., and E. Godaux. 1977. “Fast Motor Units Are Not Preferentially Activated in Rapid Voluntary Contractions in Man.” *Nature* 267 (5613): 717–19. <https://doi.org/10.1038/267717a0>.
- Duchateau, Jacques, and Roger M. Enoka. 2008. “Neural Control of Shortening and Lengthening

- Contractions: Influence of Task Constraints: Shortening and Lengthening Contractions.” *The Journal of Physiology* 586 (24): 5853–64.
<https://doi.org/10.1113/jphysiol.2008.160747>.
- Duchateau, Jacques, and Roger M. Enoka. 2011. “Human Motor Unit Recordings: Origins and Insight into the Integrated Motor System.” *Brain Research* 1409 (August): 42–61.
<https://doi.org/10.1016/j.brainres.2011.06.011>.
- Dudman, Joshua T, and John W Krakauer. 2016. “The Basal Ganglia: From Motor Commands to the Control of Vigor.” *Current Opinion in Neurobiology* 37 (April): 158–66.
<https://doi.org/10.1016/j.conb.2016.02.005>.
- Elber-Dorozko, Lotem, and Yonatan Loewenstein. 2018. “Striatal Action-Value Neurons Reconsidered.” *ELife* 7 (May): e34248. <https://doi.org/10.7554/eLife.34248>.
- Farina, Dario. 2006. “Interpretation of the Surface Electromyogram in Dynamic Contractions.” *Exercise and Sport Sciences Reviews* 34 (3): 121–27. <https://doi.org/10.1249/00003677-200607000-00006>.
- Farina, Dario, Ivan Vujaklija, Massimo Sartori, Tamás Kapelner, Francesco Negro, Ning Jiang, Konstantin Bergmeister, Arash Andalib, Jose Principe, and Oskar C. Aszmann. 2017. “Man/Machine Interface Based on the Discharge Timings of Spinal Motor Neurons after Targeted Muscle Reinnervation.” *Nature Biomedical Engineering* 1 (2): 0025.
<https://doi.org/10.1038/s41551-016-0025>.
- Floresco, Stan B. 2015. “The Nucleus Accumbens: An Interface Between Cognition, Emotion, and Action.” *Annual Review of Psychology* 66 (1): 25–52.
<https://doi.org/10.1146/annurev-psych-010213-115159>.
- Formento, Emanuele, Paul Botros, and Jose M Carmena. 2021. “Skilled Independent Control of Individual Motor Units via a Non-Invasive Neuromuscular–Machine Interface.” *Journal of Neural Engineering* 18 (6): 066019. <https://doi.org/10.1088/1741-2552/ac35ac>.
- Fuglevand, A J, D A Winter, and A E Patla. 1993. “Models of Recruitment and Rate Coding Organization in Motor-Unit Pools.” *Journal of Neurophysiology* 70 (6): 2470–88.
<https://doi.org/10.1152/jn.1993.70.6.2470>.
- Ganguly, Karunesh, and Jose M. Carmena. 2009. “Emergence of a Stable Cortical Map for Neuroprosthetic Control.” Edited by James Ashe. *PLoS Biology* 7 (7): e1000153.
<https://doi.org/10.1371/journal.pbio.1000153>.
- Ganguly, Karunesh, Dragan F Dimitrov, Jonathan D Wallis, and Jose M Carmena. 2011. “Reversible Large-Scale Modification of Cortical Networks during Neuroprosthetic Control.” *Nature Neuroscience* 14 (5): 662–67. <https://doi.org/10.1038/nn.2797>.
- Gao, Peiran, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya Ganguli. 2017a. “A Theory of Multineuronal Dimensionality, Dynamics and Measurement.” *BioRxiv*, November, 214262. <https://doi.org/10.1101/214262>.
- Garnett, R, and J A Stephens. 1981. “Changes in the Recruitment Threshold of Motor Units Produced by Cutaneous Stimulation in Man.” *The Journal of Physiology* 311 (1): 463–73.
<https://doi.org/10.1113/jphysiol.1981.sp013598>.
- Graybiel, Ann M., and Scott T. Grafton. 2015. “The Striatum: Where Skills and Habits Meet.” *Cold Spring Harbor Perspectives in Biology* 7 (8): a021691.
<https://doi.org/10.1101/cshperspect.a021691>.
- Grimby, L., and J. Hannerz. 1970. “Differences in Recruitment Order of Motor Units in Phasic and Tonic Flexion Reflex in ‘spinal Man’.” *Journal of Neurology, Neurosurgery & Psychiatry* 33 (5): 562–70. <https://doi.org/10.1136/jnnp.33.5.562>.

- Grimby, L, and J Hannerz. 1977. "Firing Rate and Recruitment Order of Toe Extensor Motor Units in Different Modes of Voluntary Contraction." *The Journal of Physiology* 264 (3): 865–79. <https://doi.org/10.1113/jphysiol.1977.sp011699>.
- Groeneboom, Nicolaas E., Sharon C. Yates, Maja A. Puchades, and Jan G. Bjaalie. 2020. "Nutil: A Pre- and Post-Processing Toolbox for Histological Rodent Brain Section Images." *Frontiers in Neuroinformatics* 14 (August): 37. <https://doi.org/10.3389/fninf.2020.00037>.
- Grover, Pulkit, and Praveen Venkatesh. 2017. "An Information-Theoretic View of EEG Sensing." *Proceedings of the IEEE* 105 (2): 367–84. <https://doi.org/10.1109/JPROC.2016.2615179>.
- Gulati, Tanuj, Ling Guo, Dhakshin S Ramanathan, Anitha Bodepudi, and Karunesh Ganguly. 2017. "Neural Reactivations during Sleep Determine Network Credit Assignment." *Nature Neuroscience* 20 (9): 1277–84. <https://doi.org/10.1038/nn.4601>.
- Haar Romeny, B. M. ter, J. J. Denier van der Gon, and C. C.A.M. Gielen. 1982a. "Changes in Recruitment Order of Motor Units in the Human Biceps Muscle." *Experimental Neurology* 78 (2): 360–68. [https://doi.org/10.1016/0014-4886\(82\)90054-1](https://doi.org/10.1016/0014-4886(82)90054-1).
- Haar Romeny, B.M. ter, J.J. Denier van der Gon, and C.C.A.M. Gielen. 1984. "Relation between Location of a Motor Unit in the Human Biceps Brachii and Its Critical Firing Levels for Different Tasks." *Experimental Neurology* 85 (3): 631–50. [https://doi.org/10.1016/0014-4886\(84\)90036-0](https://doi.org/10.1016/0014-4886(84)90036-0).
- Haar Romeny, B.M. ter, J.J. Denier van der Gon, and C.C.A.M. Gielen. 1982b. "Changes in Recruitment Order of Motor Units in the Human Biceps Muscle." *Experimental Neurology* 78 (2): 360–68. [https://doi.org/10.1016/0014-4886\(82\)90054-1](https://doi.org/10.1016/0014-4886(82)90054-1).
- Hahne, Janne M., Meike A. Schweisfurth, Mario Koppe, and Dario Farina. 2018. "Simultaneous Control of Multiple Functions of Bionic Hand Prostheses: Performance and Robustness in End Users." *Science Robotics* 3 (19): eaat3630. <https://doi.org/10.1126/scirobotics.aat3630>.
- Hamid, Arif A., Michael J. Frank, and Christopher I. Moore. 2021. "Wave-like Dopamine Dynamics as a Mechanism for Spatiotemporal Credit Assignment." *Cell* 184 (10): 2733–2749.e16. <https://doi.org/10.1016/j.cell.2021.03.046>.
- Harris, Charles R., K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, et al. 2020. "Array Programming with NumPy." *Nature* 585 (7825): 357–62. <https://doi.org/10.1038/s41586-020-2649-2>.
- Harrison, V. F., and W. B. Koch. 1972. "Voluntary Control of Single Motor Unit Activity in the Extensor Digitorum Muscle." *Physical Therapy* 52 (3): 267–72. <https://doi.org/10.1093/ptj/52.3.267>.
- Harrison, Virginia F., and O. A. Mortensen. 1962. "Identification and Voluntary Control of Single Motor Unit Activity in the Tibialis Anterior Muscle." *The Anatomical Record* 144 (2): 109–16. <https://doi.org/10.1002/ar.1091440205>.
- Harwood, B., I. Choi, and C. L. Rice. 2012. "Reduced Motor Unit Discharge Rates of Maximal Velocity Dynamic Contractions in Response to a Submaximal Dynamic Fatigue Protocol." *Journal of Applied Physiology* 113 (12): 1821–30. <https://doi.org/10.1152/jappphysiol.00879.2012>.
- Henneman, E, H P Clamann, J D Gillies, and R D Skinner. 1974. "Rank Order of Motoneurons within a Pool: Law of Combination." *Journal of Neurophysiology* 37 (6): 1338–49. <https://doi.org/10.1152/jn.1974.37.6.1338>.
- Henneman, Elwood. 1957. "Relation between Size of Neurons and Their Susceptibility to

- Discharge.” *Science* 126 (3287): 1345–47.
<https://doi.org/10.1126/science.126.3287.1345>.
- Herrmann, Uta, and Martha Flanders. 1998. “Directional Tuning of Single Motor Units.” *Journal of Neuroscience* 18 (20): 8402–16. <https://doi.org/10.1523/JNEUROSCI.18-20-08402.1998>.
- Hochberg, Leigh R., Mijail D. Serruya, Gerhard M. Friehs, Jon A. Mukand, Maryam Saleh, Abraham H. Caplan, Almut Branner, David Chen, Richard D. Penn, and John P. Donoghue. 2006. “Neuronal Ensemble Control of Prosthetic Devices by a Human with Tetraplegia.” *Nature* 442 (7099): 164–71. <https://doi.org/10.1038/nature04970>.
- Hu, Sile, Davide Ciliberti, Andres D. Grosmark, Frédéric Michon, Daoyun Ji, Hector Penagos, György Buzsáki, Matthew A. Wilson, Fabian Kloosterman, and Zhe Chen. 2018. “Real-Time Readout of Large-Scale Unsorted Neural Ensemble Place Codes.” *Cell Reports* 25 (10): 2635–2642.e5. <https://doi.org/10.1016/j.celrep.2018.11.033>.
- Hunnicutt, Barbara J, Bart C Jongbloets, William T Birdsong, Katrina J Gertz, Haining Zhong, and Tianyi Mao. 2016. “A Comprehensive Excitatory Input Map of the Striatum Reveals Novel Functional Organization.” *ELife* 5 (November): e19103. <https://doi.org/10.7554/eLife.19103>.
- Illyés, S. 1977. “The Voluntary Control of Single Motor Unit Activity.” *IFAC Proceedings Volumes* 10 (15): 86–95. [https://doi.org/10.1016/S1474-6670\(17\)66409-7](https://doi.org/10.1016/S1474-6670(17)66409-7).
- Jarosiewicz, Beata, Anish A Sarma, Daniel Bacher, Nicolas Y Masse, John D Simeral, Brittany Sorice, Erin M Oakley, et al. 2015. “Virtual Typing by People with Tetraplegia Using a Self-Calibrating Intracortical Brain-Computer Interface.” *Science Translational Medicine* 7 (313): 313ra179. <https://doi.org/10.1126/scitranslmed.aac7328>.
- Jin, Xin, and Rui M. Costa. 2010. “Start/Stop Signals Emerge in Nigrostriatal Circuits during Sequence Learning.” *Nature* 466 (7305): 457–62. <https://doi.org/10.1038/nature09263>.
- Jin, Xin, Fatuel Tecuapetla, and Rui M Costa. 2014. “Basal Ganglia Subcircuits Distinctively Encode the Parsing and Concatenation of Action Sequences.” *Nature Neuroscience* 17 (3): 423–30. <https://doi.org/10.1038/nn.3632>.
- Jolly, Eshin. 2018. “Pymer4: Connecting R and Python for Linear Mixed Modeling.” *Journal of Open Source Software* 3 (31): 862. <https://doi.org/10.21105/joss.00862>.
- Jun, James J., Nicholas A. Steinmetz, Joshua H. Siegle, Daniel J. Denman, Marius Bauza, Brian Barbarits, Albert K. Lee, et al. 2017. “Fully Integrated Silicon Probes for High-Density Recording of Neural Activity.” *Nature* 551 (7679): 232–36. <https://doi.org/10.1038/nature24636>.
- Kandel, Eric R., ed. 2013. *Principles of Neural Science*. 5th ed. New York: McGraw-Hill.
- Kapelner, Tamás, Ivan Vujaklija, Ning Jiang, Francesco Negro, Oskar C. Aszmann, Jose Principe, and Dario Farina. 2019. “Predicting Wrist Kinematics from Motor Unit Discharge Timings for the Control of Active Prostheses.” *Journal of NeuroEngineering and Rehabilitation* 16 (1): 47. <https://doi.org/10.1186/s12984-019-0516-x>.
- Kernell, Daniel, and Hans Hultborn. 1990. “Synaptic Effects on Recruitment Gain: A Mechanism of Importance for the Input-Output Relations of Motoneurone Pools?” *Brain Research* 507 (1): 176–79. [https://doi.org/10.1016/0006-8993\(90\)90542-J](https://doi.org/10.1016/0006-8993(90)90542-J).
- Klaus, Andreas, Joaquim Alves da Silva, and Rui M. Costa. 2019. “What, If, and When to Move: Basal Ganglia Circuits and Self-Paced Action Initiation.” *Annual Review of Neuroscience* 42 (1): 459–83. <https://doi.org/10.1146/annurev-neuro-072116-031033>.
- Klaus, Andreas, Gabriela J. Martins, Vitor B. Paixao, Pengcheng Zhou, Liam Paninski, and Rui

- M. Costa. 2017. “The Spatiotemporal Organization of the Striatum Encodes Action Space.” *Neuron* 95 (5): 1171-1180.e7. <https://doi.org/10.1016/j.neuron.2017.08.015>.
- Koralek, Aaron C., Xin Jin, John D. Long II, Rui M. Costa, and Jose M. Carmena. 2012. “Corticostriatal Plasticity Is Necessary for Learning Intentional Neuroprosthetic Skills.” *Nature* 483 (7389): 331–35. <https://doi.org/10.1038/nature10845>.
- Krakauer, John W., and Reza Shadmehr. 2006. “Consolidation of Motor Memory.” *Trends in Neurosciences* 29 (1): 58–64. <https://doi.org/10.1016/j.tins.2005.10.003>.
- Kukulka, Carl G., and H. Peter Clamann. 1981. “Comparison of the Recruitment and Discharge Properties of Motor Units in Human Brachial Biceps and Adductor Pollicis during Isometric Contractions.” *Brain Research* 219 (1): 45–55. [https://doi.org/10.1016/0006-8993\(81\)90266-3](https://doi.org/10.1016/0006-8993(81)90266-3).
- Kupferschmidt, David A., Konrad Juczewski, Guohong Cui, Kari A. Johnson, and David M. Lovinger. 2017. “Parallel, but Dissociable, Processing in Discrete Corticostriatal Inputs Encodes Skill Learning.” *Neuron* 96 (2): 476-489.e5. <https://doi.org/10.1016/j.neuron.2017.09.040>.
- Lazega, Emmanuel, and Tom A. B. Snijders. n.d. *Multilevel Network Analysis for the Social Sciences: Theory, Methods and Applications*. Methodos Series 12.
- Lee, Daeyeol, Hyojung Seo, and Min Whan Jung. 2012. “Neural Basis of Reinforcement Learning and Decision Making.” *Annual Review of Neuroscience* 35 (1): 287–308. <https://doi.org/10.1146/annurev-neuro-062111-150512>.
- Lemke, Stefan M., Dhakshin S. Ramanathan, Ling Guo, Seok Joon Won, and Karunesh Ganguly. 2019. “Emergent Modular Neural Control Drives Coordinated Motor Actions.” *Nature Neuroscience* 22 (7): 1122–31. <https://doi.org/10.1038/s41593-019-0407-2>.
- Luo, Thomas Zhihao, Adrian Gopnik Bondy, Diksha Gupta, Verity Alexander Elliott, Charles D Kopec, and Carlos D Brody. 2020. “An Approach for Long-Term, Multi-Probe Neuropixels Recordings in Unrestrained Rats.” *ELife* 9 (October): e59716. <https://doi.org/10.7554/eLife.59716>.
- Manning, C. D., T. A. Miller, M. L. Burnham, C. D. Murnaghan, B. Calancie, and P. Bawa. 2010. “Recovery of Human Motoneurons during Rotation.” *Experimental Brain Research* 204 (1): 139–44. <https://doi.org/10.1007/s00221-010-2295-2>.
- Markowitz, Jeffrey E., Winthrop F. Gillis, Celia C. Beron, Shay Q. Neufeld, Keiramarie Robertson, Neha D. Bhagat, Ralph E. Peterson, et al. 2018. “The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection.” *Cell* 174 (1): 44-58.e17. <https://doi.org/10.1016/j.cell.2018.04.019>.
- Marshall, Najja J., Joshua I. Glaser, Eric M. Trautmann, Elom A. Amematsro, Sean M. Perkins, Michael N. Shadlen, L.F. Abbott, John P. Cunningham, and Mark M. Churchland. 2021. “Flexible Neural Control of Motor Units.” Preprint. Neuroscience. <https://doi.org/10.1101/2021.05.05.442653>.
- Martinez-Valdes, E., F. Negro, C. M. Laine, D. Falla, F. Mayer, and D. Farina. 2017. “Tracking Motor Units Longitudinally across Experimental Sessions with High-Density Surface Electromyography: Motor Unit Tracking with High-Density EMG.” *The Journal of Physiology* 595 (5): 1479–96. <https://doi.org/10.1113/JP273662>.
- McGinty, Vincent B., Sylvie Lardeux, Sharif A. Taha, James J. Kim, and Saleem M. Nicola. 2013. “Invigoration of Reward Seeking by Cue and Proximity Encoding in the Nucleus Accumbens.” *Neuron* 78 (5): 910–22. <https://doi.org/10.1016/j.neuron.2013.04.010>.
- Meer, Matthijs A.A. van der, Adam Johnson, Neil C. Schmitzer-Torbert, and A. David Redish.

2010. “Triple Dissociation of Information Processing in Dorsal Striatum, Ventral Striatum, and Hippocampus on a Learned Spatial Decision Task.” *Neuron* 67 (1): 25–32. <https://doi.org/10.1016/j.neuron.2010.06.023>.
- Mello, Gustavo B.M., Sofia Soares, and Joseph J. Paton. 2015. “A Scalable Population Code for Time in the Striatum.” *Current Biology* 25 (9): 1113–22. <https://doi.org/10.1016/j.cub.2015.02.036>.
- Millán, Jose del R., and Jose Carmena. 2010. “Invasive or Noninvasive: Understanding Brain-Machine Interface Technology [Conversations in BME.]” *IEEE Engineering in Medicine and Biology Magazine* 29 (1): 16–22. <https://doi.org/10.1109/MEMB.2009.935475>.
- Miller, Cory T., David Gire, Kim Hoke, Alexander C. Huk, Darcy Kelley, David A. Leopold, Matthew C. Smear, Frederic Theunissen, Michael Yartsev, and Christopher M. Niell. 2022. “Natural Behavior Is the Language of the Brain.” *Current Biology* 32 (10): R482–93. <https://doi.org/10.1016/j.cub.2022.03.031>.
- Milner-Brown, H. S., R. B. Stein, and R. Yemm. 1973. “The Orderly Recruitment of Human Motor Units during Voluntary Isometric Contractions.” *The Journal of Physiology* 230 (2): 359–70. <https://doi.org/10.1113/jphysiol.1973.sp010192>.
- Naito, Akira. 2004. “Electrophysiological Studies of Muscles in the Human Upper Limb: The Biceps Brachii.” *Anatomical Science International* 79 (1): 11. <https://doi.org/10.1111/j.1447-073x.2004.00064.x>.
- Nakanishi, Masaki, Yijun Wang, Xiaogang Chen, Yu-Te Wang, Xiaorong Gao, and Tzyy-Ping Jung. 2018. “Enhancing Detection of SSVEPs for a High-Speed Brain Speller Using Task-Related Component Analysis.” *IEEE Transactions on Biomedical Engineering* 65 (1): 104–12. <https://doi.org/10.1109/TBME.2017.2694818>.
- Nardone, A., C. Romanò, and M. Schieppati. 1989. “Selective Recruitment of High-Threshold Human Motor Units during Voluntary Isotonic Lengthening of Active Muscles.” *The Journal of Physiology* 409 (1): 451–71. <https://doi.org/10.1113/jphysiol.1989.sp017507>.
- Neely, Ryan M., Aaron C. Koralek, Vivek R. Athalye, Rui M. Costa, and Jose M. Carmena. 2018. “Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia.” *Neuron* 97 (6): 1356–1368.e4. <https://doi.org/10.1016/j.neuron.2018.01.051>.
- Negro, Francesco, Silvia Muceli, Anna Margherita Castronovo, Ales Holobar, and Dario Farina. 2016. “Multi-Channel Intramuscular and Surface EMG Decomposition by Convolutional Blind Source Separation.” *Journal of Neural Engineering* 13 (2): 026027. <https://doi.org/10.1088/1741-2560/13/2/026027>.
- Nicola, S. M. 2010. “The Flexible Approach Hypothesis: Unification of Effort and Cue-Responding Hypotheses for the Role of Nucleus Accumbens Dopamine in the Activation of Reward-Seeking Behavior.” *Journal of Neuroscience* 30 (49): 16585–600. <https://doi.org/10.1523/JNEUROSCI.3958-10.2010>.
- Nise, Norman S. 2015. *Control Systems Engineering*. Seventh edition. Hoboken, NJ: Wiley.
- Nuyujukian, P., J. M. Fan, J. C. Kao, S. I. Ryu, and K. V. Shenoy. 2015. “A High-Performance Keyboard Neural Prosthesis Enabled by Task Optimization.” *IEEE Transactions on Biomedical Engineering* 62 (1): 21–29. <https://doi.org/10.1109/TBME.2014.2354697>.
- Nuyujukian, Paul, Jose Albites Sanabria, Jad Saab, Chethan Pandarinath, Beata Jarosiewicz, Christine H. Blabe, Brian Franco, et al. 2018. “Cortical Control of a Tablet Computer by People with Paralysis.” Edited by Dingguo Zhang. *PLOS ONE* 13 (11): e0204566. <https://doi.org/10.1371/journal.pone.0204566>.
- O’Doherty, J. 2004. “Dissociable Roles of Ventral and Dorsal Striatum in Instrumental

- Conditioning.” *Science* 304 (5669): 452–54. <https://doi.org/10.1126/science.1094285>.
- Orsborn, Amy L, and Bijan Pesaran. 2017. “Parsing Learning in Networks Using Brain–Machine Interfaces.” *Current Opinion in Neurobiology* 46 (October): 76–83. <https://doi.org/10.1016/j.conb.2017.08.002>.
- Orsborn, Amy L., Helene G. Moorman, Simon A. Overduin, Maryam M. Shanechi, Dragan F. Dimitrov, and Jose M. Carmena. 2014. “Closed-Loop Decoder Adaptation Shapes Neural Plasticity for Skillful Neuroprosthetic Control.” *Neuron* 82 (6): 1380–93. <https://doi.org/10.1016/j.neuron.2014.04.048>.
- Ota, Keisuke, Yasuhiro Oisi, Takayuki Suzuki, Muneki Ikeda, Yoshiki Ito, Tsubasa Ito, Hiroyuki Uwamori, et al. 2021. “Fast, Cell-Resolution, Contiguous-Wide Two-Photon Imaging to Reveal Functional Network Architectures across Multi-Modal Cortical Areas.” *Neuron* 109 (11): 1810-1824.e9. <https://doi.org/10.1016/j.neuron.2021.03.032>.
- Pak, Nikita, Joshua H. Siegle, Justin P. Kinney, Daniel J. Denman, Timothy J. Blanche, and Edward S. Boyden. 2015. “Closed-Loop, Ultraprecise, Automated Craniotomies.” *Journal of Neurophysiology* 113 (10): 3943–53. <https://doi.org/10.1152/jn.01055.2014>.
- Pandarinath, Chethan, Paul Nuyujukian, Christine H. Blabe, Brittany L. Sorice, Jad Saab, Francis R. Willett, Leigh R. Hochberg, Krishna V. Shenoy, and Jaimie M. Henderson. 2017. “High Performance Communication by People with Paralysis Using an Intracortical Brain-Computer Interface.” *eLife* 6: 1–27. <https://doi.org/10.7554/eLife.18554>.
- Papp, Eszter A., Trygve B. Leergaard, Evan Calabrese, G. Allan Johnson, and Jan G. Bjaalie. 2014. “Waxholm Space Atlas of the Sprague Dawley Rat Brain.” *NeuroImage* 97 (August): 374–86. <https://doi.org/10.1016/j.neuroimage.2014.04.001>.
- Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, et al. 2018. “Scikit-Learn: Machine Learning in Python.” arXiv. <http://arxiv.org/abs/1201.0490>.
- Pégar, Nicolas C., Alan R. Mardinly, Ian Antón Oldenburg, Savitha Sridharan, Laura Waller, and Hillel Adesnik. 2017. “Three-Dimensional Scanless Holographic Optogenetics with Temporal Focusing (3D-SHOT).” *Nature Communications* 8 (1): 1228. <https://doi.org/10.1038/s41467-017-01031-3>.
- Peters, Andrew J., Simon X. Chen, and Takaki Komiyama. 2014. “Emergence of Reproducible Spatiotemporal Activity during Motor Learning.” *Nature* 510 (7504): 263–67. <https://doi.org/10.1038/nature13235>.
- Peters, Andrew J., Julie M. J. Fabre, Nicholas A. Steinmetz, Kenneth D. Harris, and Matteo Carandini. 2021. “Striatal Activity Topographically Reflects Cortical Activity.” *Nature* 591 (7850): 420–25. <https://doi.org/10.1038/s41586-020-03166-8>.
- Pierella, Camilla, Elisa Galofaro, Alice De Luca, Luca Losio, Simona Gamba, Antonino Massone, Ferdinando A. Mussa-Ivaldi, and Maura Casadio. 2021. “Recovery of Distal Arm Movements in Spinal Cord Injured Patients with a Body-Machine Interface: A Proof-of-Concept Study.” *Sensors* 21 (6): 2243. <https://doi.org/10.3390/s21062243>.
- Puchades, Maja A., Gergely Csucs, Debora Ledergerber, Trygve B. Leergaard, and Jan G. Bjaalie. 2019. “Spatial Registration of Serial Microscopic Brain Images to Three-Dimensional Reference Atlases with the QuickNII Tool.” Edited by Manuel S. Malmierca. *PLOS ONE* 14 (5): e0216796. <https://doi.org/10.1371/journal.pone.0216796>.
- Radhakrishnan, Saritha M., Stuart N. Baker, and Andrew Jackson. 2008. “Learning a Novel Myoelectric-Controlled Interface Task.” *Journal of Neurophysiology* 100 (4): 2397–2408. <https://doi.org/10.1152/jn.90614.2008>.

- Reback, Jeff, Jbrockmendel, Wes McKinney, Joris Van Den Bossche, Tom Augspurger, Phillip Cloud, Simon Hawkins, et al. 2021. *Pandas-Dev/Pandas: Pandas 1.3.3* (version v1.3.3). Zenodo. <https://doi.org/10.5281/ZENODO.5501881>.
- Recanatesi, Stefano, Gabriel Koch Ocker, Michael A. Buice, and Eric Shea-Brown. 2019. “Dimensionality in Recurrent Spiking Networks: Global Trends in Activity and Local Origins in Connectivity.” *PLOS Computational Biology* 15 (7): e1006446. <https://doi.org/10.1371/journal.pcbi.1006446>.
- Richards, Blake A., Timothy P. Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, et al. 2019. “A Deep Learning Framework for Neuroscience.” *Nature Neuroscience* 22 (11): 1761–70. <https://doi.org/10.1038/s41593-019-0520-2>.
- Rick, Jochen. 2010. “Performance Optimizations of Virtual Keyboards for Stroke-Based Text Entry on a Touch-Based Tabletop.” In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*, 77–86. UIST '10. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/1866029.1866043>.
- Riek, S., and P. Bawa. 1992. “Recruitment of Motor Units in Human Forearm Extensors.” *Journal of Neurophysiology* 68 (1): 100–108. <https://doi.org/10.1152/jn.1992.68.1.100>.
- Rossant, Cyrille, Shabnam N Kadir, Dan F M Goodman, John Schulman, Maximilian L D Hunter, Aman B Saleem, Andres Grosmark, et al. 2016. “Spike Sorting for Large, Dense Electrode Arrays.” *Nature Neuroscience* 19 (4): 634–41. <https://doi.org/10.1038/nn.4268>.
- Rothenhoefer, Kathryn M., Vincent D. Costa, Ramón Bartolo, Raquel Vicario-Feliciano, Elisabeth A. Murray, and Bruno B. Averbeck. 2017. “Effects of Ventral Striatum Lesions on Stimulus-Based versus Action-Based Reinforcement Learning.” *The Journal of Neuroscience* 37 (29): 6902–14. <https://doi.org/10.1523/JNEUROSCI.0631-17.2017>.
- Sadtler, Patrick T., Kristin M. Quick, Matthew D. Golub, Steven M. Chase, Stephen I. Ryu, Elizabeth C. Tyler-Kabara, Byron M. Yu, and Aaron P. Batista. 2014. “Neural Constraints on Learning.” *Nature* 512 (7515): 423–26. <https://doi.org/10.1038/nature13665>.
- Sales-Carbonell, Carola, Wahiba Taouali, Loubna Khalki, Matthieu O. Pasquet, Ludovic F. Petit, Typhaine Moreau, Pavel E. Rueda-Orozco, and David Robbe. 2018. “No Discrete Start/Stop Signals in the Dorsal Striatum of Mice Performing a Learned Action.” *Current Biology* 28 (19): 3044-3055.e5. <https://doi.org/10.1016/j.cub.2018.07.038>.
- Schieltz, Holger, Niels J. Dingemans, Shinichi Nakagawa, David F. Westneat, Hassen Allegue, Céline Teplitsky, Denis Réale, Ned A. Dochtermann, László Zsolt Garamszegi, and Yimen G. Araya-Ajoy. 2020. “Robustness of Linear Mixed-effects Models to Violations of Distributional Assumptions.” Edited by Chris Sutherland. *Methods in Ecology and Evolution* 11 (9): 1141–52. <https://doi.org/10.1111/2041-210X.13434>.
- Schmalfuß, L., R. Rupp, M.R. Tuga, A. Kogut, M. Hewitt, J. Meincke, F. Klinker, et al. 2015. “Steer by Ear: Myoelectric Auricular Control of Powered Wheelchairs for Individuals with Spinal Cord Injury.” *Restorative Neurology and Neuroscience* 34 (1): 79–95. <https://doi.org/10.3233/RNN-150579>.
- Schmitzer-Torbert, N.C., and A.D. Redish. 2008. “Task-Dependent Encoding of Space and Events by Striatal Neurons Is Dependent on Neural Subtype.” *Neuroscience* 153 (2): 349–60. <https://doi.org/10.1016/j.neuroscience.2008.01.081>.
- Schultz, Wolfram. 2019. “Recent Advances in Understanding the Role of Phasic Dopamine Activity.” *F1000Research* 8 (September): 1680.

- <https://doi.org/10.12688/f1000research.19793.1>.
- Schultz, Wolfram, Peter Dayan, and P. Read Montague. 1997. “A Neural Substrate of Prediction and Reward.” *Science* 275 (5306): 1593–99. <https://doi.org/10.1126/science.275.5306.1593>.
- Schultz, Wolfram, Peter Dayan, and P Read Montague. n.d. “A Neural Substrate of Prediction and Reward,” 7.
- Segal, Richard L. 1992. “Neuromuscular Compartments in the Human Biceps Brachii Muscle.” *Neuroscience Letters* 140 (1): 98–102. [https://doi.org/10.1016/0304-3940\(92\)90691-Y](https://doi.org/10.1016/0304-3940(92)90691-Y).
- Sellers, Eric W., Theresa M. Vaughan, and Jonathan R. Wolpaw. 2010. “A Brain-Computer Interface for Long-Term Independent Home Use.” *Amyotrophic Lateral Sclerosis* 11 (5): 449–55. <https://doi.org/10.3109/17482961003777470>.
- Shanechi, Maryam M. 2019. “Brain–Machine Interfaces from Motor to Mood.” *Nature Neuroscience* 22 (10): 1554–64. <https://doi.org/10.1038/s41593-019-0488-y>.
- Shin, Eun Ju, Yunsil Jang, Soyoun Kim, Hoseok Kim, Xinying Cai, Hyunjung Lee, Jung Hoon Sul, et al. 2021. “Robust and Distributed Neural Representation of Action Values.” *ELife* 10 (April): e53045. <https://doi.org/10.7554/eLife.53045>.
- Siegle, Joshua H, Aarón Cuevas López, Yogi A Patel, Kirill Abramov, Shay Ohayon, and Jakob Voigts. 2017. “Open Ephys: An Open-Source, Plugin-Based Platform for Multichannel Electrophysiology.” *Journal of Neural Engineering* 14 (4): 045003. <https://doi.org/10.1088/1741-2552/aa5eea>.
- Staudenmann, Didier, and Wolfgang Taube. 2015. “Brachialis Muscle Activity Can Be Assessed with Surface Electromyography.” *Journal of Electromyography and Kinesiology* 25 (2): 199–204. <https://doi.org/10.1016/j.jelekin.2014.11.003>.
- Steinmetz, Nicholas A., Cagatay Aydin, Anna Lebedeva, Michael Okun, Marius Pachitariu, Marius Bauza, Maxime Beau, et al. 2021. “Neuropixels 2.0: A Miniaturized High-Density Probe for Stable, Long-Term Brain Recordings.” *Science* 372 (6539): eabf4588. <https://doi.org/10.1126/science.abf4588>.
- Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. Second edition. Adaptive Computation and Machine Learning Series. Cambridge, Massachusetts: The MIT Press.
- Thomas, C. K., B. H. Ross, and R. B. Stein. 1986. “Motor-Unit Recruitment in Human First Dorsal Interosseous Muscle for Static Contractions in Three Different Directions.” *Journal of Neurophysiology* 55 (5): 1017–29. <https://doi.org/10.1152/jn.1986.55.5.1017>.
- Thorn, Catherine A., Hisham Atallah, Mark Howe, and Ann M. Graybiel. 2010. “Differential Dynamics of Activity Changes in Dorsolateral and Dorsomedial Striatal Loops during Learning.” *Neuron* 66 (5): 781–95. <https://doi.org/10.1016/j.neuron.2010.04.036>.
- Ting, Jordyn E., Alessandro Del Vecchio, Devapratim Sarma, Nikhil Verma, Samuel C. Colachis, Nicholas V. Annetta, Jennifer L. Collinger, Dario Farina, and Douglas J. Weber. 2021. “Sensing and Decoding the Neural Drive to Paralyzed Muscles during Attempted Movements of a Person with Tetraplegia Using a Sleeve Array.” *Journal of Neurophysiology* 126 (6): 2104–18. <https://doi.org/10.1152/jn.00220.2021>.
- Ting, Jordyn E., Alessandro Del Vecchio, Devapratim Sarma, Samuel C. Colachis, Nicholas V. Annetta, Jennifer L. Collinger, Dario Farina, and Douglas J. Weber. 2021. “Sensing and Decoding the Neural Drive to Paralyzed Muscles during Attempted Movements of a Person with Tetraplegia Using a Sleeve Array.” Preprint. Neurology. <https://doi.org/10.1101/2021.02.24.21250962>.

- Tonin, Luca, and José del R. Millán. 2021. “Noninvasive Brain–Machine Interfaces for Robotic Devices.” *Annual Review of Control, Robotics, and Autonomous Systems* 4 (1): annurev-control-012720-093904. <https://doi.org/10.1146/annurev-control-012720-093904>.
- Tuthill, John C., and Eiman Azim. 2018. “Proprioception.” *Current Biology* 28 (5): R194–203. <https://doi.org/10.1016/j.cub.2018.01.064>.
- Veuthey, T. L., K. Derosier, S. Kondapavulur, and K. Ganguly. 2020. “Single-Trial Cross-Area Neural Population Dynamics during Long-Term Skill Learning.” *Nature Communications* 11 (1): 4057. <https://doi.org/10.1038/s41467-020-17902-1>.
- Virtanen, Pauli, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, et al. 2020. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python.” *Nature Methods* 17 (3): 261–72. <https://doi.org/10.1038/s41592-019-0686-2>.
- Wakeling, James M. 2009. “The Recruitment of Different Compartments within a Muscle Depends on the Mechanics of the Movement.” *Biology Letters* 5 (1): 30–34. <https://doi.org/10.1098/rsbl.2008.0459>.
- Wang, Wei, Jennifer L. Collinger, Alan D. Degenhart, Elizabeth C. Tyler-Kabara, Andrew B. Schwartz, Daniel W. Moran, Douglas J. Weber, et al. 2013. “An Electrographic Brain Interface in an Individual with Tetraplegia.” Edited by Shawn Hochman. *PLoS ONE* 8 (2): e55344. <https://doi.org/10.1371/journal.pone.0055344>.
- Willett, Francis R., Donald T. Avansino, Leigh R. Hochberg, Jaimie M. Henderson, and Krishna V. Shenoy. 2021. “High-Performance Brain-to-Text Communication via Handwriting.” *Nature* 593 (7858): 249–54. <https://doi.org/10.1038/s41586-021-03506-2>.
- Williams, Alex H., Ben Poole, Niru Maheswaranathan, Ashesh K. Dhawale, Tucker Fisher, Christopher D. Wilson, David H. Brann, et al. 2020. “Discovering Precise Temporal Patterns in Large-Scale Neural Recordings through Robust and Interpretable Time Warping.” *Neuron* 105 (2): 246–259.e8. <https://doi.org/10.1016/j.neuron.2019.10.020>.
- Wolpaw, J. R., and D. J. McFarland. 2004. “Control of a Two-Dimensional Movement Signal by a Noninvasive Brain-Computer Interface in Humans.” *Proceedings of the National Academy of Sciences* 101 (51): 17849–54. <https://doi.org/10.1073/pnas.0403504101>.
- Yagishita, Sho, Akiko Hayashi-Takagi, Graham C.R. Ellis-Davies, Hidetoshi Urakubo, Shin Ishii, and Haruo Kasai. 2014. “A Critical Time Window for Dopamine Actions on the Structural Plasticity of Dendritic Spines.” *Science* 345 (6204): 1616–20. <https://doi.org/10.1126/science.1255514>.
- Yin, Henry H, Shweta Prasad Mulcare, Monica R F Hilário, Emily Clouse, Terrell Holloway, Margaret I Davis, Anita C Hansson, David M Lovinger, and Rui M Costa. 2009. “Dynamic Reorganization of Striatal Circuits during the Acquisition and Consolidation of a Skill.” *Nature Neuroscience* 12 (3): 333–41. <https://doi.org/10.1038/nn.2261>.
- Yoshida, Keitaro, Iku Tsutsui-Kimura, Anna Kono, Akihiro Yamanaka, Kenta Kobayashi, Masahiko Watanabe, Masaru Mimura, and Kenji F. Tanaka. 2020. “Opposing Ventral Striatal Medium Spiny Neuron Activities Shaped by Striatal Parvalbumin-Expressing Interneurons during Goal-Directed Behaviors.” *Cell Reports* 31 (13): 107829. <https://doi.org/10.1016/j.celrep.2020.107829>.
- Zhuang, Katie Z., Nicolas Sommer, Vincent Mendez, Saurav Aryan, Emanuele Formento, Edoardo D’Anna, Fiorenzo Artoni, et al. 2019. “Shared Human–Robot Proportional Control of a Dexterous Myoelectric Prosthesis.” *Nature Machine Intelligence* 1 (9): 400–411. <https://doi.org/10.1038/s42256-019-0093-5>.

Zuylen, E. J. van, C. C. Gielen, and J. J. Denier van der Gon. 1988. "Coordination and Inhomogeneous Activation of Human Arm Muscles during Isometric Torques." *Journal of Neurophysiology* 60 (5): 1523–48. <https://doi.org/10.1152/jn.1988.60.5.1523>.