NORTHWESTERN UNIVERSITY

Deciphering the Geometry of Primary Motor Cortical Manifolds: Observations from Naturalistic Movements and Implications for Intracortical Brain-Computer Interfaces

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Biomedical Engineering

By

Ege Altan

EVANSTON, ILLINOIS

September 2023

Abstract

Each neuron in the primary motor cortex (M1) is like a musician in an orchestra, contributing to a larger harmony under the constraint of a "neural manifold"—a geometric score describing the correlated signals produced by the neural musicians that drive movement. Despite the widespread recognition of the importance of M1 manifolds, the connection between their geometry and motor behavior remains unclear. Is the simple geometry a fundamental characteristic of M1, or is it an artifact resulting from the constrained laboratory tasks typically used in these investigations? Answering this question may enhance our understanding of how M1 orchestrates various movements. It also has practical implications for intracortical brain-computer interfaces (iBCIs), tools that translate neural activity into control of external devices for individuals with paralysis. My first study focused on understanding the geometric properties of neural manifolds using simulated neural signals. In my second study, I used this methodology to compare the geometry of M1 manifolds of monkeys during naturalistic, unconstrained behaviors to that during laboratory-based behaviors. Surprisingly, the geometry of M1 manifolds was only slightly more complex during naturalistic behaviors. Finally, I used M1 manifolds to improve the "decoders" of iBCIs, the brain-to-behavior maps that typically allow the control of external devices, to instead enable the control of the user's own muscles. I used two strategies based on the neural manifold concept: first, a decoder that mapped neural signals in the M1 manifold of an individual with quadriplegia directly to the muscle activity of a monkey, and second, adapting a decoder built solely from monkey data for human use. My findings 1) show that the simple geometry of M1 manifolds is not just an artifact of laboratory constraints but rather a computational strategy that provides a harmonized tune for reliable movement control across tasks and 2) illuminate the potential of iBCIs to empower individuals with paralysis to regain control of their own muscles, even on the real-world stage outside the laboratory.

ACKNOWLEDGMENTS

In the complex maze of possible scientific pursuits, I have been fortunate to discover my unique niche that aligns with my aspirations. I am deeply grateful to many remarkable individuals who have enriched my journey with their boundless dimensions of support, guidance, and love.

I want to express my sincere gratitude to my thesis committee members, Dr. Sara Solla, Dr. Jorge Nocedal, and for a brief period, Dr. Sandro Mussa-Ivaldi. They have been my intellectual guiding stars. Dr. Solla, in particular, has been more than a committee member; she has been a close collaborator, a third advisor, and a role model. Her enthusiasm for every topic we worked on, her knack for distilling complex ideas, and her constructive criticism have been invaluable. I aspire to emulate her scientific acumen and pedagogical prowess. With their unique perspectives, Dr. Jorge Nocedal and Dr. Sandro Mussa-Ivaldi have consistently inspired me to challenge conventional thinking and explore novel approaches.

I want to thank my doctoral advisors, Dr. Lee Miller and Dr. Eric Perreault, who have been the bedrock of my academic journey together. Their unwavering support, faith in me, and commitment to fostering my growth have undoubtedly made me a better researcher and communicator.

I am grateful to my colleagues in the Miller Lab and the Neuromuscular Control Lab for their collaboration, feedback, and friendship. A special mention goes to Dr. Ali Farshchian, whose journey through a cold Chicago winter day to recruit me to the Miller Lab was pivotal. I also want to thank Dr. Fabio Rizzoglio, who worked closely with me on the third study, and Dr. Xuan Ma, who has been a mentor and reliable shoulder to lean on.

I want to extend my heartfelt thanks to my friends and family for their unwavering support. I am grateful to my mom, dad, and brother for supporting me even from overseas. In addition, the

friendship of Dr. Vatsala Goyal and soon-to-be Drs. Thomas Plaisier and Drew Beauchamp has been a source of strength.

I cannot express my gratitude enough to my best friend and partner for life, Becca. Her radiating smile has been my sunshine on cloudy days, and her positivity my anchor in stormy seas. She has been my rock, safe harbor, and most significant source of inspiration.

And finally, to my cat Leroy, who has been my most reliable study partner, especially during the isolating days of the pandemic. Thank you for your companionship and comfort.

List of Abbreviations

ADAM	Adaptive moment
AE	Autoencoder
ALS	Amyotrophic lateral sclerosis
ANOVA	Analysis of variance
APB	Abductor pollicis brevis
ASIA	American Spinal Cord Injury Association
BRAIN	Brain Research through Advancing Innovative Neurotechnologies
CCA	Canonical correlations analysis
CD	Correlation dimension
DANCo	Dimensionality from angle and norm concentration
DI	Dorsal interossei
ECRb	Extensor carpi radialis brevis
ECRl	Extensor carpi radialis longus
ECU	Extensor carpi ulnaris
EDCr	Extensor digitorum communis radialis
EMG	Electromyogram
FA	Factor analysis
FCR	Flexor carpi radialis
FCU	Flexor carpi ulnaris
FDA	Food and Drug Administration
FDI	Flexor digitorum indicis
FDP	Flexor digitorum profundus
FDS	Flexor digitorum superficialis
FES	Functional electrical stimulation
FSA	Fisher separability analysis
FSR	Force sensitive resistor
GPFA	Gaussian-process factor analysis
HMM	Hidden Markov model

IACUC	Institutional Animal Care and Use Committee
iBCI	Intracortical brain-computer interface
JAE	Joint autoencoder
LBMLE	Levina-Bickel maximum likelihood estimation
LDS	Linear dynamical system
LFADS	Latent factor analysis through dynamical systems
LLE	Locally linear embedding
LSTM	Long short-term memory
M1	Primary motor cortex
MLE	Maximum likelihood estimation
MSE	Mean squared error
PA	Parallel analysis
PCA	Principal component analysis
PD	Preferred direction
PFC	Prefrontal cortex
PMd	Dorsal premotor cortex
PR	Participation ratio
PT	Pronator teres
RMS	Root mean square
RNN	Recurrent neural network
SCI	Spinal cord injury
SNR	Signal-to-noise ratio
SUP	Supinator
TNN	Two nearest neighbors
UMAP	Uniform manifold approximation and projection
V1	Primary visual cortex
V2	Secondary visual cortex
VAF	Variance accounted for
WHO	World Health Organization

DEDICATION

I dedicate this thesis to monkeys that contributed to my work and the countless other research animals that have participated in scientific studies throughout history. Their sacrifices have expanded the boundaries of human health and advanced scientific knowledge, for which we must be grateful. It is my sincerest hope that the ongoing progress of science and technology will pave the path toward compassionate and ethical alternatives that eliminate the need for these sacrifices.

TABLE OF CONTENTS

List of Figures

Main Figures

Supplementary Figures

List of Tables

CHAPTER 1: INTRODUCTION

This chapter will begin with the clinical basis for my interest in understanding the function of the primary motor cortex (M1). Then, I will set the stage for the investigations that I have carried out with a brief overview of the anatomy and physiology of M1 and insights gained from landmark studies on its function in driving movement. While these studies have revealed much of what we know about the function of the primary motor cortex, for many years they were restricted largely to single-neuron analyses, which have limitations in understanding the function of populations of neurons. Next, I will discuss the advancements in neural recording technology and the subsequent paradigm shift from a single-neuron to a population-level view of cortical activity in the form of “neural manifolds,” a geometric representation of the collective activity of neurons. I will share examples from recent computational neuroscience theories on how the geometry of these neural manifolds, namely their dimensionality and nonlinearity, may be indicative of computational processing within and across brain areas. I will evaluate the benefits and drawbacks of tools used for computing the manifolds and their geometry. Finally, I will provide a summary that illustrates the knowledge gaps that the central chapters of this thesis aim to fill.

Clinical significance of relating primary motor cortical (M1) activity to hand function

Humans have incredible hand dexterity; from grabbing a cup of coffee to playing the guitar, we use our hands to interact intricately with our environment. Disease and accidents that result in loss of hand function, such as spinal cord injury, have devastating consequences. According to the World Health Organization (WHO), between 250,000 and 500,000 people suffer a spinal cord injury yearly (World Health Organization 2006, 2013). Restoration of hand function is of the utmost importance for individuals with quadriplegia, as highlighted by a national survey (Anderson 2004). Nearly 80% of individuals with quadriplegia would consent to brain surgery to regain some use of their hands (Blabe et al. 2015). Designing a technology capable of restoring

biomimetic control to all 27 degrees of freedom of the human hand is challenging, necessitating both novel approaches and an improved conceptual understanding of the brain's control of movement.

Functional electrical stimulation (FES)

One approach that helps restore function to individuals with paralysis is called functional electrical stimulation (FES). FES is a method employed to electrically activate paralyzed muscles, allowing individuals with paralysis to regain essential functions like bladder and respiratory control and even the recovery of voluntary movement (Lynch and Popovic 2008). FES systems deliver electrical pulses that contract muscles, crudely mimicking the physiological process of muscle contraction in non-disabled individuals (Bajd and Munih 2010; Ibitoye et al. 2016). FES can be used to stimulate the nerves or, in cases of denervation, the individual muscle fibers. However, when stimulating the muscles directly, a significantly larger electrical amplitude is needed to contract the muscles and the recruitment order of motor units gets reversed, causing the non-optimal recruitment of the fatigable muscle fibers over fatigue-resistant fibers (Bickel, Gregory, and Dean 2011; Gobbo et al. 2014).

The first FDA-approved FES system for grasp is the Freehand, which helps restore movement to individuals with C5-C6 level spinal cord injury (Hobby, Taylor, and Esnouf 2001). Patients use their shoulders to control the system, as the shoulder muscles remain largely unaffected at this injury level. Sensors across both shoulders together enable movements like lateral pinch and palmar grasp. Similar FES systems have been developed that enable forearm and elbow movements as well, including the Implantable Stimulator-Telemeter-12 (Kilgore et al. 2008), ReGrasp (Prochazka 2017), Handmaster (Venugopalan et al. 2015), and Stimgrasp (Barelli, Avelino, and Castro 2022).

The several FES systems for grasping have their limitations. These systems typically offer only a small number of movements. Further, they do not facilitate individuated finger movements, which may be attributed to challenges in stimulating the small intrinsic muscles of the hand. Many FES systems employ a basic set of control signals that activates a preprogrammed set of stimulation commands that generate a desired movement. Such systems lack the sophistication needed to activate the muscles related to finger movement individually. Moreover, FES systems feature unintuitive control methods, such as shoulder shrugging to activate the hand, imposing a cognitive burden on users. Electrical stimulation that drives the muscles somewhat mimics the underlying physiologic process, but voluntary control does not follow the underlying cognitive process. Such unintuitive control systems are impractical for individuals with high-level spinal cord injury who may lack any residual limb or neck movement to operate them.

Intracortical brain-computer interfaces (iBCIs)

A promising alternative that aims to restore movement to individuals with paralysis is the intracortical brain-computer interface (iBCI) paradigm. iBCIs capture signals from the motor areas of the brain, process them and convert them into instructions for output devices to help execute desired tasks. The primary objective of iBCIs is to provide a functional replacement. These functions could involve operating computers and robots or even restoring movement by completely bypassing the natural neuromuscular pathway from the brain to the muscles. Central to these iBCIs is the “neural decoder,” an algorithm that links neural activity to external movement variables.

Through animal and human research, iBCIs have made significant progress over the last few decades. Commercial neural interface enterprises like Neuralink, Paradromics, and Synchron attempt to bring this promising yet currently lab-bound technology to the market (Musk and Neuralink 2019; Obaid et al. 2020; Opie 2021). The overwhelming majority of both research-based and commercial iBCIs typically control the kinematics of stereotypic tasks (Hochberg et

al. 2006, 2012; Collinger et al. 2013; Wodlinger et al. 2015; Andersen, Aflalo, and Kellis 2019; Rastogi et al. 2020). In these models, researchers employ “observation-based” neural decoders, where the training data usually consists of motor cortical signals from users observing and attempting to imitate endpoint trajectories of a cursor or a robot.

Unfortunately, this observation-based decoding is not feasible for biomimetic iBCIs that decode force or muscle activity, simply because these variables cannot be observed. This is a significant limitation for iBCIs—the iBCI performance must be enhanced to approach the dependability of natural muscle-based function (Shih, Krusienski, and Wolpaw 2012). To address this limitation, researchers are exploring the possibility of restoring control of a user’s own limb by combining FES and iBCIs. In this model, users can reanimate their limbs with thoughts by using FES on the paralyzed muscles (Popović 2003; Bryden et al. 2005; Bouton et al. 2016; Ajiboye et al. 2017). However, only one proof-of-concept study has used neural decoders that predict muscle activation as input to FES for limb reanimation (Ethier et al. 2012). In this study, the neural decoder was computed between neural signals from the primary motor cortex and the muscle activity of an able-bodied monkey performing the ball grasp task. Then, the researchers temporarily chemically blocked the neural pathway between the brain and muscles, mimicking the paralysis that results from spinal cord injury. The monkey was able to operate the prebuilt neural decoder, which produced muscle-like predictions that were fed into the FES system, allowing the monkey to perform the ball grasp task with his thoughts successfully.

For individuals with paralysis with no motor output, building biomimetic, muscle-based decoders presents a challenge. Unlike observable kinematic trajectories on a computer screen, variables such as muscle activation and force cannot be observed. However, despite the lack of observable motor output, relevant information concerning muscle activation may exist within the primary motor cortex even during attempted movements. By exploring this connection, we can,

in theory, build muscle-based decoders. Therefore, it is critical to investigate the relationship between M1 activity and muscle activity resulting in hand movement even for paralyzed individuals, as it holds the potential to dramatically improve the quality of life for millions of individuals living with paralysis.

In Chapter 4, I propose two methods for building muscle-based iBCIs for individuals with paralysis, the output of which can be used as control commands for FES systems that aim at restoring muscle function. This system relies on the computation of the “neural manifold,” for which I make extensive use of the tools I developed and tested in Chapters 2 and 3.

Primary motor cortex (M1) anatomy and physiology

Despite extensive investigation, a definitive understanding of how the motor cortex controls movement remains unclear. Multiple theories have emerged over many decades of motor neuroscience research, each proposing distinct ideas about the relationship between motor cortical neural activity and movement. These theories have predominantly been based on single-neuron recordings from M1, focusing on specific features or components of a given movement. First, I will review the anatomy of the primary motor cortex and highlight examples from landmark studies that have shaped our understanding of its physiology.

Brief anatomy of M1

The primary motor cortex (M1), or Brodmann Area 4 (Brodmann 1909), controls muscle contractions and movement. M1 was first identified by its low threshold for eliciting movements through electrical stimulation (Fritsch 1870; Ferrier 1873). Located along the precentral gyrus and extending into the central sulcus, M1 is the primary motor output of the cerebral cortex to the spinal cord, with direct projections to spinal circuits via the pyramidal tracts. The neurons that synapse to spinal circuits are known as the corticospinal neurons (Rathelot and Strick 2009). Most corticospinal neurons synapse on spinal interneurons and have an indirect effect on

the spinal motoneurons that ultimately drive muscles that generate force and movement. Among them are the largest neurons in the central nervous system, which can reach as large as 100 μm in diameter (Purves et al. 2008). These large corticospinal neurons, also known as Betz cells, can directly send their axons to anterior horn cells in the spinal cord, synapsing to motoneurons that drive the muscles (Betz 1874).

The spatial arrangement of neurons in M1 exhibits the functional organization of a topographical map of the body (M. H. Schieber 1999; M. H. Schieber and Hibbard 1993; M. H. Schieber 2001). For example, the medial part of M1 corresponds to leg and proximal arm movements. In contrast, the lateral regions are related to the distal arm and facial movements (Penfield and Boldrey 1937).

There is evidence from primate studies that neurons in the sulcus are more closely associated with dexterous finger movements (Poliakov and Schieber 1999; Marc H. Schieber 1996; M. H. Schieber 1999, 2001; M. H. Schieber and Hibbard 1993). However, it may be difficult to record neurons from the sulcus as the recording electrodes need to be further away from the surface of the cortex. The data I will analyze in the upcoming chapters were collected using Utah multielectrode arrays from the neurons in the “hand area” on the gyrus, with neurons changing their activity patterns with respect to movements and generated forces about the wrist, hand, and to some extent, fingers.

A brief history of the physiology of M1

The precise function of M1 during motor behavior has been debated for decades: is the primary motor cortical activity associated with kinematic position and velocity, or does it encode muscle activity and force? How about the representation of posture and context? Although contemporary population-level understanding of the primary motor cortex acknowledges that many different classes of function are compatible, a complete understanding of how M1

generates movement still remains uncertain. Decades-long research on single-neuron analyses of the motor cortex has shown evidence for its many movement-related functions.

One of the first pieces of evidence for the relationship of M1 with force and muscle activity came in the 1960s. In his groundbreaking study, Evarts investigated the relationship between M1 and movement variables by recording the pyramidal tract neurons in monkeys performing wrist deflection tasks (Evarts 1968). These neurons were more strongly correlated with force and rate of change of force than with displacement about the wrist. Fetz and Baker showed that increased muscle activation typically coincided with bursts of neuronal activity (Fetz and Baker 1969). Another study published a year later by Humphrey and colleagues supported Evarts' findings by predicting force most accurately from a small population of neurons, though position, velocity, and rate of change of force were also predictable to a lesser extent (Humphrey, Schmidt, and Thompson 1970). Toward the end of the 1970s, Thach classified M1 neurons into categories based on their correlation with different movement variables, such as the force of muscular activity, joint position, and direction of movement (Thach 1978). In the beginning of the 1980s, Cheney and Fetz found strong correlations between cortical neurons and active torque during isometric tasks and tasks against various levels of elastic loads (Cheney and Fetz 1980).

Shortly after, a new theory emerged, shifting the focus from muscle-like representations to the relationship between motor cortical neural activity and kinematic variables (Georgopoulos et al. 1982, 1983; Georgopoulos, Schwartz, and Kettner 1986). Georgopoulos and colleagues found that most of the motor cortical neurons that modulated their activity with movement had a "preferred direction." They reported that neurons changed their firing rates in an orderly manner according to the kinematics of the movement and maximally fired at their preferred direction (Georgopoulos et al. 1982). Even though there were concerns that the existence of preferred directions of neurons does not necessarily imply that M1 encodes kinematic variables (Mussa-Ivaldi 1988), numerous studies adopted this Georgopoulos-esque framework, highlighting the

correlation between motor cortical neural activity and kinematic variables (Alexander and Crutcher 1990; Caminiti, Johnson, and Urbano 1990; Paninski et al. 2004; Schwartz and Moran 2000). The dominant view at this time was that M1 encoded the kinematics of reaching movements more strongly than force or muscle activation.

These kinematic studies, along with earlier experiments relating M1 activity to force and EMG, show the ability of M1 to encode various movement variables. Although there is no consensus on which variable is most reliably represented in M1, some relatively recent studies have shifted the focus back to muscles by demonstrating the superior robustness of M1 to EMG relationship over kinematics (S. H. Scott and Kalaska 1997; Sergio, Hamel-Pâquet, and Kalaska 2005; Morrow, Jordan, and Miller 2007; Cherian, Krucoff, and Miller 2011). On one hand, the stronger association of M1 activity with muscle activation is unsurprising because the neurons ought to activate the muscles to generate movement. On the other hand, given that the role of M1 is to produce movement and not necessarily represent it (Stephen H. Scott 2004; Cisek 2006), it is sufficient for M1 to only encode higher-level kinematic goals; such functional organization would allow the details about muscle activation to be carried out by the remainder of the neural circuitry. Despite the lack of agreement among researchers, the fact remains that many variables related to movement are represented in M1, including information related to the activation of muscles.

Advancements in neural recording technology have enabled a population-level view of cortical activity

The importance of studying neurons in groups rather than in isolation has been recognized since the early days of neural recordings (Perkel, Gerstein, and Moore 1967a, 1967b). Yet, many of the studies described in the previous section relied on the investigation of single neurons from the primary motor cortex, partly due to limitations in recording technology. Today, we no longer have the same limitation. We can simultaneously record from hundreds to a few

thousands of neurons using multielectrode arrays with recording technologies such as Utah arrays (Nordhausen, Maynard, and Normann 1996; Hong and Lieber 2019) and Neuropixels (Steinmetz et al. 2018, 2021). We can even record from up to a million neurons using optical imaging techniques (Demas et al. 2021). Compared to multielectrode array recordings, the latter approach has significant tradeoffs between spatial and temporal resolution (Siegle et al. 2021; Huang et al. 2021).

Urai and colleagues' recent analysis indicated that the traditional "Moore's Law" concerning the number of neurons that can be recorded simultaneously has been surpassed by modern recording technologies (Stevenson and Kording 2011). Today, we can make whole-brain recordings from roundworms (*Caenorhabditis elegans*) with roughly 1,000 neurons (Schrödel et al. 2013) and larval zebrafish with over 100,000 neurons (X. Chen et al. 2018). Based on this trend, it is anticipated that capturing the activity of over a 100 million neurons in the entire mouse brain could be achievable within the next twenty to one hundred years (Urai et al. 2022). The Brain Research Through Advancing Innovative Neurotechnologies (BRAIN) Initiative of the National Institute of Health, among other significant endeavors, is committed to accelerating this progress in the future (Saxena and Cunningham 2019).

Large-scale recordings have led to a shift in focus from hand-picked neurons to more comprehensive and unbiased surveys of neural responses. As simultaneous large-scale measurements at a single-neuron resolution have become more feasible (Schrödel et al. 2013; Ahrens et al. 2013; Marblestone et al. 2013; Mann, Gallen, and Clandinin 2017; Aimon et al. 2019; Yamamoto and Yuste 2020; Kleinfeld et al. 2019), the ability to record from many neurons at once has increased the statistical power, reduced the number of research animals required and enabled capturing rare responses or cell types (P. H. Li et al. 2015). These large-scale recordings have also shown that neural activity related to movement is more widely distributed

than previously thought (Musall et al. 2019; Stringer, Pachitariu, Steinmetz, Reddy, et al. 2019; Salkoff et al. 2020).

Despite these strengths, large-scale neural data have also presented large-scale challenges in understanding brain function. In the next section, I will review the newer framework that transcends some of these limits. Central to this newer framework is the “neural manifold,” the low-dimensional geometry that constrains the collective activity of recorded neurons. Chapters 2, 3, and 4 are all based on computation and interpretation of the neural manifolds in M1 from large-scale (roughly 100 neurons) multi-electrode array recordings from monkeys and humans.

Neural manifolds: from a single-neuron to a population-level view of the brain

How should we connect large-scale neural data to behavior and cognition? One possibility is to apply the single-neuron framework to large-scale recordings. Even though single neurons are the basic units of computation in the brain, viewing them individually is limited for multiple reasons. First, single-neuron analyses of potentially millions of recorded neurons are cumbersome; also, investigating the individual tuning properties of each of the millions of neurons to innumerable movement-related variables would be difficult. Second, neural computations such as movement planning and decision-making may only be evident at the level of neural populations; single neurons may not show an obvious tuning to stimulus or task variables (Steinmetz et al. 2019; Mante et al. 2013; Shenoy, Sahani, and Churchland 2013; Russo et al. 2020; Saxena et al. 2022). Therefore, Georgopoulos-esque single-neuron analyses are limited when the function of neural populations is in question, begetting the need for a more comprehensive framework.

Even though contemporary electrophysiology allows simultaneous recordings from many thousands of neurons, the number of recorded neurons is still a mere drop in the ocean

compared to the total number of primary motor cortical neurons that are activated during movement. For example, millions of motor cortical neurons become activated when executing even a simple reaching behavior (P. Gao and Ganguli 2015; Hennig et al. 2018; Russo et al. 2020). Yet, the activity of these small recorded populations of neurons can explain a large portion of task-relevant variability (Cunningham and Yu 2014; P. Gao et al. 2017). How can this be possible?

This baffling observation is in part due to a critical empirical observation about the brain. Large-scale recordings revealed that the brain does not utilize all the potential degrees of freedom provided by the neurons for behavior generation because neural activity tends to be correlated. One dominant theory about the source of this covariation is due to the constraints imposed by the wiring between neurons (Sadtler et al. 2014). This covariation of neural activity limits the possible patterns of collective neural activity; there are a lot of unexplored regions within the state space spanned by the activity of each individual neuron, as neural activity occupies fewer states than it would if each neuron contributed independently or randomly to the population activity (see Cunnigham and Yu 2014 for a graphical representation). The portion of the neural state space containing the observed states is called the “neural manifold” (see Gallego et al. 2017 for a comprehensive review). While the term “manifold” is a mathematical term with precise definitions of topological continuity, certain local properties, and smoothness (J. Lee 2010), in the context of neuroscience the term is often used loosely to describe the geometry within which the neural population activity is contained.

Activity on the neural manifolds is indicative of processing within and across neural populations

The neural manifold framework serves as an effective tool to connect the activity of populations of neurons to behavior. Many of the earlier examples of how the neural manifolds help us to understand neural computation came from studies on the motor cortices. Santhanam and

colleagues provided one of the earliest findings of a neural manifold associated with movement control (Santhanam et al. 2009). They analyzed population activity in the dorsal premotor cortex (PMd) during a delayed center-out reach task and discovered that a three-dimensional manifold was sufficient to identify target-specific clusters of activity during the delay period. Subsequently, Churchland et al. demonstrated a systematic decrease in trial-to-trial variability in the activity on the neural manifolds associated with PMd following stimulus onset (Churchland, Yu, et al. 2010). Since then, there has been a surge of studies that showed that population-level activity on the neural manifolds is critical for understanding the function of neural circuits across many brain areas (Stopfer, Jayaraman, and Laurent 2003; DiCarlo and Cox 2007; Churchland et al. 2012; Mante et al. 2013; Kaufman et al. 2014; Sadtler et al. 2014; Gallego et al. 2017; Chung, Lee, and Sompolinsky 2018; Remington et al. 2018; Low et al. 2018; Rubin et al. 2019; Chaudhuri et al. 2019; Russo et al. 2020; Nieh et al. 2021; Libby and Buschman 2021; Chandak and Raman 2021; Ehrlich and Murray 2022; Gardner et al. 2022).

One could argue that the conclusions drawn from the studies mentioned above do not necessarily require the concept of neural manifolds. At first glance, a neural manifold may be misconstrued as a mere geometric simplification, seemingly providing little to no insight into the underlying computational principles within the recorded population. However, the importance of neural manifolds becomes evident when we face challenges in interpreting the intricate activation patterns of individual neurons, which can mask the fundamental computational principles at play.

Let us take a delayed reaching task as an example, where the motor cortex must prepare for a movement first and execute it after a certain delay. In both the preparation and execution stages, individual neurons display complex firing patterns. If our analysis is limited to studying the responses of these individual neurons and their respective correlations with the two stages,

we fall short in comprehending how neural populations manage to prepare for movement without prematurely setting it into action.

This is where the power of neural manifolds come into play; neural manifolds provide an understanding of such processes by unveiling the computational principles that may be overshadowed by the activities of individual neurons. To accurately answer the question of how neural populations prime themselves for movement without causing it, a thorough exploration of neural manifolds becomes essential.

Churchland, Shenoy, and colleagues (Churchland, Cunningham, et al. 2010; Churchland et al. 2012) employed the neural manifold concept to provide the groundwork for the phenomenon of how the rich and complex neural activity in both PMd and M1 during movement planning does not generate movement. To explain this further, Kaufman and colleagues identified a low-dimensional neural manifold from M1 recordings and linearly related these latent variables to the activity of muscles (Kaufman et al. 2014). They divided the neural manifold into an *output-potent* space, controlling muscle activity, and an *output-null* space, not directly affecting muscle activity. They showed that preparatory activity lies in the output-null space, serving as an initialization point from which the population activity evolves to generate the desired movement. Elsayed and colleagues expanded this analysis, demonstrating that preparatory and movement activities lie in orthogonal spaces within the manifold, with population dynamics evolving from one to the other going from preparation to movement (Elsayed et al. 2016).

The distinction between output-potent and output-null spaces was employed by Slutzky and colleagues to investigate the long-term stability of brain-computer interfaces (Flint et al. 2016). They discovered that not all recorded neurons needed to be uniformly stable for BCI control and that neural activity in the output-potent space was significantly more stable than that in the output-null space.

In addition to within-area processing, the neural manifolds framework has also provided valuable insights into multi-region communication. Some studies have identified specific activity patterns in one brain area inherited from an upstream region, referred to as a "communication subspace" (Semedo et al. 2019). For instance, in the macaque visual system, only a small subset of the primary visual cortex (V1) population variability drives variability in the secondary visual cortex (V2). This V1-to-V2 predictive dimension is largely non-overlapping with the V2-to-V1 feedback subspace. This example shows that V1 can route selective activity and reduce unwanted co-fluctuations in downstream areas; some of the activity patterns are communicated to downstream structures, while others remain private. It would not be possible to uncover this mechanism of processing and computation by looking at single neurons in isolation. In short, the neural manifold framework has been essential for understanding larger, multi-region neural recordings and promises to provide further understanding of brain-wide communication.

A recent interest has been to investigate the similarity of neural manifolds across conditions. Are they completely unique to a specific task and neural population, or do they share common traits across different behaviors, time spans, and individuals? One study found that the activity on the neural manifolds associated with skilled hand and arm tasks was to some extent preserved, despite clear differences in the activity of recorded neurons and muscles (Gallego et al. 2018). Another study found the neural manifold associated with head direction in mice was consistent even when the animal was actively foraging or asleep, two very distinct cognitive states (Chaudhuri et al. 2019). Similarly, the activity on the neural manifolds was consistent across time for a given behavior despite changes in the recorded neurons across sessions (Gallego et al. 2020).

Interestingly, the similarity of neural manifolds extends beyond tasks, cognitive state, and time. For example, there is some preliminary evidence that neural manifolds could be similar across species executing the same motor behavior. The neural manifolds computed from the motor

cortical recordings from human participants exhibited rotational trajectories (Pandarinath et al. 2015) similar to those observed in monkeys (Churchland et al. 2012). Similarly, a study showed that latent activity on the neural manifolds of monkeys and mice performing the same behavior was similar and argued that manifolds arise from constraints on brain development from an evolutionary perspective (Safaie et al. 2022).

The neural manifolds and the associated population activity offer valuable insights into the brain's process of generating movement and representing task-relevant information. Despite the complexity of patterns at the level of individual neurons, these manifolds remain remarkably consistent across various behaviors and over time. Crucially, there is preliminary evidence suggesting that neural manifolds might be stable across different individuals and across species, despite each existing in distinct coordinate frames spanned by their respective neurons. These consistencies suggest that neural manifolds serve as a fundamental framework for probing the computational principles within neural populations.

In Chapter 4, drawing from the stability of neural manifolds across individuals, I will investigate the feasibility of cross-user EMG decoding, which can potentially address the limitations of observation-based kinematic iBCI decoders. I will first test two cross-user decoding approaches from one monkey to another and then attempt cross-user decoding between a monkey and a human. My investigations in Chapter 4 will highlight the potential of the neural manifolds for improving iBCI technology for transforming the lives of people living with paralysis through intuitive muscle-based control.

Dimensionality reduction methods for identifying the neural manifold

Dimensionality reduction is a crucial step in computing the latent variables that define the neural manifolds, as it is a tool for simplifying a large number of correlated signals, such as the spiking activity of neurons. The goal of dimensionality reduction is to represent the activity of N

individual neurons using a smaller number of d latent variables (Cunningham and Yu 2014). The low-dimensional latent variables consist of a relatively small number of components that capture the majority of population variance (Stopfer et al. 1997; Sadtler et al. 2014). Since d is a smaller number than N, this operation is called dimensionality *reduction.*

Broadly, there are two types of dimensionality reduction methods: linear and nonlinear. Based on the underlying assumption on the nonlinearity of the neural manifold, the appropriate method might be linear or nonlinear. Nevertheless, both classes of dimensionality reduction often necessitate an *a priori* knowledge of the presumed number of latent variables, d, that define the neural manifold. The number of latent signals is also known as the "dimensionality" of the neural manifold. Before discussing different types of dimensionalities, and their importance regarding neural computation, I will first briefly overview linear and nonlinear dimensionality reduction methods that have been commonly used to compute latent variables that define the neural manifolds (see Cunningham and Yu 2014, for a detailed review of dimensionality reduction in large-scale neural recordings).

Linear dimensionality reduction methods, such as principal component analysis (PCA) and factor analysis (FA), are often the first choice for computing the low-dimensional activity associated with neural manifolds. PCA identifies an orthonormal set of direction in neural state space that capture the most variance in the neural data, while FA focuses on preserving shared variance across neurons and discarding independent variance for each neuron, which is considered to be spiking variability (Churchland, Yu, et al. 2010). Time series methods, like Gaussian process factor analysis (Yu et al. 2009), linear dynamical systems (A. C. Smith and Brown 2003; Kulkarni and Paninski 2007; Macke et al. 2011; Buesing, Macke, and Sahani 2012), and hidden Markov models (Abeles et al. 1995; Seidemann et al. 1996; Ponce-Alvarez et al. 2012; Bollimunta, Totten, and Ditterich 2012), further leverage the sequential nature of neural data to provide a more refined account of temporal dynamics in population activity. These linear

methods have been favored due to their computational simplicity, interpretability, and robustness to noise.

However, there is no prior knowledge that the neural manifold is linear. In fact, there is growing evidence that neural manifolds exhibit nonlinear geometries (Chaudhuri et al. 2019; De and Chaudhuri 2022; Gardner et al. 2022). In this scenario, the linear dimensionality reduction methods would yield unfaithful representations of the underlying nonlinear nature of the neural manifold. Fortunately, plenty of nonlinear dimensionality reduction methods exist.

Nonlinear dimensionality reduction methods, such as Isomap (Tenenbaum, de Silva, and Langford 2000), locally linear embedding (Roweis and Saul 2000; Broome, Jayaraman, and Laurent 2006; Mitchell-Heggs et al. 2023; Saha et al. 2013), and uniform manifold approximation and projection (McInnes, Healy, and Melville 2018; Gardner et al. 2022) are designed to capture low-dimensional structure in data that may not be well represented by linear relationships between latent variables and observed neural signals. Although these methods can potentially reveal more complex relationships and hidden structures in neural data, they come with their own set of challenges. For example, they rely on local neighborhood estimation between the states in the neural state space, which can be problematic if the high-dimensional space is not densely sampled (Terrell and Scott 1992; Yianilos 2000; Bengio, Delalleau, and Le Roux 2005). In addition, there is evidence that nonlinear methods can be more sensitive to noise and could be harder to interpret (Boots and Gordon 2012).

The widespread adoption of deep learning across domains of science has come to neuroscience as well. These days, many use nonlinear dimensionality reduction methods based on deep learning to compute the latent variables that define the neural manifolds. These deep learning methods, although computationally more expensive, have been effective in identifying the latent variables associated with neural manifolds. Feedforward or recurrent autoencoders (Hinton and Salakhutdinov 2006) have been successfully used to extract nonlinear manifolds

that explain significant variation in neuronal spiking activity (Sedler, Versteeg, and Pandarinath 2022).

Compared to traditional dimensionality reduction methods, the deep learning framework offers several notable advantages. First, the deep learning based architectures allow for the capture of highly nonlinear relationships in the data, which can be essential for uncovering potentially complex and nonlinear neural manifolds. Second, the deep learning framework is modular, which easily allows for explicitly modeling the temporal relationships within the neural data. For example, a static feedforward architecture could be replaced with a recurrent one, which would be important for capturing the inherent temporal structure and predicting the future states of neural activity should that be of interest.

One widely-used deep network model that gained considerable traction in neuroscience is Latent Factor Analysis via Dynamical Systems (LFADS) (Pandarinath, O'Shea, et al. 2018). The architecture of LFADS consists of several components, including an encoder, a decoder, and a controller, all three comprised of recurrent neural networks (RNNs). The encoder converts the neural data into a set of initial conditions to a nonlinear dynamical system. The controller infers the inputs needed to fully specify the dynamical system. Finally, the decoder reads out the dynamical system to predict the input neural data. LFADS is trained using variational inference techniques for robustness to noise (Im et al. 2017; Kingma and Welling 2019).

In Chapters 2 and 3, when I compute M1 manifolds associated with movement, I will use nonlinear autoencoders. The Joint Autoencoder network that I will present in Chapter 2 is also a form of a nonlinear autoencoder.

Intrinsic and embedding dimensionality of neural manifolds

There remain ambiguities about the structure of neural manifolds and how their structure relates to function. For example, neural manifolds have geometric properties, including dimensionality,

d, and nonlinearity. The linear and nonlinear dimensionality reduction methods outlined in the earlier section require knowledge, or an estimate, of d to faithfully compute the latent signals that evolve in the neural manifolds.

However, it is unclear how the nonlinear shape or dimensionality of the neural manifold is tied to neural function. For instance, the notion of "low" in the context of what is considered “low-dimensional” is not clear. Additionally, there remain unanswered questions regarding how behavioral complexity and task constraints influence the nonlinearity and dimensionality of the neural manifold and how these geometric parameters of the manifold relate to computational principles.

Recent theoretical work related the dimensionality of neural activity to the dimensionality of a given task (P. Gao and Ganguli 2015; P. Gao et al. 2017). This work was largely done in the motor cortices, and the low dimensionality of M1 could be tied to the simplicity of motor tasks done in constrained laboratory environments. According to this theory, the task parameters and constraints that restrict task-irrelevant movement provide an upper bound to the observed dimensionality in M1. Indeed, studies investigating neural manifolds associated with constrained laboratory tasks have reported that the dimensionality of M1 manifolds ranges between 10-20 (Sadtler et al. 2014; P. Gao et al. 2017; Russo et al. 2018; Perich, Gallego, and Miller 2018), a vastly smaller number than the millions of M1 neurons that change their activity during movement. These findings led to the description of M1 manifolds as low-dimensional.

A fundamental question surrounding this observation is whether the observed low dimensionality of M1 is a simple byproduct of constraints in the laboratory. I attempt to answer this question in Chapter 3, where I will compute neural manifolds associated with unconstrained, natural behaviors and compare them to those associated with constrained movements in the laboratory.

One critique of the theory proposed by Gao and Ganguli stems from the brain's need to process not only task-relevant but also task-irrelevant information. This processing could expand the intrinsic and embedding dimensionality of the neural manifolds beyond the inherent complexity of the task (Ebitz and Hayden 2021). The integration of multiple sensory, cognitive, and task-relevant variables, even in seemingly simple decision-making tasks, could introduce complexities and blur the relationship between task complexity and the corresponding dimensionality of the neural manifolds.

In a recent study where mice were tasked to choose an image out of two, their decisions incorporated factors such as the locations of the images, previous rewards, and past choices (C. S. Chen et al. 2021). Although the authors did not explicitly calculate the associated neural manifold dimensionality, they claimed that it would exceed the one-dimensional complexity of the task.

Another study aimed to explicitly connect task dimensionality with neural dimensionality in population recordings from hippocampal neurons (Low et al. 2018). Using methods that could account for nonlinear relationships, they found that the hippocampal manifolds not only captured task-relevant information but also consistently contained one extra dimension. This observation was true for both one- and two-dimensional navigation tasks they investigated: the manifold was two- and three-dimensional, respectively. The researchers proposed that the extra dimension enabled neuronal population activity to trace different paths on the neural manifold even for identical behavioral conditions and highlighted the presence of an internal cognitive process.

Given the unclear relationship between the geometry of neural manifolds and computational processing, Jazayeri and Ostojic introduced just last year a new theory that is already gaining rapid recognition (Jazayeri and Ostojic 2021). In this theory, the authors split the notion of dimensionality into two complementary concepts: *intrinsic* and *embedding dimensionality*. In this model, intrinsic dimensionality reflects the nature of the information encoded in collective neural

activity, while embedding dimensionality reveals how that information is processed by neural circuits. While they are related, these two notions of dimensionality provide distinct insights into how neural populations represent and process information. Contrasting intrinsic and embedding dimensionality of neural manifolds may offer a promising perspective for understanding brain function.

According to this theory, the intrinsic dimensionality of neural activity is largely determined by sources of information related to incoming stimuli, ongoing movements, and other cognitive latent variables such as prior experiences and future expectations involving the task. However, these sources of information are not represented equally across neural populations. For example, in early sensory areas, intrinsic dimensionality would be more closely associated with incoming visual stimuli in the primary visual cortex (Stringer et al. 2021; Zhao and Park 2017; Ringach 2019), or the organization of chemical odors in the olfactory areas (Pashkovski et al. 2020; Chandak and Raman 2021). In contrast, the intrinsic dimensionality of the neurons in the motor cortex would be more closely associated with ongoing movements (Churchland, Yu, et al. 2010; Churchland et al. 2012; Sussillo et al. 2015; Kato et al. 2015) than variables like context and upcoming movements, the latter of which are well represented in simultaneously recorded neurons from the supplementary motor cortex (Russo et al. 2020; Zimnik and Churchland 2021). Therefore, investigating the intrinsic dimensionality of neural manifolds could reveal how a population encodes information.

Intrinsic dimensionality, linked to the potentially nonlinear structure in collective neural activity, provides insights into how a population represents stimuli, movements, and other important latent variables. However, it does not give much detail about how these variables might be processed. The downstream relay of latent variables relies on their positioning, or *embedding*, within the neural state space. Embedding dimensionality measures the linear dimensions that

govern these latent variables and can potentially inform insights into downstream processing mechanisms (Badre et al. 2021).

From this perspective, the embedding dimensionality informs the computational trade-off between generalizability and expressivity (or flexibility) of neural representations (Rigotti et al. 2013; Fusi, Miller, and Rigotti 2016; Maass 2016; Cohen et al. 2020). When the embedding dimensionality is high, it facilitates the transmission of information encoded by a neural population downstream without interference, which enhances the flexibility of how downstream populations can access the encoded information.

To illustrate this, imagine a neural population that encodes three distinct states: apples, penguins, and pears. If these states are encoded in a large number of dimensions, downstream populations could employ simple linear decoders to extract information from arbitrary groups of states. For instance, it would be easy to access apples and penguins together while excluding pears (Ebitz and Hayden 2021). In contrast, if these states are embedded in fewer dimensions, it would be harder to linearly distinguish the group containing apples and penguins without also including pears. Therefore, a neural population offering a high level of flexibility would have higher embedding dimensionality than a population lacking it, even if populations may encode the same latent variables.

The hypotheses about brain information processing can be viewed as different forms of embedding. For example, if distinct sets of neurons perform unique functional roles, the latent variables would exist in subspaces governed by these subsets of neurons alone. Conversely, when different information sources are processed by groups of neurons that overlap, the resulting embeddings no longer align with the principal axes of the state space. This type of processing, also known as mixed selectivity (Rigotti et al. 2013), has been strongly associated with the high embedding dimensionality observed in brain areas related to cognition (Fusi, Miller, and Rigotti 2016; Langdon, Genkin, and Engel 2023).

Interpreting embedding dimensionality requires careful thought. This concept is especially relevant when information is relayed downstream linearly; one could imagine that even if apples, penguins, and pears are embedded in very few dimensions, a highly nonlinear decoder could group apples and penguins while excluding pears. Although the evidence revised earlier regarding the linear relay of information downstream is abundant, it is important to remember that this could be an oversimplification of how populations communicate. The linearity assumption might be too simplistic and may fail to truly represent the brain's intricate and potentially nonlinear processing methods. As such, in our attempt to link embedding dimensionality to how the brain processes information, we must proceed with caution, as these concepts have not yet been fully elucidated.

In summary, although our understanding of this novel theory is still growing and developing, emerging evidence indicates that comparing the intrinsic and embedding dimensionality of neural structures can help us better describe how neurons carry and process information. Intrinsic dimensionality provides insight into the type of information being processed—the "what." On the other hand, embedding dimensionality relates to the way this information is arranged for further processing—the "how." By examining these two aspects of dimensionality within and between brain regions during certain behaviors, we can gain insights into the computations performed by these neural populations.

In Chapter 3, I will directly apply the concepts of intrinsic and embedding dimensionality to investigate the geometry of M1 manifolds and explore how it might process task constraints—is the low intrinsic and embedding dimensionality of M1 due to the constrained tasks in the laboratory setting, or rather it is part of a computational strategy? The answer to this question has implications for translating M1-to-muscle iBCIs outside the laboratory. In Chapter 4, I will use the estimates of embedding dimensionality to accurately interpret muscle-related

information from M1 structures across different individuals, which will serve as a proof-of-concept study aimed at developing intracortical iBCIs to work across users.

Estimating intrinsic and embedding dimensionality

Estimating the intrinsic and embedding dimensionality of neural manifolds can be a challenge when dealing with datasets that contain noise and nonlinearity (Camastra and Staiano 2016). In the context of neural recordings, noise can stem from various sources, including sensory transduction, voltage-gated ion channels, synapses, and experimental factors associated with the recording setup (Faisal 2008). To complicate matters further, the accuracy of dimensionality estimators also relies on the number of available features (i.e., the number of recorded neurons) and the number of samples (Camastra 2002, Camastra 2003), the latter of which is often limited in experimental settings. To uncover neural manifolds and connect their geometry to function, we need robust methods that can handle these challenges.

Fortunately, large-scale neuroscience is not the only field in need of robust dimensionality estimators. “Big data” comes in many forms and can often be described with lower-dimensional representations without significant information loss (Wolfe 2013; Fan, Han, and Liu 2014; Snášel et al. 2017; Z. Sun et al. 2020). To address this need, many groups have developed dimensionality estimators and benchmarked them on synthetic datasets with known ground truth geometric properties (Camastra and Staiano 2016; Campadelli et al. 2015; Denti et al. 2022; Pope et al. 2021; Ansuini, Laio, and Macke 2019; Rozza et al. 2012; C. Li et al. 2018; Kégl 2002; Einbeck and Kalantana 2013; Lombardi et al. 2011; Facco et al. 2017; Ceruti et al. 2014; Albergante, Bac, and Zinovyev 2019).

We can broadly classify dimensionality estimators as linear and nonlinear based on their assumptions on the relationship between the low-dimensional latent signals and the high-dimensional data. For my purposes, the data features are the activity of simultaneously

recorded neurons. Linear and nonlinear methods would be more appropriate for estimating the embedding and intrinsic dimensionality of the neural manifold, respectively. In Chapter 2, I will benchmark a battery of linear and nonlinear dimensionality estimators on simulated neural data with known ground truth properties. I will expand on the linear and nonlinear dimensionality estimation methods in detail in the Methods section of Chapter 2, but I will present the intuition behind how they work in the next two sections. In Chapters 3 and 4, I will apply the most accurate dimensionality estimation algorithms we found in simulated data to experimental M1 recordings from monkeys and humans.

Estimating the embedding dimensionality of the neural manifold

When discussing dimensionality *reduction*, I referred to PCA as a popular linear strategy. PCA can also be used for dimensionality *estimation*, particularly in relation to the embedding dimensionality of the neural manifold. Although PCA offers a structured approach to dimensionality reduction, it does not offer a definitive strategy for identifying the embedding dimensionality of the neural manifold. Often, it necessitates 1) arbitrarily selecting a variance threshold that we aim to retain in the low-dimensional data representation, 2) ordering the principal components from the most to least amount of variance of the original data they individually capture, and 3) counting the number of principal components that are needed to achieve the threshold.

The initial step of setting a variance threshold directly influences the number of principal components that are retained. A lower variance threshold would lead to fewer components and influence the final estimate of the embedding dimensionality of the underlying manifold. However, more sophisticated PCA-based algorithms such as Participation Ratio (PR) (P. Gao et al. 2017) and Parallel Analysis (PA) (Horn 1965; Buja and Eyuboglu 1992) offer a more principled approach, as these methods circumvent the need for an arbitrarily selected variance threshold. PR uses the relative decline in the variance explained for each successive principal

component to yield an estimate. On the other hand, PA requires repeatedly shuffling data along the sample axis independently for each feature and applying PCA on the shuffled data to generate null distributions for the explained variance by each principal component. Then, one would identify the principal components in the original data that account for more variance than their corresponding null distributions. Therefore, PA yields an embedding dimensionality estimate by determining the number of principal components that contribute significantly to the data's overall variance beyond what might be expected by chance alone.

In Chapter 2, I will first generate neural data from simulated neural manifolds with known ground truth and then evaluate the performance of these rigorous PCA-based methods against the commonly used practice of arbitrarily setting a variance threshold and counting the number of leading principal components.

Estimating the intrinsic dimensionality of the neural manifold

Linear algorithms such as PCA with a variance threshold and its variants like PR and PA may perform well in estimating the embedding dimensionality of the neural manifold. However, they overestimate the intrinsic dimensionality. The extent of this overestimation depends on the extent of nonlinearity of the neural manifold (Camastra and Staiano 2016; Tenenbaum, de Silva, and Langford 2000; Hinton and Salakhutdinov 2006). In these cases, nonlinear dimensionality estimators become more appropriate for estimating the intrinsic dimensionality of the neural manifold.

The nonlinear dimensionality estimators usually exploit the patterns in the local distances between data points and yield an estimate based on an implicit comparison to the patterns seen in well-studied synthetic manifolds with known dimensionality. For example, one of these patterns is the distributions of the distances between neighboring points in the manifold, which is shown to be informative of the intrinsic dimensionality (A. N. Gorban and Tyukin 2018). In

Chapter 2, I will benchmark some of the well-studied intrinsic dimensionality estimators such as Correlation Dimension (Grassberger and Procaccia 1983; Grassberger 1983), Levina-Bickel Maximum Likelihood Estimation (Levina and Bickel 2004), Two Nearest Neighbors (Facco et al. 2017), and Fisher Separability Analysis (Albergante, Bac, and Zinovyev 2019) on simulated neural datasets.

Summary

From simple actions like grasping a cup to complex movements such as playing the guitar, our ability to perform highly dexterous motor tasks is driven by M1. Understanding how M1 generates these movements has been a central part of modern neuroscience since the 19th century. Contemporary recording techniques, such as multielectrode array recordings, have allowed us to simultaneously record the firing activity of hundreds to thousands of neurons on this quest. While there are millions of active M1 neurons, our recordings of just a small fraction of these neurons have been remarkably predictive of variables related to movement, such as kinematics, force, and muscle activity. This observation raises a critical question: how is it possible that we can meaningfully understand the relationship between M1 activity and motor behavior despite substantially undersampling the active neurons?

The answer to this question lies in the correlated activity of neurons; rather than exhibiting entirely unique activity patterns, neurons fire together in patterns. If we construct a coordinate system where each axis corresponds to the firing activity of individual neurons, the time-varying activity of the population is constrained to a much smaller geometric space spanning fewer dimensions than the number of recorded neurons, leaving most of the coordinate system unexplored. This low-dimensional structure encapsulating the collective activity of neurons is referred to as the “neural manifold.”

Over the last decade, neural manifolds have been a common observation across various brain areas including M1. These manifolds have offered a powerful framework for understanding computations within and across populations of neurons, addressing some of the limitations of a single-neuron framework. For example, the complex activity of individual M1 neurons clouded our understanding of how the neurons in M1 prepare for movement without executing it. The neural manifold framework showed us that the preparatory and movement-related epochs are organized in orthogonal directions within the neural manifold, a possible explanation of how a given set of neurons can prepare for movements before causing them.

Despite the widespread use of neural manifolds, there are gaps in our understanding of the relationship between them and motor behavior. This thesis aims to address some of these gaps. For example, we do not know how the dimensionality and nonlinearity of these manifolds are related to different behaviors, each requiring different computational processes.

The first gap lies in the limitations in the accurate estimation of the dimensionality and nonlinearity of neural manifolds. Chapter 2 of the thesis is dedicated to improving this estimation process using simulated neural datasets with known dimensionality, nonlinearity, and levels of noise. In this chapter, I will introduce a deep-learning-based denoising algorithm, termed the "Joint Autoencoder," and devise a pipeline to systematically measure the dimensionality and nonlinearity of neural manifolds. The development and implementation of these improved methods constitute the primary objective of this chapter.

Chapter 3 will address the question of whether the observed low dimensionality of M1 manifolds functions as a general principle that transcends the boundaries of laboratory conditions or if it is merely an artifact of constraints in laboratory tasks. To fill this knowledge gap, this chapter will investigate the M1 activity of monkeys during natural behaviors beyond the typical laboratory conditions, such as grasping small treats like blueberries and walking on perch bars. The objective here is to compare the dimensionality of the neural manifolds during these natural

behaviors to those observed under controlled laboratory tasks to gain insight into the computational principles of M1.

Finally, Chapter 4 will focus on applying the concept of low-dimensional neural manifolds to the field of iBCIs. Current iBCIs have an important limitation because they require observable variables, like the kinematics of a cursor on a screen, for building decoders. This approach is particularly problematic for biomimetic, muscle-based iBCIs because muscle activity cannot be observed. To overcome this challenge, this chapter will introduce two novel strategies grounded on the neural manifold concept. The primary objective of Chapter 4 is to enhance the functionality of muscle-based iBCIs with two cross-user decoding approaches, paving the way to empower individuals with paralysis to regain control of their muscles using their neural activity.

CHAPTER 2: ESTIMATING THE DIMENSIONALITY OF THE MANIFOLD UNDERLYING MULTI-ELECTRODE NEURAL RECORDINGS

Ege Altan[1,2,*], Sara A. Solla[1,3], Lee E. Miller[1,2,4,5], Eric J. Perreault[2,4,5]

[1] Department of Physiology, Northwestern University, Chicago, IL, United States of America

[2] Department of Biomedical Engineering, Northwestern University, Evanston, IL, United States of America

[3] Department of Physics and Astronomy, Northwestern University, Evanston, IL, United States of America

[4] Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL, United States of America

[5] Shirley Ryan AbilityLab, Chicago, IL, United States of America

Foreword

The work in this chapter was published in *PLOS Computational Biology* on November 29, 2021.

Abstract

It is generally accepted that the number of neurons in a given brain area far exceeds the number of neurons needed to carry any specific function controlled by that area. For example, motor areas of the human brain contain tens of millions of neurons that control the activation of tens or at most hundreds of muscles. This massive redundancy implies the covariation of many neurons, which constrains the population activity to a low-dimensional manifold within the space of all possible patterns of neural activity. To gain a conceptual understanding of the complexity of the neural activity within a manifold, it is useful to estimate its dimensionality, which quantifies the number of degrees of freedom required to describe the observed population activity without significant information loss. While there are many algorithms for dimensionality estimation, we do not know which are well suited for analyzing neural activity. The objective of this study was to evaluate the efficacy of several representative algorithms for estimating the dimensionality of linearly and nonlinearly embedded data. We generated synthetic neural recordings with known

intrinsic dimensionality and used them to test the algorithms' accuracy and robustness. We emulated some of the important challenges associated with experimental data by adding noise, altering the nature of the embedding of the low-dimensional manifold within the high-dimensional recordings, varying the dimensionality of the manifold, and limiting the amount of available data. We demonstrated that linear algorithms overestimate the dimensionality of nonlinear, noise-free data. In cases of high noise, most algorithms overestimated the dimensionality. We thus developed a denoising algorithm based on deep learning, the "Joint Autoencoder," which significantly improved subsequent dimensionality estimation. Critically, we found that all algorithms failed when the intrinsic dimensionality was high (above 20) or when the amount of data used for estimation was low. Based on the challenges we observed, we formulated a pipeline for estimating the dimensionality of experimental neural data.

Author Summary

The number of neurons that we can record from has increased exponentially for decades; today we can simultaneously record from thousands of neurons. However, the individual firing rates are highly redundant. One approach to identifying important features from redundant data is to estimate the dimensionality of the neural recordings, which represents the number of degrees of freedom required to describe the data without significant information loss. Better understanding of dimensionality may also uncover the mechanisms of computation within a neural circuit. Circuits carrying out complex computations might be higher-dimensional than those carrying out simpler computations. Typically, studies have quantified neural dimensionality using one of several available methods despite a lack of consensus on which method would be most appropriate for neural data. In this work, we used several methods to investigate the accuracy of simulated neural data with properties mimicking those of actual neural recordings. Based on these results, we devised an analysis pipeline to estimate the dimensionality of neural recordings. Our work will allow scientists to extract informative features from a large number of

highly redundant neurons, as well as quantify the complexity of information encoded by these neurons.

Introduction

Studies that simultaneously record the activity of many neurons have shown that cortical neural activity is highly redundant (Saxena and Cunningham 2019). In primary motor cortex (M1), redundancy arises as tens of millions of neurons control tens or at most hundreds of muscles. This redundancy implies significant covariation in the activity of many neurons, which confines the population neural activity to a low-dimensional manifold embedded in the neural space of all possible patterns of neural population activity (Cunningham and Yu 2014; Elsayed and Cunningham 2017; P. Gao and Ganguli 2015; P. Gao et al. 2017; Trautmann et al. 2019; Williams et al. 2018; Williamson et al. 2019; Gallego et al. 2017). Low-dimensional manifolds have also been observed in a variety of other cortical regions (Rigotti et al. 2013; Mazor and Laurent 2005; Mante et al. 2013; Churchland et al. 2012; Harvey, Coen, and Tank 2012; Wärnberg and Kumar 2019; Gallego et al. 2020; Williamson et al. 2016; Mazzucato, Fontanini, and La Camera 2016). Reliable algorithms for identifying these manifolds and characterizing their dimensionality are increasingly important as our ability to record from large populations of neurons increases (Stevenson and Kording 2011). The dimensionality of the manifold describing the coordinated firing of a set of neurons quantifies the number of degrees of freedom needed to describe population activity without significant information loss (Camastra and Staiano 2016; J. A. Lee and Verleysen 2007). Projecting the observed firing patterns onto the manifold yields a low-dimensional set of latent signals that can simplify the interpretation of population neural activity (Cunningham and Yu 2014; Pang, Lansdell, and Fairhall 2016; Gallego et al. 2017). Low-dimensional latent signals can facilitate the manipulation or the extraction of signals for brain-computer interfaces, a rehabilitative technology that converts neural signals into control

commands to restore movement to paralyzed patients (Pandarinath et al. 2017; Degenhart et al. 2020b).

Unfortunately, it is surprisingly difficult to estimate the dimensionality of neural manifolds, particularly in the realistic condition of a noisy, nonlinear embedding. There is evidence of a nonlinear mapping between the recorded neural activity and the associated low-dimensional latent signals (Rigotti et al. 2013; Y. Gao et al. 2016; Wu et al. 2017; Batty et al. 2022). Noise propagates from the level of sensory transduction and amplification, the opening and closing of voltage-gated ion channels, and builds up at the level of synapses, causing neural firing to be a stochastic process (Faisal, Selen, and Wolpert 2008). The two effects, nonlinearity and noise, combine to pose significant challenges to existing dimensionality estimation algorithms. The accuracy of the estimators also depends on the amount of available data (Camastra and Vinciarelli 2002; Camastra 2003), which is limited in most experimental paradigms. If we wish to identify the manifolds associated with experimentally measured neural activity, we need methods that are robust in the presence of these challenges.

The methods that have been proposed for estimating the dimensionality of neural manifolds can be broadly categorized into linear or nonlinear algorithms, based on assumptions about the nature of the mapping between the low-dimensional representation of the latent signals and the high-dimensional space of neural activity. The most commonly used linear method for dimensionality reduction is Principal Component Analysis (PCA), based on identifying mutually orthogonal directions in the empirical neural space of recorded activity; these directions are monotonically associated with the largest data variance. PCA provides a hierarchical description in which the data projected onto the manifold subtended by the leading principal components become closer and closer to the recorded data as the dimensionality of the linear manifold is increased towards the dimensionality of the empirical neural space. Although PCA provides a useful and systematic tool for variance-based dimensionality reduction, it does not specify how

to uniquely identify the dimensionality of the manifold: the typical implementation requires the choice of an arbitrary variance threshold. Other PCA-based algorithms such as Participation Ratio (PR) (Mazzucato, Fontanini, and La Camera 2016; P. Gao et al. 2017) and Parallel Analysis (PA) (Horn 1965; Buja and Eyuboglu 1992) provide more principled prescriptions for linear dimensionality estimation, by incorporating criteria for determining an optimal number of leading principal components to use when constructing the low-dimensional manifold.

Linear dimensionality estimation algorithms may work well for linear datasets, but are likely to overestimate the dimensionality of a manifold arising from a nonlinear mapping between the low-and high-dimensional spaces (Camastra and Staiano 2016; J. A. Lee and Verleysen 2007; Tenenbaum, de Silva, and Langford 2000; Hinton and Salakhutdinov 2006). In contrast, nonlinear methods (e.g., Correlation Dimension (Grassberger and Procaccia 1983; Kalantan and Einbeck 2012; Einbeck and Kalantana 2013), Levina-Bickel Maximum Likelihood Estimation (Levina and Bickel 2004), Two Nearest Neighbors (Facco et al. 2017), and Fisher Separability Analysis (Albergante, Bac, and Zinovyev 2019)) may provide accurate dimensionality estimates for both linearly and nonlinearly embedded data.

Most dimensionality estimation methods have been tested in the absence of noise even though it is known that linear and nonlinear methods overestimate dimensionality when the data is noisy (Camastra and Staiano 2016). The robustness of dimensionality estimation algorithms to noise remains to be characterized.

The objective of this study was to characterize the accuracy of several dimensionality estimation algorithms when applied to high-dimensional recordings of neural activity. We evaluated previously proposed algorithms on synthetic datasets of known dimensionality to identify conditions under which each method succeeded and/or failed. Specifically, we evaluated how the algorithms handled the nature of the embedding (linear or nonlinear), the amount of noise added to the simulated neural data, and the amount of data available. We found increasing

levels of noise to be a challenge for all tested algorithms. We therefore also evaluated different approaches for reducing noise prior to performing dimensionality estimation, including the “Joint Autoencoder”, a method we developed based on deep learning techniques. Together, our results allowed us to propose a methodological pipeline for estimating the intrinsic dimensionality of high-dimensional datasets of recorded neural activity.

Methods

Simulation of neural signals

We generated the synthetic data used to evaluate the various dimensionality estimation algorithms as follows. First, we created d signals by randomly selecting (d x M) samples from an empirical distribution of firing rates that we obtained from multi-electrode array recordings of neural activity in the macaque primary motor cortex (M1) made while the monkey was performing a center-out task (Gallego et al. 2018). The firing rates were binned at 50 ms. The sampling was done randomly across all recorded neurons and time bins within successful trials. Our goal was to generate M samples of d-dimensional latent variables that were uncorrelated with each other and individually uncorrelated over time; we verified that these randomly selected signals were indeed uncorrelated, as intended. These signals provided a set of variables of known dimension d that preserved the first-order firing statistics of the neural activity recorded in M1. Our procedure aimed at generating simulated data that reproduces possible states of activity of a neural population, without considering the order in which these states might be visited; in other words, we focused on population statistics as opposed to population dynamics.

These signals provided a d-dimensional latent set used to construct synthetic high-dimensional data sets **(Figure 2-1)**. We allowed d to vary from 3 to 40. For the analyses where a fixed value of d was used we chose d=6, to approximate the characteristics of real data collected in our laboratory **(Table 2-1)**. The d-dimensional latent signals were first smoothed using a Gaussian

kernel (s.d.: 50 ms), and then multiplied by a N x d mixing matrix W with entries that were randomly selected from a zero-mean Gaussian distribution with unit variance. This resulted in a dataset X composed of M samples, each of them N-dimensional. We chose N=96 to reproduce the number of signals recorded by the multi-electrode array used to obtain the original experimental data. The activity in each of the N=96 simulated channels was scaled to the range from zero to one to compensate for variability in firing rates across neurons and across time. The effect of non-uniform firing rate variances across channels was considered separately (see "Effect of non-uniform variances across channels" in Results).

A nonlinear embedding was implemented by processing each simulated neural recording in X with an exponential activation function:

$$f(X) = \frac{e^{\alpha X}-1}{e^{\alpha}-1} \qquad \text{(Equation 2-1)}$$

The choice of an exponential nonlinearity was based on results from Generalized Linear Models, for which the statistics of the modeled variable determines the nonlinear link function (Nelder and Wedderburn 1972). In our case, the variables of interest are spike counts. Under the assumption that these variables follow a Poisson-like distribution, the appropriate choice of link function is the logarithm (Nelder and Wedderburn 1972). The inverse of the link function, the exponential, is the appropriate nonlinear function for relating the linear combination of explanatory covariates, the latent signals, to the variables of interest, the simulated firing rates. The exponential activation function used in our simulations allowed us to control the degree of nonlinearity by varying the single parameter α, and to ensure that the range of the nonlinearly embedded synthetic data remained between zero and one. Finally, we added independent Gaussian noise to each of the channels in X, to generate signals with known signal-to-noise ratio. This choice of noise model provides a simple and widely used mechanism for simulating stochastic processes (Cunningham et al. 2009).

The various steps in this procedure allowed us to generate datasets of known intrinsic dimensionality, embedding type (linear/nonlinear), and signal-to-noise ratio.

Figure 2-1: Generation of simulated datasets. First, latent neural signals were obtained by randomly sampling the firing rates of primary motor cortical recordings. The number of latent signals determined the intrinsic dimensionality of the dataset. Then, the dimensionality of the dataset was increased through linear combinations effected by multiplication with a weight matrix W. The entries of W were sampled from a zero-mean Gaussian distribution with unit variance. Then, the resulting signals were then scaled to the [0,1] range by dividing them by their maximum value. This procedure yielded noise-free, linear datasets. In nonlinear simulations, the signals were then activated nonlinearly using the exponential function in Equation 1 (red box in diagram). In noisy simulations, zero-mean Gaussian noise with variance specified by a predetermined signal-to-noise ratio was added to the signals. This procedure yielded linear or nonlinear, noisy datasets with known signal-to-noise ratio.

Dimensionality estimation algorithms

We evaluated two classes of dimensionality estimation algorithms, those that assumed a linear embedding and those that also allowed for a nonlinear embedding.

Linear algorithms

Linear algorithms map high-dimensional data to a lower dimensional, linear subspace. Principal Component Analysis (PCA) is often used for linear dimensionality estimation in neuroscience (Cunningham and Yu 2014; P. Gao and Ganguli 2015; Williams et al. 2018; Kaufman et al. 2014; Gallego et al. 2018; Sadtler et al. 2014). All the linear algorithms that we tested (summarized below) are based on PCA but use different criteria for dimensionality estimation.

Principal Component Analysis with a variance cutoff

PCA creates a low-dimensional representation of the data by sequentially finding orthogonal directions that explain the most remaining variance. Unit vectors that identify those directions, the PCA eigenvectors $\{v_i\}$, provide an orthonormal basis for the N-dimensional data space. The eigenvectors are labeled in decreasing order of the variance associated with each direction, given by the eigenvalues $\{\lambda_i\}$. The simplest way to use PCA for dimensionality estimation is to find the number of principal components required to reach a predetermined threshold of cumulative variance. The selection of a variance threshold can be rather arbitrary, and a range of thresholds have been used in the literature. In this study, we used a threshold of 90%, which yielded accurate estimates of dimensionality for the noise-free linear datasets.

Participation Ratio (PR)

This approach provides a principled way of finding a variance threshold when the ground truth is not known (Mazzucato, Fontanini, and La Camera 2016; P. Gao et al. 2017). PR uses a simple formula based on the eigenvalues:

$$PR = \frac{\left(\sum_{i=1}^{N} \lambda_i\right)^2}{\sum_{i=1}^{N} (\lambda_i)^2} \qquad \text{(Equation 2-2)}$$

If the leading eigenvalue carries all the variance ($\lambda_i \neq 0$ for $i = 1$ and $\lambda_i = 0$ for all $i \geq 2$), then PR = 1. At the other extreme, if all eigenvalues are equal, the variance is spread evenly across all the dimensions, and PR=N. The actual value of PR interpolates between these two extreme conditions to estimate the intrinsic dimensionality, and thus the number of principal components to be kept (P. Gao et al. 2017).

Parallel Analysis (PA)

Much like the Participation Ratio, Parallel Analysis is a principled approach to finding a variance threshold (Horn 1965; Buja and Eyuboglu 1992). Parallel Analysis generates null distributions

for the eigenvalues by repeatedly shuffling each of the N dimensions of the data separately. The shuffling step ensures that the remaining correlations across the different dimensions of the data are due to chance. We repeated the shuffling procedure 200 times, resulting in a null distribution for each eigenvalue based on 200 samples. The eigenvalues that exceeded the 95th percentile of their null distribution were identified as significant; the number of significant eigenvalues determined the number of dimensions to be kept. Although this method has not been directly applied to neural data, similar approaches based on finding null distributions of eigenvalues have been used for neural dimensionality estimation (Machens, Romo, and Brody 2010).

Nonlinear algorithms

Nonlinear algorithms can in principle estimate the dimensionality of either linearly or nonlinearly embedded data. Unlike the linear algorithms we tested, the nonlinear algorithms need not rely on a global model for the probability distribution from which the data are assumed to be drawn (in the case of PCA, the model is a multivariate Gaussian distribution). Instead, many nonlinear algorithms estimate intrinsic dimensionality directly from local geometric properties of the data. Common local properties include distance and separability of each data point relative to its neighbors. Although nonlinear algorithms are not yet commonly used in neuroscience, they have been used to estimate dimensionality in several other fields that produce high-dimensional datasets (Campadelli et al. 2015).

Correlation Dimension (CD)

Correlation Dimension estimates dimensionality by calculating how the number of data samples that fall within a hypersphere change as a function of its radius. This method, originally developed in 1983 (Grassberger and Procaccia 1983), has benefitted from recent efforts to improve computational speed and accuracy (Kalantan and Einbeck 2012; Einbeck and Kalantana 2013). Although there are only a few applications of Correlation Dimension analysis

to neural data (Kobayashi et al. 2000; Boon et al. 2008), it is widely used in other disciplines (Kalantan and Einbeck 2012).

Levina-Bickel Maximum Likelihood Estimation (LBMLE)

The Levina-Bickel Maximum Likelihood Estimation method (Levina and Bickel 2004) is an extension of Correlation Dimension that uses a maximum likelihood approach to estimate distances between data points. This method has been successfully applied to some of the benchmark datasets used in machine learning, such as the Faces (Tenenbaum, de Silva, and Langford 2000) and Hands datasets (Kégl 2002).

Two Nearest Neighbors (TNN)

The Two Nearest Neighbors method also uses the distance between data points to estimate dimensionality (Facco et al. 2017). However, unlike Levina-Bickel Maximum Likelihood Estimation, it considers only the first and second neighbors of each point. The ratio of the cumulative distribution of second-neighbor to first-neighbor distances is a function of data dimensionality. By focusing on shorter distances, the method avoids unwanted effects resulting from density changes across the manifold. This method has been successfully applied to synthetic datasets of hyperspheres with known dimensionality (Facco et al. 2017), and to real-world datasets including molecular simulations (Pinamonti et al. 2017) and images of hand-written digits (Tenenbaum, de Silva, and Langford 2000).

Fisher Separability Analysis (FSA)

High-dimensional datasets exhibit simple geometric properties such as the likely orthogonality of two randomly picked directions. These properties have recently been characterized as the *blessings of dimensionality* (Alexander N. Gorban, Makarov, and Tyukin 2020), in contrast to the well-known concept of the *curse of dimensionality*. A useful example is the increasing ease with which a hyperplane can separate any given sample in a dataset from all other samples as the

dimensionality of the dataset increases. Fisher separability is a computationally efficient, simple, and robust method to assess such separability (Fisher 1936; A. N. Gorban and Tyukin 2018). Dimensionality can be estimated in terms of the probability that a point in the dataset is Fisher separable from the remaining points (Albergante, Bac, and Zinovyev 2019). The probability distribution of Fisher separability allows the dimensionality of both linear and nonlinear manifolds to be estimated. This method has been applied to study the mutation profiles of the genes resulting in tumors as a means to evaluate therapeutic approaches (Le Morvan, Zinovyev, and Vert 2017).

Denoising algorithms

Noise that is uncorrelated across channels will lead to dimensionality estimates that approach the number of channels as the level of noise increases. To mitigate this overestimation problem, we implemented two approaches to denoise neural data. Both rely on an initial estimate of an upper bound dimensionality D, for which we used Parallel Analysis. To quantify the performance of the denoising algorithms, we reported Variance Accounted For (VAF) between the denoised signals and the noise-free signals, the latter providing the ground truth.

PCA denoising

The linear approach to denoising was based on PCA. Once the value of D was determined, we used the D leading principal components to reconstruct the original data. PCA-based denoising is based on the assumption that most of the noise is relegated to the discarded, low-variance principal components.

Joint Autoencoder denoising

We also used a neural network for denoising **(Figure 2-2**). For this purpose, we divided the 96-dimensional simulated dataset X into two 48-dimensional partitions: X_1 and X_2. These partitions were each mapped by the compressive halves of the respective autoencoders to the D-

dimensional subspaces Z_1 and Z_2 respectively. These compressed subspaces were used to obtain reconstructed versions of X_1 and X_2, respectively denoted $\hat{X}_1$ and $\hat{X}_2$, using the expansive halves of the corresponding autoencoders. The cost function C used to train the Joint Autoencoder network not only minimized the reconstruction error for X_1 and X_2, but also the difference between Z_1 and Z_2:

$$C = MSE\left(X_1, \hat{X}_1\right) + MSE\left(X_2, \hat{X}_2\right) + MSE\left(Z_1, Z_2\right) \qquad \text{(Equation 2-3)}$$

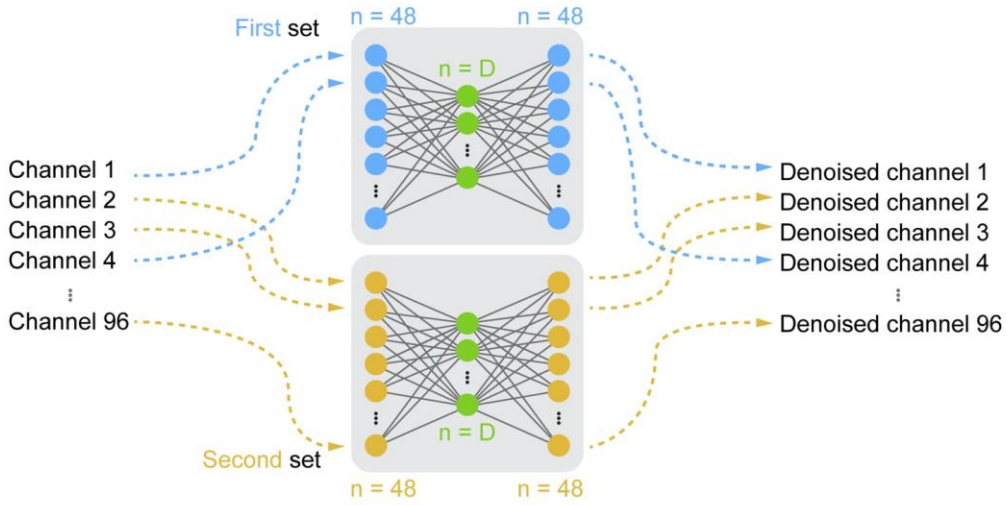

Figure 2-2: Architecture of the Joint Autoencoder. Channels of the 96-dimensional simulated datasets were randomly partitioned into two sets of signals (blue and yellow). Each 48-dimensional set was reconstructed through the corresponding D-dimensional subspace, Z_1 and Z_2 (green). The reconstructed outputs of the networks were the denoised channels.

This design assumes that each of the partitions X_1 and X_2 contains the information necessary to robustly identify the underlying D-dimensional signals Z_1 and Z_2, but not the independent noise components that will differ between the two partitions. We trained the Joint Autoencoder using the ADAM optimizer with a learning rate $\eta = 0.001$ and dropout regularization on the input layer with $p = 0.05$. The use of Rectified Linear Unit (ReLU) activation functions in all layers ensured

that the autoencoder network would both operate on and output non-negative signals while allowing for nonlinear embeddings. Our choice of using the ReLU activation function was motivated by its documented success in modeling a wide variety of nonlinearities for deep learning applications (Glorot, Bordes, and Bengio 2011; LeCun, Bengio, and Hinton 2015). In addition, the strict nonnegativity of the ReLU function mimics that of real neural recordings.

Ethics statement

All surgical and experimental procedures that yielded the multi-electrode array recordings from non-human primates (Gallego et al. 2018), which formed the basis of our simulated neural signals, were approved by Institutional Animal Care and Use Committee (IACUC) of Northwestern University. The subject was monitored daily. The subject's diet consisted of standard laboratory animal diet, fresh fruits, and vegetables, and was provided with access to various types of enrichment.

Statistical analyses

We used Monte Carlo simulations to generate up to 10 replications of synthetic data sets, each corresponding to microelectrode array recording data from an experimental session. We noted the number of replications (n) in the figure captions where applicable. Our choice of the number of replications is reasonable compared to the number of experimental sessions that we would expect to see in experiments with monkeys (Gallego et al. 2018; Russo et al. 2018; Perich, Gallego, and Miller 2018). The simulations differed by their random number generator seed, which dictated the pseudorandom sampling procedures required for generating the signals. There were three sampling steps in our simulations **(Figure 2-1)**. First was the creation of the low-dimensional latent signals, which were sampled from an empirical firing rate distribution. The second was the entries of the mixing matrix W, which were sampled from a zero-mean Gaussian distribution with unit variance. The third was the additive noise, sampled from a zero-

mean Gaussian distribution with variance determined by the specified signal-to-noise ratio. We used bootstrapping with 10,000 iterations to compute the statistic of interest and computed its confidence interval using $\alpha = 0.05$. We used Bonferroni correction for multiple comparisons.

Results

Despite the large number of available algorithms for dimensionality estimation, there has been no systematic study of how well-suited they are for the analysis of neural data. Here we test several representative algorithms on synthetic datasets for which the intrinsic dimensionality is known, to assess their ability to estimate the true dimensionality of the data across a range of simulated conditions relevant to neuroscience. These assessments resulted in a recommended procedural pipeline for estimating the intrinsic dimensionality of a set of neural recordings.

Dimensionality of noise-free datasets

We first considered the simplest case: how accurately can we determine the dimensionality of linearly embedded, noise-free datasets? To answer this question, we applied the six selected algorithms to datasets with dimensionality $d = 6$. We focused on $d = 6$ as this was the dimensionality estimate of actual multi-electrode array recordings found when using the methods investigated here. In this scenario, all tested linear and nonlinear algorithms estimated the true dimensionality accurately **(Figure 2-3)**. Under noise-free conditions, the nonlinear algorithms were as accurate as the linear ones on linearly embedded datasets.

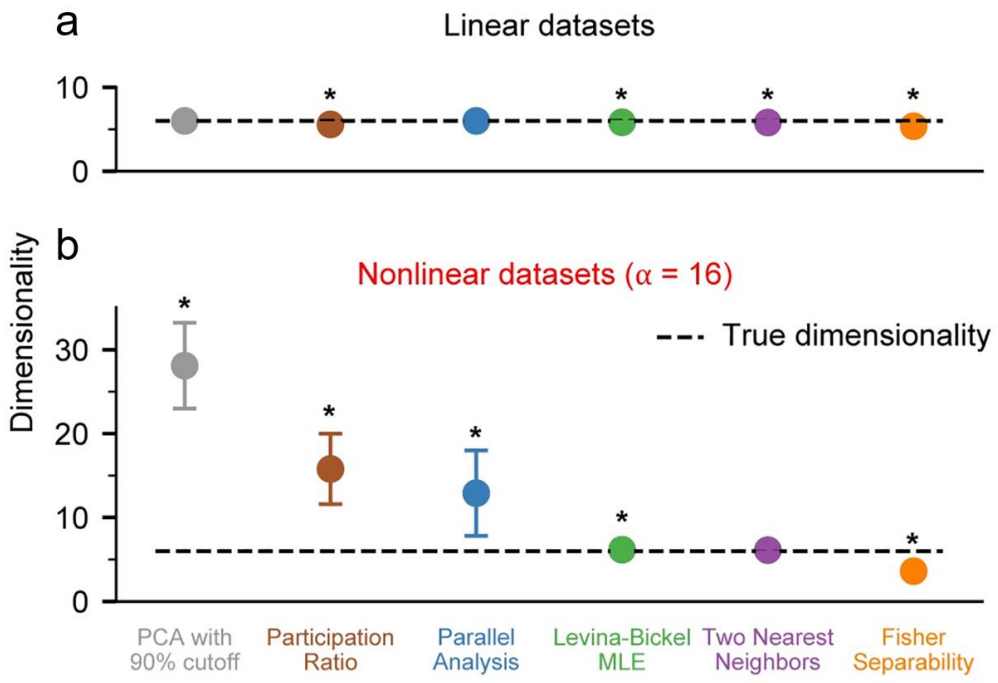

Figure 2-3: Dimensionality of noise free datasets. a) We applied PCA with 90% variance cutoff (PCA90, gray), Participation Ratio (PR, brown), Parallel Analysis (PA, blue), Levina-Bickel Maximum Likelihood Estimation (LBMLE, green), Two Nearest Neighbors (TNN, purple), and Fisher Separability Analysis (FSA, orange) to linearly embedded, $\mathrm{d} = 6$ datasets (n=10). b) Same as in a, but for nonlinearly embedded datasets. Circles indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant difference of the mean from the true dimensionality of 6 at the significance level of α=0.05.

Next, we evaluated all algorithms on nonlinearly embedded noise-free datasets, also for $d = 6$. Nonlinearities were introduced as in Equation 1, using $\alpha = 16$. In this case, the three linear algorithms dramatically overestimated the true dimensionality, with errors reaching more than 400% of the true value **(Figure 2-3b)**. In contrast, the nonlinear algorithms performed well; the Levina-Bickel Maximum Likelihood Estimation and the Two Nearest Neighbors methods were more accurate than Fisher Separability Analysis, which slightly underestimated the true dimensionality.

Effect of non-uniform variances across channels

The normalization of our simulated datasets restricted channel activity to the [0,1] interval, thus imposing a large degree of variance similarity across channels. In contrast, variances of real neural recordings can vary as much as 10-fold from channel to channel. To evaluate the performance of the dimensionality estimation algorithms considered here in the presence of non-uniform variances across channels, we scaled each channel of simulated neural data by a randomly chosen real number between 1 and 10. We found that most algorithms yielded lower estimates of dimensionality when applied to the rescaled data in comparison to the estimates obtained when the algorithms were applied to the data before rescaling **(Figure 2-4a and b)**. However, note that both Levina-Bickel Maximum Likelihood Estimation and the Two Nearest Neighbors yielded remarkably accurate dimensionality estimates when applied to rescaled data.

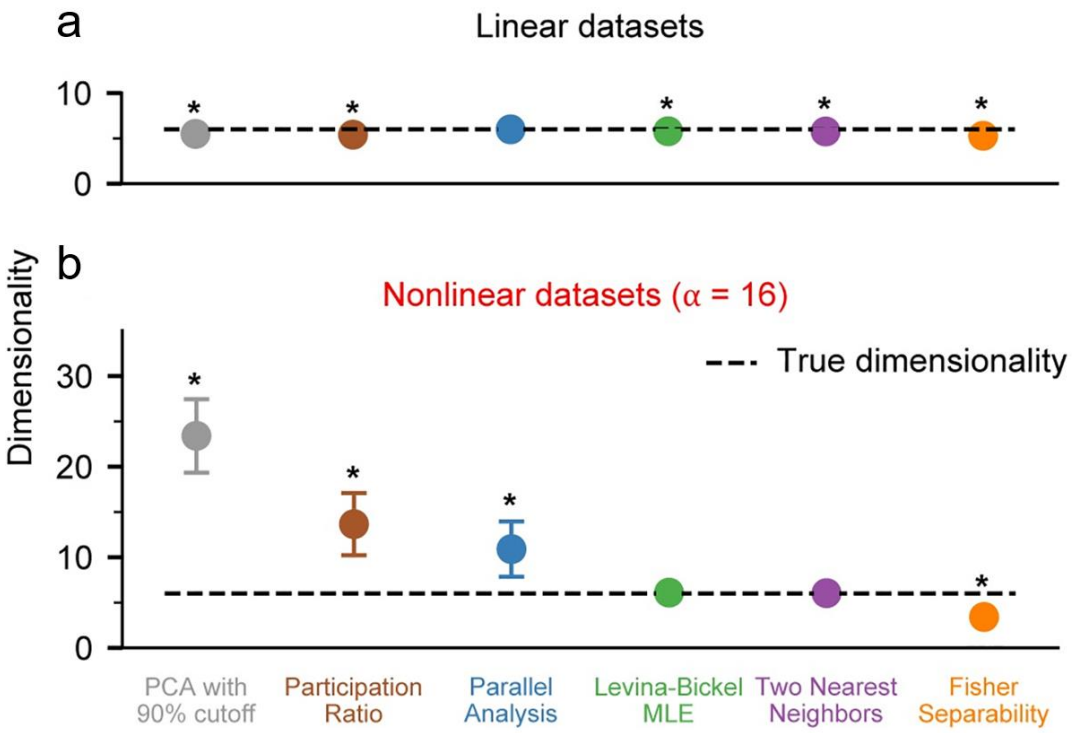

Figure 2-4: Dimensionality of noise free datasets with unequal variance. a) We applied PCA with 90% variance cutoff (PCA90, gray), Participation Ratio (PR, brown), Parallel Analysis (PA,

blue), Levina-Bickel Maximum Likelihood Estimation (LBMLE, green), Two Nearest Neighbors (TNN, purple), and Fisher Separability Analysis (FSA, orange) to linearly embedded, d = 6 datasets (n=10) after randomly scaling each of the N=96 channels. b) Same as in a, but for nonlinearly embedded datasets. Circles indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant differences of the mean from the true dimensionality of 6 at the significance level of α=0.05.

Because of the superior accuracy of Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors, we focused on these two methods for the remainder of the nonlinear analyses. We also retained Parallel Analysis as a benchmark for some of the analyses, as it was the most accurate linear method for estimating the dimensionality of nonlinearly embedded data.

Effect of true dimensionality on algorithm accuracy

We next evaluated how the true intrinsic dimensionality of the noise-free data influenced algorithm accuracy. Can any intrinsic dimensionality be reliably estimated? We found that the answer is no: the accuracy of all algorithms suffered when the intrinsic dimensionality of the synthetic data was too high. Parallel Analysis was accurate on linear datasets with $d < 20$, but inaccurate on nonlinear datasets of all dimensions, as expected **(Figure 2-5)**. Below about d = 6, Levina-Bickel Maximum Likelihood Estimation and Two-Nearest Neighbors were accurate on both linear and nonlinear datasets. However, Levina-Bickel Maximum Likelihood Estimation began to underestimate the dimensionality of both linearly embedded **(Figure 2-5a)** and nonlinearly embedded **(Figure 2-5b)** datasets for $d > 6$. This underestimation increased with increasing d. For nonlinear datasets, the estimate saturated at d = 13, where underestimation began to get much worse. These results revealed that the intrinsic dimensionality of nonlinearly embedded datasets is hard to estimate reliably when it is large.

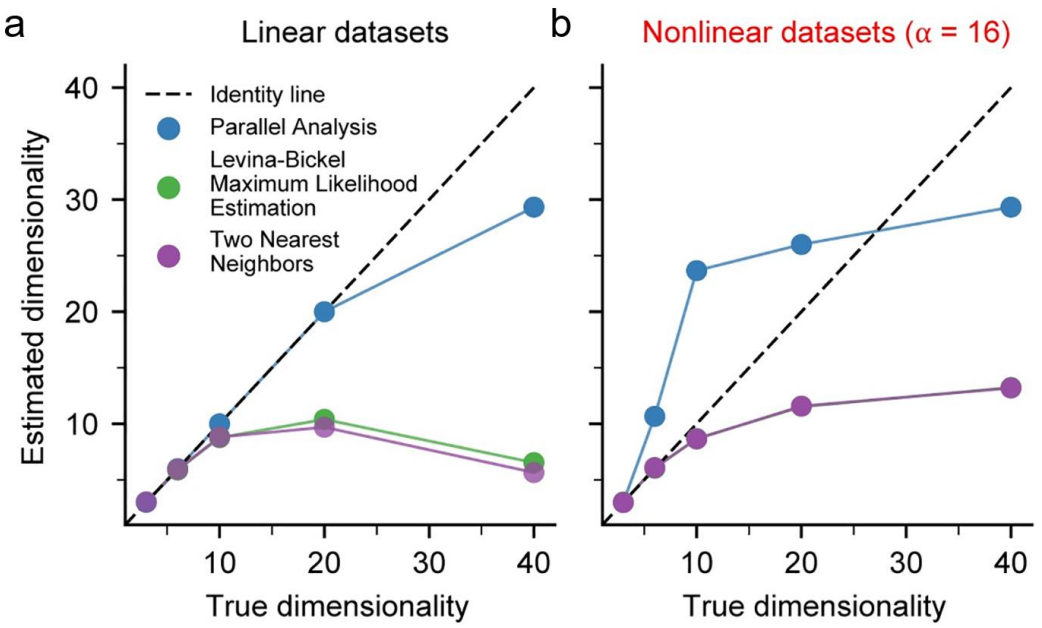

Figure 2-5: Effect of increasing true dimensionality on dimensionality estimates. a) The dimensionality of noise free, linear datasets (n=3) was assessed using Parallel Analysis (PA), Levina-Bickel Maximum Likelihood Estimation (LBMLE), and Two Nearest Neighbors (TNN). Dashed line indicates the identity line. b) Same as a, but for nonlinear datasets. The curve for TNN precisely overlays that of LBMLE, causing it to be obscured.

Effect of the level of nonlinearity

We next evaluated how the degree of nonlinearity influenced the accuracy of the dimensionality estimation algorithms. We controlled the degree of nonlinearity by varying the parameter α in Equation 1; this parameter controls the slope of the exponential activation function used to generate the nonlinearly embedded datasets. We found that both Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors provided accurate dimensionality estimates for all tested levels of nonlinearity **(Figure 2-6**). Surprisingly, even Parallel Analysis was accurate up to levels of nonlinearity corresponding to $\alpha \approx 8$, where it started to overestimate the intrinsic dimensionality. These results revealed that Levina-Bickel Maximum Likelihood

Estimation and Two Nearest Neighbors provide accurate dimensionality estimates for wide levels of nonlinearity, whereas Parallel Analysis is accurate only for low levels of nonlinearity.

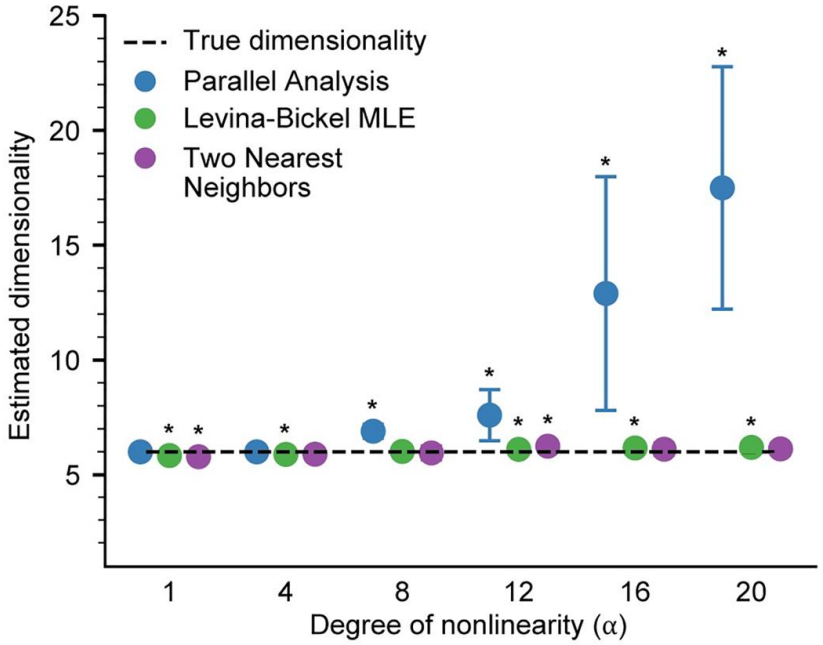

Figure 2-6: Effect of changing the degree of nonlinearity. Dimensionality of nonlinear datasets (n=10) with varying levels of nonlinearity, controlled by the α parameter (See Methods), was assessed using Parallel Analysis (PA), Levina-Bickel Maximum Likelihood Estimation (LBMLE), and Two Nearest Neighbors (TNN). Circles indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant differences of the mean from the true dimensionality of 6 at the significance level of α=0.05.

Amount of data required for estimating dimensionality

Ideally, algorithms would require only small amounts of data, so that the intrinsic dimensionality could be estimated even during transient behaviors and for a small number of recording channels. We thus evaluated the amount of data required to estimate the dimensionality of datasets with d = 6, by varying both the number of samples M and the number of recording channels N.

On linear datasets, the accuracy of Parallel Analysis depended only on the number of channels: the algorithm was accurate if 20 or more channels were available **(Figure 2-7a**). In contrast, the accuracy of both Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors also depended on the number of samples **(Figure 2-7b and c**). Around M = 600, requiring about 30 seconds of data binned at 50 ms, was sufficient for accurate estimates of intrinsic dimensionality using either of these two nonlinear methods.

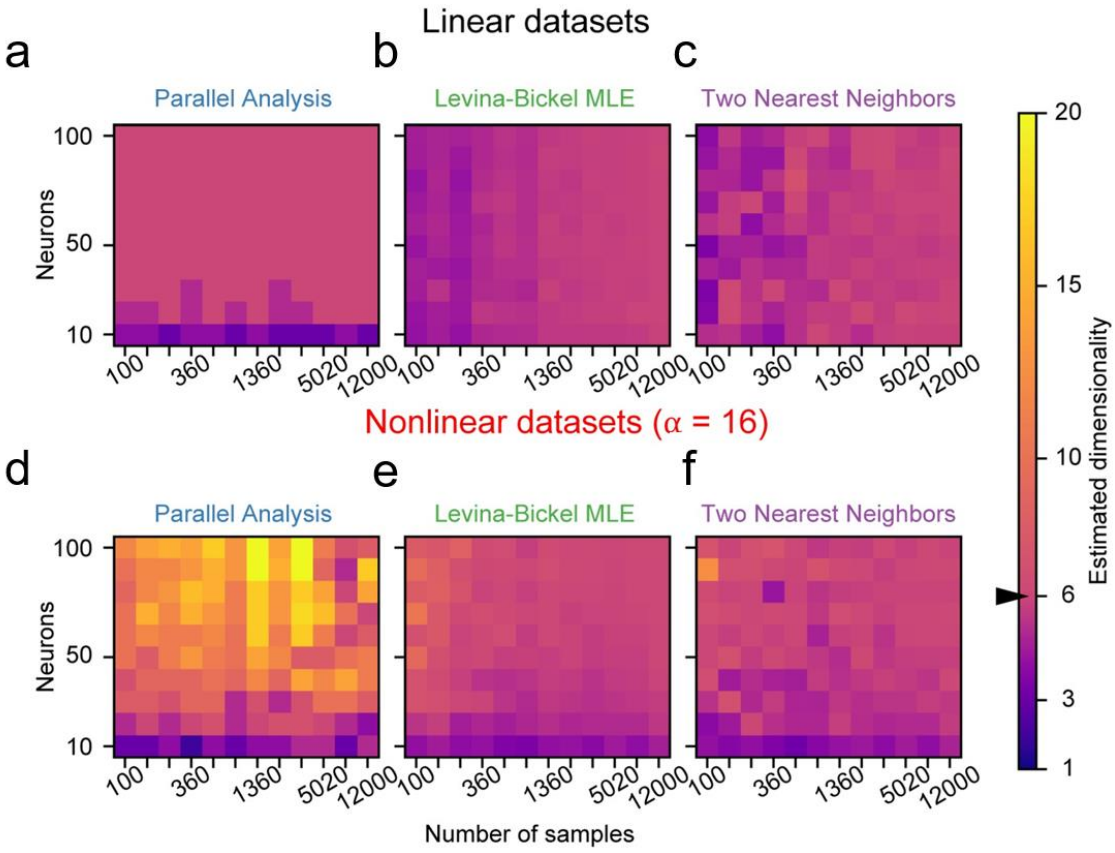

Figure 2-7: Amount of data required by dimensionality estimators. Amount of data required by a) Parallel Analysis (PA), b) Levina-Bickel Maximum Likelihood Estimation (LBMLE), and c) Two Nearest Neighbors (TNN) on linear datasets. Data length was logarithmically scaled between M = 100 and M = 12,000 samples. The correct dimensionality d = 6 is shown in pink. Light colors indicate overestimation and dark colors indicate underestimation of dimensionality. d, e, and f) Same as a, b, and c, respectively, but for nonlinear datasets.

As expected for highly nonlinear datasets (α = 16, d = 6), Parallel Analysis was not accurate **(Figure 2-7d**) regardless of the amount of data. Both Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors were accurate provided that data from more than 50 channels were available **(Figure 2-7e and f**). Furthermore, while Levina-Bickel Maximum Likelihood estimation required around 600 samples of data for accurate dimensionality estimates, Two Nearest Neighbors required more than twice as many samples. These results would also depend on the actual dimensionality d; here we focused on $d = 6$.

Evaluating and reducing the effects of noise

Any experiment will include some amount of noise in the recorded signals. As expected, all tested algorithms overestimated intrinsic dimensionality in the presence of noise **(Figure 2-8**). For any given noise level, estimation errors for the linear datasets **(Figure 2-8a)** were a bit smaller than those for the nonlinear datasets **(Figure 2-8b)**. Adding noise with a power of only 1% of that of the signal (SNR = 20 dB) caused Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors to overestimate the dimensionality of the nonlinear data by ~200% **(Figure 2-8b)**. PA yielded consistent overestimation errors across all nonzero levels of noise for both linear and nonlinear data.

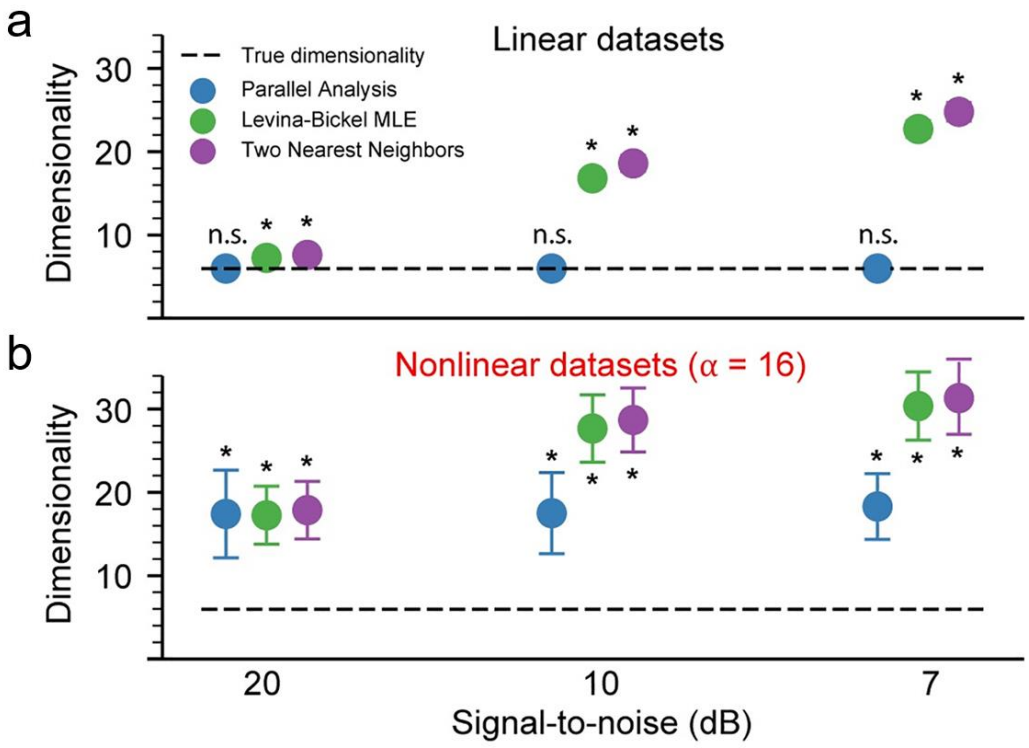

Figure 2-8: Effect of noise on dimensionality estimates. Estimated dimensionality of linear (a) and nonlinear (b) datasets (n=10) with 20 dB, 10 dB, and 7 dB signal-to-noise ratio was assessed using Parallel Analysis (PA), Levina-Bickel Maximum Likelihood Estimation (LBMLE), and Two Nearest Neighbors (TNN). Circles indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant differences of the mean from the true dimensionality of 6 at the significance level of α=0.05.

We evaluated two algorithms for mitigating the effects of noise prior to estimating dimensionality: a PCA-based linear method and a Joint Autoencoder nonlinear neural network (see Methods). Both methods were quite effective for denoising the linear datasets **(Figure 2-9a)**, with the PCA-based approach slightly better than the Joint Autoencoder at the higher noise levels. For linear datasets, dimensionality estimates following PCA-based denoising were highly accurate, yielding correct estimates of the true intrinsic dimension even for high-noise signals **(Figure 2-9b)**. The Joint Autoencoder was significantly more effective for denoising the

nonlinear datasets **(Figure 2-9c)**. Joint Autoencoder denoising on nonlinear datasets resulted in dimensionality estimates that still increasingly overestimated with increasing noise, but at a much slower rate than without denoising **(Figure 2-9D)**. The highest noise level we tested (20%; SNR = 7 dB) caused the dimensionality to be overestimated by about 100%**.** These results were consistent for different degrees of nonlinearity. The more nonlinear the data, the more appropriate it was to use the Joint Autoencoder for denoising.

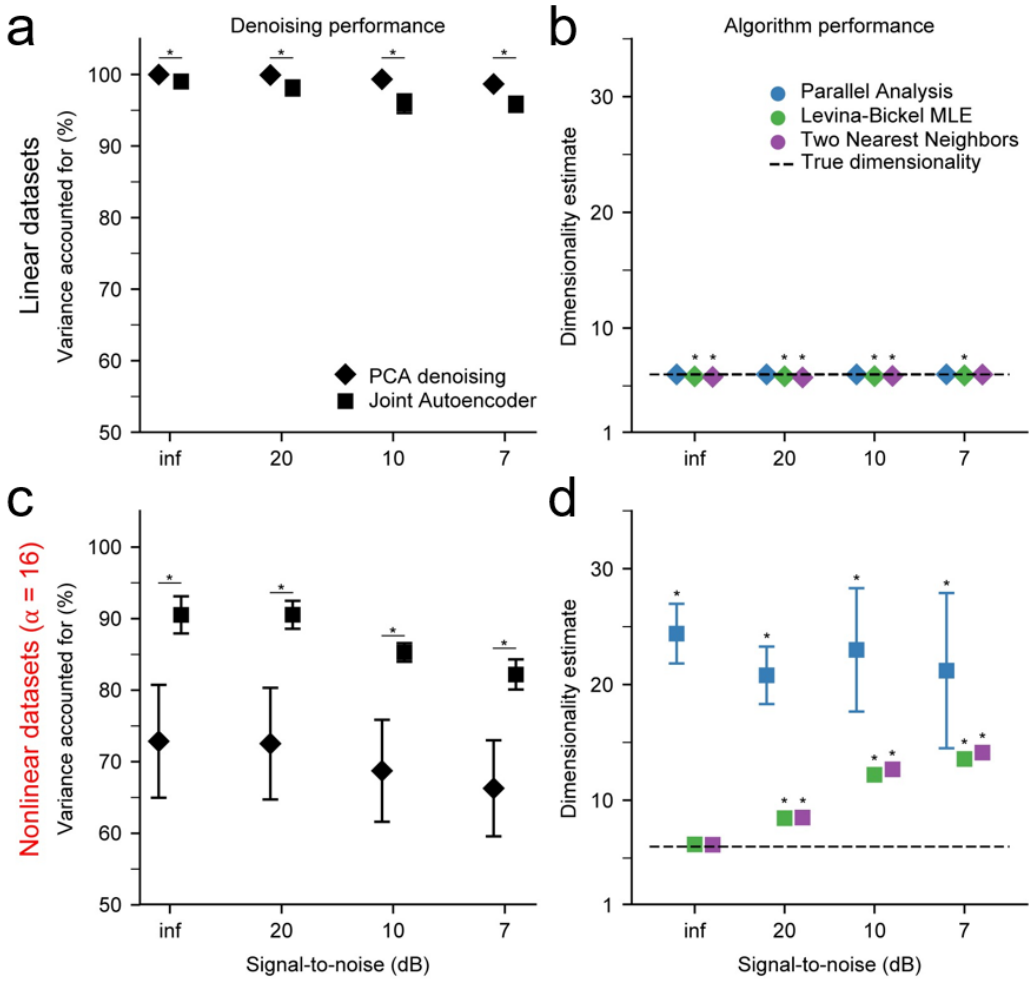

Figure 2-9: Performance of PCA and Joint Autoencoder (JAE) denoising algorithms. a) PCA and JAE denoising applied to linear datasets (n=10) with varying signal-to-noise ratio. Symbols indicate the mean and error bars indicate the standard deviation of the VAF between noise-free and denoised signals. Asterisks indicate significant difference between mean values at the

significance level of α=0.05. b) Dimensionality estimation on linear datasets after PCA denoising. Dimensionality was estimated using Parallel Analysis (PA), Levina-Bickel Maximum Likelihood Estimation (LBMLE), and Two Nearest Neighbors (TNN). Symbols indicate the mean and error bars indicate the standard deviation of the dimensionality estimates. Asterisks indicate significant differences of the mean from the true dimensionality of 6 at the significance level of α=0.05. c) Same as in a, but for nonlinear datasets. d) Same as in b, but for nonlinear datasets after JAE denoising.

Discussion

This study evaluated techniques for estimating the intrinsic dimensionality of high-dimensional neural recordings. We considered representative linear and nonlinear algorithms, testing their performance on synthetic datasets that captured properties of neural recordings likely to affect dimensionality estimation. The tested datasets had known intrinsic dimensionality, known levels of noise, and embeddings that were either linear or nonlinear. Our results demonstrated that none of the tested algorithms work for all possible scenarios, but they yielded important insights for when estimates of intrinsic dimensionality are likely to be valid and when they are not. As expected, we found that linear estimation methods are generally not as accurate as nonlinear methods when the mapping between the low-dimensional latent space and the high-dimensional space of neural recordings is nonlinear. Surprisingly, the linear method Parallel Analysis estimated the dimensionality of mildly nonlinear datasets well though it failed for more highly nonlinear embeddings. In contrast, the nonlinear methods worked well on both linear and highly nonlinear datasets but failed once the intrinsic dimensionality of the data became too high.

Noise was a challenge for all methods, causing dimensionality to be overestimated even for signal-to-noise ratios as low as 20 dB (1% noise variance). We presented two approaches for

denoising the data so as to improve the accuracy of the dimensionality estimation. These were a linear PCA-based approach and a novel nonlinear, deep learning approach that we call the Joint Autoencoder. Both denoising approaches attempted to remove signal components that were not shared across the data channels. To achieve this, the PCA-based approach simply removed Principal Components with low variance, whereas the Joint Autoencoder identified an underlying manifold that was common to two randomly sampled sets of channels. Both approaches relied on a linear, upper-bound estimate of the intrinsic dimensionality. Denoising by either method substantially improved subsequent dimensionality estimation, but the Joint Autoencoder was substantially more effective in denoising nonlinear datasets. For linear datasets, dimensionality estimates using Parallel Analysis, Levina-Bickel Maximum Likelihood Estimation, and Two Nearest Neighbors were accurate after PCA-denoising. In the nonlinear case, dimensionality estimates using the same three methods were similarly accurate after JAE-denoising.

Implications for evaluation of experimental recordings

Due to its computational efficiency and ease of interpretation, most studies have used PCA with an arbitrary variance cutoff to estimate the dimensionality of M1 neural recordings (Gallego et al. 2018; Kaufman et al. 2014; Sadtler et al. 2014; P. Gao and Ganguli 2015; Williamson et al. 2016). While we have shown that some of the linear methods can be quite effective, simply eliminating non leading PCs based on a cumulative variance cutoff was the least accurate of the algorithms that we tested. Parallel Analysis, the most accurate linear method, performed as well or even better than some of the more advanced and computationally demanding nonlinear methods. Therefore, PA should suffice as a quick and effective approach to estimating dimensionality, even for mildly noisy and nonlinear datasets.

Some of the linear methods used in neuroscience studies rely on a structure of repeated trials in the data (Machens, Romo, and Brody 2010; Williams et al. 2018). These methods use the

regularity of repeated trials in a supervised scenario to identify neural dimensions associated with specific experimentally controlled conditions. Such supervised methods cannot be applied to data obtained during non-stereotyped, non-repeating behaviors. All of the methods that we assess in this study are unsupervised and thus applicable to datasets with no repeated trial structure.

Despite the simplicity of linear algorithms, estimating the dimensionality of nonlinear manifolds requires nonlinear algorithms. There is some evidence that neural manifolds may be nonlinear. Recent studies have shown that nonlinear methods for inferring behavioral parameters from M1 neural manifolds are superior to linear methods (Sussillo et al. 2016; Pandarinath, O'Shea, et al. 2018; Pandarinath, Cora Ames, et al. 2018; Farshchian et al. 2018). This suggests that the underlying neural manifold representing motor intent may be nonlinear, and that linear dimensionality estimation methods may be inadequate when estimating the intrinsic dimensionality of primary motor cortical recordings. Studies that investigated the dimensionality of M1 using linear methods most likely overestimated its true intrinsic dimensionality.

Nonlinear algorithms were more accurate than linear algorithms for nonlinear datasets of dimensionality below 10. However, nonlinear methods underestimated dimensionalities above 10. This is a critical concern for experimental recordings, since a low dimensionality estimate from a nonlinear method might be inaccurate if the true dimensionality were large. Multiple studies using linear methods have reported an estimated dimensionality of M1 of around 10 for simple, well-practiced behaviors (Sadtler et al. 2014; Perich, Gallego, and Miller 2018; P. Gao et al. 2017). Our results show that linear methods provide an upper bound to the estimate of intrinsic dimensionality as long as the true dimensionality of the data is below 20. If the intrinsic dimensionality of M1 is substantially higher for more dexterous use of arm and hand than for the scenarios that have typically been studied, the nonlinear methods investigated here may underestimate it.

One method for addressing this concern would be to use nonlinear methods to reduce the dimensionality of a dataset to that of its nonlinear dimensionality estimate, and then to assess the amount of variance that the nonlinear low-dimensional representation captures. If the VAF is high, the data may be truly nonlinearly low dimensional. If, on the other hand, the VAF is low, the true intrinsic dimensionality could be higher than estimated. For the latter case, a practical approach would be to report only the linear dimensionality estimate and emphasize that it only provides an upper bound to the true dimensionality.

We currently lack techniques for reliably assessing datasets with high intrinsic dimensionality, at least when considering practical situations with limited data. There have been some theoretical studies of the amount of data needed for accurate estimation of dimensionality (L. A. Smith 1988; Eckmann and Ruelle 1992). Correlation Dimension, the method on which many nonlinear algorithms are based, requires that the number of data samples M be on the order of $10^{d/2}$ (Camastra and Vinciarelli 2002). The total amount of data can be increased by either recording from more channels or for a longer duration. Studies that investigated the dimensionality of the primary visual cortex (V1) found that the eigenvalue spectrum of the neural signals obtained from approximately one thousand neurons decayed as a power law (Stringer, Pachitariu, Steinmetz, Carandini, et al. 2019; Stringer, Pachitariu, Steinmetz, Reddy, et al. 2019). These findings would not have been possible if recording from a hundred neurons, which would not have revealed the long, slow-decaying tail of the eigenvalue distribution. One interpretation of these findings is that the linear dimensionality of V1 is arbitrarily large. However, an alternative interpretation is that the neural data are embedded in a very nonlinear manifold, causing the intrinsic dimensionality to be overestimated by the linear methods used in these studies.

The stochastic nature of neural firing and the noise associated with experimental measurements will also cause the intrinsic dimensionality to be overestimated. The two denoising approaches that we presented are simple and effective. Depending on the assumptions about the underlying

structure of firing patterns, alternative denoising approaches may be useful. For example, if the temporal relationship between the firing patterns of the population neural activity is of interest, one could use denoising methods that explicitly attempt to model these dynamics, such as Latent Factor Analysis through Dynamical Systems (LFADS), prior to estimating the dimensionality (Pandarinath, O'Shea, et al. 2018).

For the past five decades since the time of Evarts' early experiments (Evarts 1966), assessing the relationship between behaviors and single neuron signals recorded from the brain has been a mainstay of motor systems research. Although the focus of our study was on the dimensionality of the neural manifold to which the population activity is confined, the natural next step in the analysis is to investigate the dynamics of the signals within the neural manifold and their relation to behavior.

Limitations of the study

While we tried to replicate essential features of experimental data, there are certain characteristics that we did not try to model in our simulations. For example, we only considered additive Gaussian isotropic noise, for simplicity. Experimental recordings might include non-additive, non-isotropic, or non-Gaussian noise. In such cases, PCA may not be an appropriate approach to denoising, even for linearly embedded data. Methods such as factor analysis or extensions such as Gaussian-Process Factor Analysis (Yu et al. 2009), and preprocessing steps such as square-root transforms or pre-whitening could be used instead.

We scaled the firing rates of each channel to be in the [0,1] range. The arbitrary scaling of firing rates provided a simple means for the nonlinear datasets to have the same range as their linear counterparts, as the activation function that we used mapped the [0,1] range onto itself. However, this modeling restriction does not reflect experimental neural firing data, since the range of neural firing can differ significantly even across neurons of the same type. We have

illustrated how heterogeneity in the range of firing rates affects the reliability of dimension estimation algorithms. Soft-normalization approaches that are commonly used in neuroscience (dividing a neuron's firing rate by its range plus a small constant, e.g. (Churchland et al. 2012)), would results in the amplification of signals with low variance and would cause variance-based algorithms to result in higher dimensionality estimates.

The latent signals used to generate simulated firing rates have the same first-order statistics as the actual data from which they were sampled. This unrealistic scenario is addressed through the random rescaling of individual simulated channels, to reflect the heterogeneity in the range of firing rates observed in actual neural recordings. The latent signals corresponding to the rescaled data no longer share common statistics. This scenario allows us to address an important problem: that low-variance latent signals can be informative (Yan et al. 2020). As demonstrated in our study, the use of nonlinear methods for dimensionality estimation ameliorates the problems that arise when neglecting low-variance signals as purely noisy.

Recommended analysis pipeline

Based on our results, we recommend the following approach for estimating the dimensionality of neural recordings **(Figure 2-10)**. First, obtain an upper-bound estimate D of the intrinsic dimensionality of the data. We found that Parallel Analysis works well for this purpose, being both computationally efficient and the most accurate linear method in our tests. Next, the signals should be denoised. Our denoising approach worked by projecting the neural signals into a subspace of dimensionality D equal to the upper-bound dimensionality estimate, and then reconstructing them based on these projections. A PCA based reconstruction is easy to implement and interpret and may be preferable if computational efficiency is important. A nonlinear denoising algorithm, such as the Joint Autoencoder we proposed, should also be used to assess the nonlinearity of the manifold. The usefulness of the denoising step was quantified through the VAF between the reconstructed signals, assumed to be denoised, and the noise-

free synthetic signals before noise was added to them. Our results showed that for nonlinear datasets this VAF was higher for the Joint Autoencoder than it was for PCA. However, this VAF cannot be computed for experimental data, for which we do not have access to the noise-free signals. In this scenario, the reconstruction VAF between noisy inputs and the denoised reconstructed outputs may be useful for detecting nonlinear manifolds: a higher reconstruction VAF for Joint Autoencoder denoising than for PCA denoising would signal a nonlinear manifold. A reconstruction VAF that prefers the Joint Autoencoder indicates that this denoising method yields better denoised signals. Once the signals are denoised, and the linearity or nonlinearity of the manifold is established, either a linear or nonlinear dimensionality estimation method should be used depending on the comparative performance of the corresponding denoising algorithms. The most accurate linear method we tested was Parallel Analysis. Of the nonlinear methods, Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbors were the most accurate; Levina-Bickel Maximum Likelihood Estimation required fewer data samples.

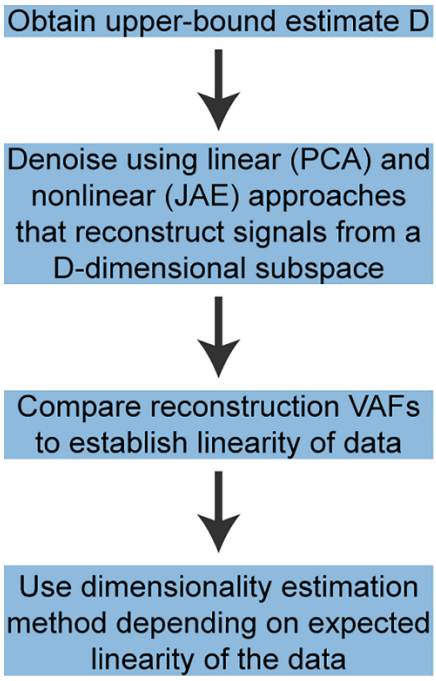

Figure 2-10: Recommended analysis pipeline for estimating the dimensionality of multi-electrode array recordings. First, obtain the upper-bound dimensionality estimate D using a

linear algorithm. Parallel Analysis works well for this purpose. Next, denoise the data using both linear (PCA based) and nonlinear (JAE based) denoising approaches and compare their reconstruction VAFs. Higher VAF for the PCA based denoising would signal a linear manifold. In contrast, higher VAF for the JAE based denoising would signal a nonlinear manifold. Finally, once the signals have been denoised using the appropriate denoising method based on the determined linearity of the manifold, estimate the dimensionality of the denoised signals. Parallel Analysis is appropriate for linear manifolds. Levina-Bickel Maximum Likelihood Estimation and Two Nearest Neighbor are the most accurate nonlinear algorithms that we tested.

Conclusions

Estimating the dimensionality of neural data is challenging. In this study, we tested several available algorithms and determined the conditions under which estimating dimensionality may be particularly difficult or even impractical. Noise is a confounding factor and must be removed prior to dimensionality estimation. Most existing studies have estimated intrinsic dimensionality using linear methods that are computationally efficient and easy to interpret. We showed that linear methods provide an upper-bound to the intrinsic dimensionality, and in cases of high noise, may even provide better estimates than nonlinear methods, although neither linear nor nonlinear methods will yield accurate estimates in this scenario. Nonlinear algorithms were more accurate for nonlinear datasets when noise was adequately removed. Finally, algorithms failed when the intrinsic dimensionality was high. It may be impractical or impossible to estimate the dimensionality of neural data when it is above ~20. However, estimation of the dimensionality of neural activity in the primary motor cortex may be possible, as many studies have reported its linear dimensionality to be within the practical limits for accurate estimation by the methods we tested.

Acknowledgments

The authors would like to thank Juan Á. Gallego for collecting the experimental data which served as the basis for our simulations.

Supporting Information

In order to apply our recommended analysis pipeline to actual neural data, we used neural activity recorded from the primary motor cortex (M1) of a macaque monkey (Monkey J) during the execution of an isometric center-out task (see (Gallego et al. 2018) for details). We used three datasets, J1, J2, and J3, corresponding to three different experimental sessions. The recorded neural data consisted of 96-channels of spiking activity, binned at 50 ms and smoothed using a Gaussian kernel with zero mean and a s.d of 50 ms to obtain firing rates. We selected five-minute segments of concatenated, successful trials taken from each of the three datasets **(Table 2-1)**. This is the actual neural data to which we applied the recommended analysis pipeline.

The analysis pipeline, described in the main body of the paper (see “Recommended analysis pipeline” in “Discussion”), consists of three steps: 1) Obtaining the upper bound dimensionality, 2) denoising, and 3) estimating dimensionality. We chose Parallel Analysis (PA) for Step 1. The estimated upper bound dimensionalities are shown in the Table below. The next step was to denoise the datasets using PCA and JAE, by reconstructing the data after compression through a low-dimensional bottleneck. The VAF between the original and reconstructed data using JAE-based denoising was consistently higher than that obtained when using PCA-based denoising, signaling nonlinearity in all three datasets. The final step was to apply the two nonlinear dimensionality estimators, Levina-Bickel Maximum Likelihood (MLE) and Two Nearest Neighbors (TNN) to the denoised datasets. The MLE and TNN estimates are shown in the Table below. Based on these dimensionality estimates, we chose an intrinsic dimensionality d=6 for most of the simulated neural activity.

	Step 1 Upper bound dimensionality	Step 2 VAF after linear and nonlinear denoising		Step 3 Estimate dimensionality	
Dataset	**PA**	**PCA VAF**	**JAE VAF**	**MLE**	**TNN**
J1	11	53%	61%	6.8	5.8
J2	11	51%	59%	6.7	5.5
J3	9	56%	62%	4.8	4.6

Table 2-1: Application of the recommended analysis pipeline to three sets of real neural recordings. The parallel analysis (PA) estimates of the dimensionality are shown for each of the datasets J1, J2, and J3. These values determined the bottleneck dimensionality to be used for denoising each dataset. The PCA-based denoising yielded reconstructions with 53%, 51%, and 56% VAF. The JAE-based denoising was slightly better for all datasets, with 61%, 59%, and 62% VAF. The better performance of the JAE-based denoising is indicative of modest nonlinearity in all three datasets. Once each dataset had been denoised using JAE, the corresponding dimensionalities were estimated using MLE and TNN. These results motivated our choice of d=6 for the intrinsic dimensionality of most of our simulated datasets.

CHAPTER 3: LOW-DIMENSIONAL NEURAL MANIFOLDS FOR THE CONTROL OF CONSTRAINED AND UNCONSTRAINED MOVEMENTS

Ege Altan[1,2], Xuan Ma[1], Lee E. Miller[1,2,3,4], Eric J. Perreault[2,3,4], Sara A. Solla[1,5,*]

[1] Department of Neuroscience, Northwestern University, Chicago, IL, United States of America

[2] Department of Biomedical Engineering, Northwestern University, Evanston, IL, United States of America

[3] Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL, United States of America

[4] Shirley Ryan AbilityLab, Chicago, IL, United States of America

[5] Department of Physics and Astronomy, Northwestern University, Evanston, IL, United States of America

Foreword

This work has been submitted to *Nature Biomedical Engineering* on May 25, 2023.

Abstract

Across many brain areas, neural population activity appears to be constrained to a low-dimensional manifold within a neural state space of considerably higher dimension. Recent studies of the primary motor cortex (M1) suggest that the activity within the low-dimensional manifold, rather than the activity of individual neurons, underlies the computations required for planning and executing movements. To date, these studies have been limited to data obtained in constrained laboratory settings where monkeys executed repeated, stereotyped tasks. An open question is whether the observed low dimensionality of the neural manifolds is due to these constraints; the dimensionality of M1 activity during the execution of more natural and unconstrained movements, like walking and picking food, remains unknown. We have now found similarly low-dimensional manifolds associated with various unconstrained natural behaviors, with dimensionality only slightly higher than those associated with constrained laboratory behaviors. To quantify the extent to which these low-dimensional manifolds carry

task-relevant information, we built task-specific linear decoders that predicted EMG activity from M1 manifold activity. In both settings, decoding performance based on activity within the estimated low-dimensional manifold was the same as decoding performance based on the activity of all recorded neurons. These results establish functional links between task-specific manifolds and motor behaviors, and highlight that both constrained and unconstrained behaviors are associated with low-dimensional M1 manifolds.

Introduction

The number of neurons involved in planning and executing a motor behavior, no matter how simple or complex, far exceeds the number of variables relevant to behavior. This mismatch makes the signals across neurons highly redundant and raises the question of how the activity of a population of neurons represents the relatively few variables relevant to behavior (P. Gao and Ganguli 2015; Gallego et al. 2017; Saxena and Cunningham 2019). Contemporary studies across many brain areas report that the activity of populations of neurons is constrained to low-dimensional subregions of the neural space known as neural manifolds (Stopfer, Jayaraman, and Laurent 2003; DiCarlo and Cox 2007; Churchland et al. 2012; Mante et al. 2013; Kaufman et al. 2014; Sadtler et al. 2014; Gallego et al. 2017; Chung, Lee, and Sompolinsky 2018; Remington et al. 2018; Low et al. 2018; Rubin et al. 2019; Chaudhuri et al. 2019; Russo et al. 2020; Nieh et al. 2021; Libby and Buschman 2021; Chandak and Raman 2021; Ehrlich and Murray 2022; Gardner et al. 2022): the computations required for planning and executing behaviors appear to be carried out through the patterns of neural activity within the manifolds rather than through the independent activity of individual neurons (Churchland and Shenoy 2007; Kaufman et al. 2014; Chaudhuri et al. 2019; Rubin et al. 2019; Vyas et al. 2020; Libby and Buschman 2021).

Despite widespread use of the term “low dimensionality”, its precise meaning and significance remains unclear. Specifically, it raises questions such as: What qualifies as low dimensional? How does the behavioral complexity affect the dimensionality of the neural manifold? Should the nonlinearity in the neural manifold be considered when estimating its dimensionality? Recent definitions from Jazayeri and Ostojic’s comprehensive review aim to clarify this ambiguity with a focus on the neural computational principles associated with both *intrinsic* and *embedding dimensionality* in neural recordings (Jazayeri and Ostojic 2021). A simple example illustrates the distinction between intrinsic and embedding dimensionalities (Jazayeri and Ostojic 2021). A ring has an intrinsic dimension of 1. If it is flat, it lies in an (x_1, x_2) plane, and its embedding dimension is 2. If the ring is made of a flexible material, it can be locally pinched away from the (x_1, x_2) plane. A full description of the ring now requires an additional direction x_3; its embedding dimension is 3. If the ring existed within a 100-dimensional space, there are still 97 directions along which the ring could be locally pinched away from the (x_1, x_2) plane. Each such pinch involving a new direction increases the embedding dimension by one without affecting the intrinsic dimension. This simple example illustrates the intuition that the discrepancy between embedding and intrinsic dimensions is a proxy for the degree of nonlinearity of the manifold (Altan et al. 2021).

The geometrical distinction between intrinsic and embedding dimensions leads to recent hypotheses about their distinct functionalities: the intrinsic dimensionality quantifies the number of independent latent variables needed to describe the collective activity of a population of neurons, while the embedding dimensionality is related to how collective information is processed and relayed downstream (Jazayeri and Ostojic 2021). For instance, recent theoretical studies posit that the low embedding dimensionality of the relevant manifolds allows for accurate brain-to-behavior maps (P. Gao and Ganguli 2015; P. Gao et al. 2017). A low embedding dimensionality implies that the relevant population dynamics can be captured by

sampling the activity of a relatively small number of neurons. This observation has two implications. At a conceptual level involving our understanding of information processing in the brain, a low embedding dimension may facilitate the downstream readout of neural activity across brain areas; at a practical level involving our efforts to monitor and interpret neural population activity, a low embedding dimension allows the relevant population activity to be captured by the current recording technologies (Urai et al. 2022).

The comparative analysis of intrinsic and embedding dimensionality is thus a critical concept to 1) quantify the redundancy in neural representations, 2) identify the number of latent components of collective activity in the neural population, 3) investigate the relation between these latent variables and behavioral task variables, sensory inputs, context, planning, and future expectations, and 4) quantify the extent to which behaviors can be decoded using current neural recording technologies.

In this study, we focus on primary motor cortical (M1) neural manifolds, which have been previously found to be low dimensional (Sadtler et al. 2014; Gallego et al. 2017; Perich, Gallego, and Miller 2018; Gallego et al. 2018; Altan et al. 2021). However, these earlier studies were conducted in laboratory settings where highly trained monkeys performed simple, constrained reach and grasp tasks. In contrast, natural behaviors outside the laboratory do not share these constraints. It remains unclear whether the observed low dimensionality of M1 manifolds is simply a byproduct of the constraints associated with laboratory tasks or if it reflects an intrinsic property of population dynamics in M1 (P. Gao and Ganguli 2015). If the M1 manifolds corresponding to unconstrained behaviors were found to be similarly low dimensional, it would suggest that low dimensionality is a general organizational principle about neural population activity in M1. In this view, the relatively few latent variables that characterize a manifold would be sufficient for characterizing the cortical control of movement. These population signals are to be read out by the downstream neural circuitry that ultimately causes movement.

In intracortical Brain-Computer Interface (iBCI) applications, we operate under hardware limitations that vastly undersample the population of motor cortical neurons contributing to a specific behavior. When M1 population activity is constrained to a low-dimensional manifold, as is the case for constrained laboratory tasks, this limitation is overcome when decoding M1 neural activity within iBCIs. The question that we address here is whether the confinement of M1 population dynamics to low-dimensional manifolds also applies to unconstrained, naturalistic behaviors for which the development of iBCIs is highly desirable.

Our primary objective in this work was to characterize the intrinsic and embedding dimensionality of M1 manifolds corresponding to unconstrained behaviors such as grasping small treats while standing in the cage and quadrupedal locomotion over perch bars, and to compare them to the intrinsic and embedding dimensionalities of manifolds corresponding to constrained behaviors in the laboratory. In this analysis, we first applied denoising algorithms to the neural signals; this mitigates the overestimation effects of noise on the dimensionality estimates (Campadelli et al. 2015; Camastra and Staiano 2016; Altan et al. 2021). Following denoising, we computed both intrinsic and embedding dimensionalities of the neural activity. We found that the intrinsic and embedding dimensionality of neural manifolds were slightly higher in unconstrained settings, but still extremely low.

Our secondary objective was to characterize the geometry of the manifolds associated with unconstrained behaviors. To this end, we assessed the nonlinearity of the neural representations in both constrained and unconstrained settings. While we found evidence that the low-dimensional manifolds associated with unconstrained behaviors were nonlinear, most latent dynamics involved exploring nearly linear regions within the neural manifolds.

Our final objective was to investigate whether the low-dimensional M1 manifolds carry sufficient information about behavior to decode simultaneously recorded electromyograms (EMGs). In both laboratory and cage settings, we demonstrated that EMGs could be decoded from the level

of activation of relatively few latent variables within neural manifolds, and that the performance of these decoders was as good as that obtained when decoding EMGs from the activity of all recorded neurons.

Our study illustrates that the low dimensionality of primary motor cortical manifolds is not exclusive to constrained laboratory tasks but is also present in unconstrained motor behaviors within a cage environment. Although the manifolds associated with unconstrained tasks have slightly higher intrinsic and embedding dimensionalities than those associated with constrained tasks, their existence indicates that the low dimensionality of M1 manifolds is characteristic of the population dynamics of M1 rather than a mere consequence of laboratory constraints. Our study advances our understanding of the computational strategies implemented to achieve M1's role in processing and representing muscle-related information. In addition, we expect these findings to facilitate the extension of neural prosthetics and brain-machine interfaces to unconstrained settings.

Methods

Recordings, tasks, and data preprocessing

We trained two 9-10 kg monkeys (Macaca mulatta) to perform tasks in two different environments. The first is defined as the "in-lab" environment, where the monkeys were seated on a standard primate chair and trained to perform a grasping task **(Figure 3-1a**). The task required them to reach and grasp a force-instrumented device located 30 cm in front of their shoulder using either left hand (monkey G) or right hand (monkey P). The shape of the device during an experimental session determined the type of grasping: a cylinder for power grasps with the palm and the fingers, a small rectangular cuboid for key grasps with the thumb and the edge of the index finger, and a small rectangular cuboid recessed within a thin slot for precision grasps with the tips of the thumb and the index finger. A pair of force sensitive resistors (FSRs)

were attached on the sides of the devices to measure the grasping forces the monkeys applied. A monitor was placed above the device to display such forces with a cursor; the position of the cursor along the vertical and horizontal axes was determined by the sum and the difference of the FSR outputs, respectively. In each trial, the monkeys were initially required to keep the hand resting on a touch pad for a random time (0.5-1.0 s). A successful holding triggered the onset of one of three possible rectangular targets on the screen and an auditory go cue. The monkey was required to place the cursor into the target and to hold it there for 0.6 s by increasing and maintaining the grasping force applied on the device.

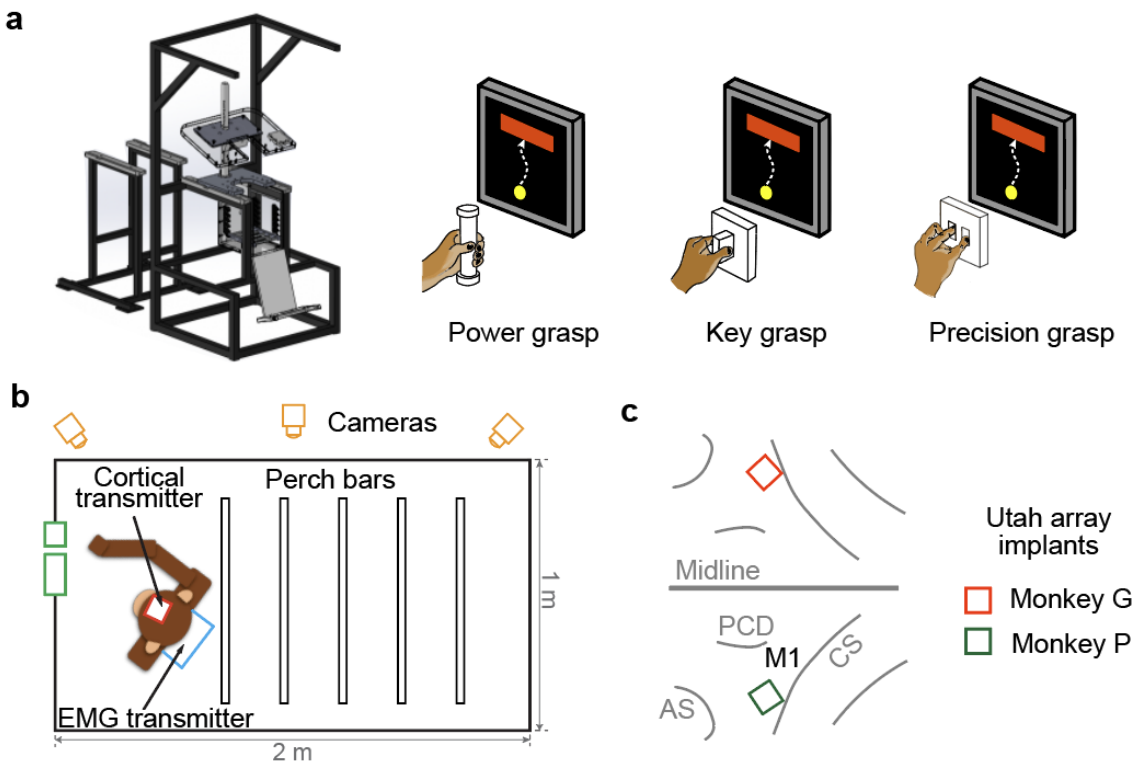

Figure 3-1: Tasks and recordings. a) Primate chair in which the monkeys performed a series of power, key, and precision grasp tasks in the constrained laboratory environment. b) The unconstrained cage environment where the monkeys moved freely. In the cage, monkeys did quadrupedal locomotion by grasping and walking over the perch bars (bar walk task). The monkeys also received treats from the experimenters (treat task). Monkeys were monitored using multiple synchronized cameras from which the individual behavior types were manually

segmented, allowing for the corresponding segmentation of neural and EMG data. c) 96-electrode Utah arrays were implanted in the hand area of the primary motor (M1) cortex of each monkey (monkey G: right hemisphere, monkey P: left hemisphere).

The second environment is defined as the “in-cage” environment, where the monkeys were placed inside a 2×1×1 m plastic cage **(Figure 3-1b**). There were five bars spanning the width of the cage mounted 10 cm above the floor; the monkey grasped these bars while walking back and forth the length of the cage. The monkey typically would do a series of power grasps on these perches as he moved. We called this behavior *bar walk*. There are small holes on the door of the cage, through which experimenters could present small food pellets as treats to the monkeys inside; the monkeys typically used a precision grasp with the thumb and the index finger to take them. We called this behavior *treat grasp*. We identified single in-cage bar walk and treat grasp segments from the continuous recordings based on synchronized video recordings, and defined those segments as “successful trials”.

We implanted a 96-electrode array with 1.5 mm shaft length (Blackrock Microsystems, Salt Lake City, UT) in the hand area of motor cortex (M1) of each monkey **(Figure 3-1c**; monkey G: right, monkey P: left). Neural signals were collected using a Cerebus system (Blackrock Microsystems, Salt Lake City, UT). For the in-lab recordings, the neural signals were amplified by a Cereplex-E headstage. For the in-cage recordings, the neural signals were amplified, digitized at 30 kHz, and transmitted by a Cereplex-W wireless headstage. The neural signals were then digitally band-pass filtered (250 – 5000 Hz). Spikes were detected using a threshold set at -5.5 times the root-mean-square (RMS) amplitude of the signal on each channel, and the time stamp and a 1.6 ms snippet of each signal surrounding the threshold crossing were recorded. We used multiunit threshold crossings on each channel instead of well isolated single neurons in all our data analyses. We applied a Gaussian kernel (S.D.: 100 ms) to the spike counts in 50 ms, non-overlapping bins to obtain a smooth firing rate as function of time for each

channel. We excluded channels with an average firing rate < 0.5 Hz during either in-lab or in-cage movements.

We also implanted intramuscular electromyographic (EMG) leads in 23 arm, forearm, and hand muscles in the left arm of monkey G and in the right arm of monkey P, contralateral to the M1 implant. For the in-lab recording sessions of monkey G, we collected EMG signals using a multi-channel differential amplifier and the analog input channels of the Cerebus system. For the in-cage sessions of monkey G and all sessions of monkey P, we collected EMG signals using a micro multi-channel amplifier (RHD2132, Intan Tech., Los Angeles, CA) and a wireless transmitter (RCB-W24A, DSP Wireless Inc., Haverhill, MA). The micro amplifier and wireless transmitter were both placed in a backpack on the monkey's jacket. The EMG signals were amplified, band-pass filtered (4-pole, 50 - 500 Hz), and sampled at 2000 Hz. The EMGs were subsequently digitally rectified and low-pass filtered (4-pole, 10 Hz, Butterworth), and subsampled to 20 Hz (50 ms bins) to match the neural spike counts.

For each monkey, we divided the neural and EMG data collected for a given task into smaller datasets of non-overlapping data segments of 75 seconds duration. We chose 75 seconds based on simulations (Altan et al. 2021) and on a preliminary analysis (**Supplementary Figure 3-1**) that assessed the effect of the amount of temporal data on dimensionality estimates. All neural and EMG data channels were normalized to their 95th percentile value.

Embedding and intrinsic dimensionality estimation algorithms

We used two algorithms for dimensionality estimation: Parallel Analysis (Horn 1965; Buja and Eyuboglu 1992; Franklin et al. 1995) and Two Nearest-Neighbors (Facco et al. 2017). These methods were selected due to their superior performance in evaluating the dimensionality of simulated linear and nonlinear manifolds, respectively (Altan et al. 2021).

Parallel Analysis (PA) is an embedding dimensionality estimator based on the eigenvalues of the covariance matrix for the dataset (Horn 1965). Unlike other embedding dimensionality estimators that rely on eigenvalue computation, PA does not rely on a predetermined variance threshold, but rather relies on counting the number of eigenvalues that exceed their respective values in a null distribution. This method has been shown to be an accurate estimator of embedding dimensionality in many fields (Dinno 2009).

To estimate the embedding dimensionality of a given M-by-N dataset with M samples of N recorded signals, the PA algorithm proceeds as follows. First, it repeats the following process K times: for each feature, the data are independently shuffled using a random permutation along the corresponding column, the temporal axis; this shuffling breaks the correlation across features. Then, the eigenvalues of the features covariance matrix for the shuffled data are obtained and sorted from largest to smallest. A null distribution is created for each eigenvalue from the K shuffles. We used $K = 200$ repetitions and fixed the random number generator seed for reproducibility.

Next, we computed the 95th percentile from the null distribution of each eigenvalue; this serves as a significance threshold. Finally, we counted the number of original eigenvalues that were larger than their respective 95th percentile null distributions. This number was D_{PA}, the embedding dimensionality estimate obtained by PA. While the 95th percentile cutoff was an arbitrary choice, PA is quite robust to a reasonable choice of threshold to characterize significant deviations to the null, noise-based eigenvalue distribution (Dinno 2009). In contrast, the usual practice of choosing a threshold for the variance-accounted-for (VAF) in a PCA-based approach often results in significant changes in the estimate of the embedding dimensionality when the threshold is changed by a few percentage points, such as choosing a 90th instead of a 95th cutoff.

To estimate the intrinsic dimensionality, we used Two Nearest-Neighbors (TNN), an estimator based on the local adjacency of data points (Levina and Bickel 2004; Facco et al. 2017; Allegra et al. 2020). Specifically, this method is based on the ratio μ of distances to the second vs the first nearest neighbors of a given point. The second-to-first nearest neighbors distance ratio is Pareto distributed with a unitary scale parameter and a shape parameter equal to the intrinsic dimensionality (Denti et al. 2022). The intrinsic dimensionality can be approximated using the following equation:

$$D_{TNN} = -\frac{\log(1-F(\mu))}{\log(\mu)} \quad \text{(Equation 3-1)}$$

Here, F(μ) is the empirical cumulative distribution of the ratio of second-to-first nearest neighbor distances for each data point. In this study, we used the publicly available python package called `scikit-dimension` (Bac et al. 2021) to compute the TNN estimate of intrinsic dimensionality.

Denoising algorithms

Noise is a confounding factor for both intrinsic and embedding dimensionality estimates (Campadelli et al. 2015; Camastra and Staiano 2016; Altan et al. 2021). Random noise across channels will lead to increased dimensionality estimates that might even approach the total number of recorded signals. To address this issue, we implemented two denoising approaches that rely on an initial estimate of an upper bound dimensionality $D = D_{PA}$ obtained using Parallel Analysis (PA).

The first approach, PCA denoising, is a linear method based on Principal Component Analysis (PCA, **Figure 3-2a**). After determining the value of D using PA, we used the D leading principal components to reconstruct the original data. This approach assumes that most of the noise is present in the low-variance principal components that have been discarded.

The second approach, Joint Autoencoder (JAE) denoising, is a neural network-based method **(Figure 3-2b**). We divided the noisy neural signals with N features into two disjoint subsets, each including $N/2$ features, and used the compressive halves of two autoencoders to map each subset into a D-dimensional subspace. The reconstructed versions of the subsets of dimension $N/2$ resulted from the expansive halves of the respective autoencoders. The JAE was trained using a cost function that minimized the mean-squared error associated with the reconstructions of each of the two subsets, as well as the mean-squared error between the two intermediate D-dimensional latent signals. This approach assumes that each subset contains the necessary information to identify the underlying D-dimensional signals while the noise components are common across the two subsets.

To evaluate the performance of the denoising algorithms, we calculated the coefficient of determination (R^2) between the denoised and noisy neural signals **(Figure 3-2c and 3-2d**). We validated these denoising approaches in a previous study based on simulated neural datasets with known levels of noise (Altan et al. 2021).

Robustness of the dimensionality estimates relative to the ambient dimensionality of the underlying neural space

Accurately interpreting the dimensionality of neural data requires consideration of the amount of data used for dimensionality estimation. Prior studies have shown that the amount of data used, measured by both the number of samples M and number of features N, can affect the accuracy and precision of the dimensionality estimates (Cunningham and Yu 2014; P. Gao and Ganguli 2015; Camastra and Staiano 2016; Altan et al. 2021; Jazayeri and Ostojic 2021). To ensure stable dimensionality estimates, we determined that datasets roughly one minute in length were sufficient for our analyses (see **Supplementary Figure 3-1**).

Similarly, we conducted analyses to determine the robustness of our estimates with respect to the number of neurons. Specifically, we aimed to assess whether our dimensionality estimates

reached an asymptote as the number of neurons increased, indicating reliable estimates. We gradually sampled increasing numbers of neurons and repeatedly computed the dimensionality estimates 20 times for each number. We sampled 5, 10, 15, 20, 25, 30, 35, and the maximum number of recorded neurons for each dataset. If the estimates reached an asymptote, we concluded that additional neurons would not contribute meaningfully to the estimated dimensionality, indicating that our dimensionality estimates were robust. Conversely, if the estimates continued to increase with the number of neurons, this suggested that more neurons would be necessary to obtain reliable estimates.

Quantifying the extent of nonlinearity

We used two measures to compute the extent of nonlinearity in the neural recordings; these measures arise from two complementary perspectives. The first measure, that we called *manifold nonlinearity index*, focused on the nonlinearity of the geometry of the neural manifold. The second measure, that we called *local flatness index*, focused on how much of the population activity is concentrated in approximately linear regions within the manifold.

To compute the manifold nonlinearity index, we used the ratio of embedding-to-intrinsic dimensionality of the neural manifold (Altan et al. 2021). The manifold nonlinearity index is similar to the dimensionality gain metric used in an artificial neural network study that sought to extract predictive latent signals (Recanatesi et al. 2021).

To compute the local flatness index, we compared the Euclidean and geodesic distances between every pair of data points in the state space of each dataset (**Supplementary Figure 3-2**). Euclidean distance is the length of the shortest straight line between two points, while geodesic distance takes into account the curvature of the neural manifold by measuring the length of the shortest path that lies within the manifold and connects the two points. To calculate the geodesic distance from one point to all others, we started from a nearest-neighbor cloud

limited to 200 samples. Distances to those 200 closest activity patterns were computed along the straight lines joining the central point to its neighbors within the cloud. It is from these short linear segments that global geodesics are constructed.

A large discrepancy between the Euclidean and geodesic distances indicates that the population activity explores curved, nonlinear regions within the manifold. A higher degree of overlap suggests that the neural activity is largely confined to nearly linear regions within the manifold. To determine the extent to which the neural activity evolves in a linearizable region of the neural manifold, we computed the empirical distribution of both Euclidean and geodesic distances for each dataset. We used 20 bins to cover the full range shared between these two distance distributions; we then converted these distributions to normalized densities. Within each bin, the smaller of the two density values, multiplied by the bin size and summed over all bins provides the value for the local flatness index. **Supplementary Figure 3-2** shows a graphical illustration of how we computed the local flatness index for a randomly chosen bar walk dataset from Monkey P.

Computing the activity on the low-dimensional neural manifolds

To project the population neural activity onto the low-dimensional neural manifolds, we employed two different techniques, each with different underlying assumptions about the linearity of the data: Principal Component Analysis (PCA) and a nonlinear autoencoder. PCA is a widely used linear technique based on finding the orthogonal directions in neural space that correspond to maximum variance in the data. In contrast, the nonlinear autoencoder is a neural network-based technique that can capture complex nonlinear relationships in the data and is especially useful for projecting population activity onto the low-dimensional nonlinear neural manifolds. The autoencoder consisted of five feedforward layers (input, first hidden, bottleneck, second hidden, and output layers) with ReLU activation and was trained using the ADAM

optimizer on the reconstruction error, the mean-squared error between output and input; default training parameters were used, and training lasted for 50 epochs.

For both techniques, we used the TNN estimates of the intrinsic dimensionality to determine the manifold dimensions. For PCA, this meant that we retained a number of leading principal components equal to D_{TNN} rounded to the nearest integer. For the autoencoder, we made the number of neurons in the bottleneck layer equal to D_{TNN} rounded to the nearest integer.

Decoding electromyograms (EMGs)

Our setup allowed for the simultaneous recording of neural signals from the primary cortex and EMG signals from muscles as the monkeys engaged in various motor behaviors. We used linear regression to decode EMG signals from neural signals; this basic, interpretable, and linear method is widely used in the field of brain-computer interfaces (Carmena et al. 2003b; Glaser et al. 2020). Each set of EMG signals was decoded using three types of inputs: all available neurons, latent signals obtained from PCA, and latent signals obtained from the bottleneck layer of the autoencoder. We used five-fold cross validation for each approach and reported all five test folds for a given neural-to-EMG dataset pair.

Statistical analyses

We reported our computations in the mean ± standard deviation format. We used Welch's t-test for statistical comparisons when the distribution of the statistic of interest was normal. We used Wilcoxon rank sum test in the case when the normality assumption was violated (**Supplementary Figure 3-3**). When we compared EMG decoding accuracies from PCA embeddings, AE embeddings, and all available neurons, we used repeated measures ANOVA **(Figure 3-6**). We reported the p-values obtained from the statistical test in the figure legends and used Bonferroni correction for multiple comparisons where applicable.

Ethics statement

All surgical and experimental procedures that yielded the datasets comprising of multi-electrode neural recordings and intramuscular electromyogram (EMG) signals from non-human primates were approved by the Institutional Animal Care and Use Committee of Northwestern University. The two monkeys were monitored daily to ensure their well-being and health. Their diet consisted of a standard laboratory animal diet supplemented with fresh fruits and vegetables to provide optimal nutrition. Additionally, the monkeys were provided with access to various types of enrichment to promote mental stimulation and overall well-being.

Results

Reducing the effects of noise using linear and nonlinear methods

Both neurons and neural recordings are noisy (Faisal, Selen, and Wolpert 2008), which can cause an overestimation of the dimensionality (Camastra 2003; Facco et al. 2017; Altan et al. 2021; Jazayeri and Ostojic 2021) of recorded neural activity. To mitigate the effect, we first denoised the recorded neural signals using two approaches that differed by their assumption about the linearity of the manifold to which the data is mostly confined. The linear denoising method was based on Principal Component Analysis (PCA, **Figure 3-2a**). The nonlinear denoising method was a neural network-based approach called the Joint Autoencoder (JAE, **Figure 3-2b**).

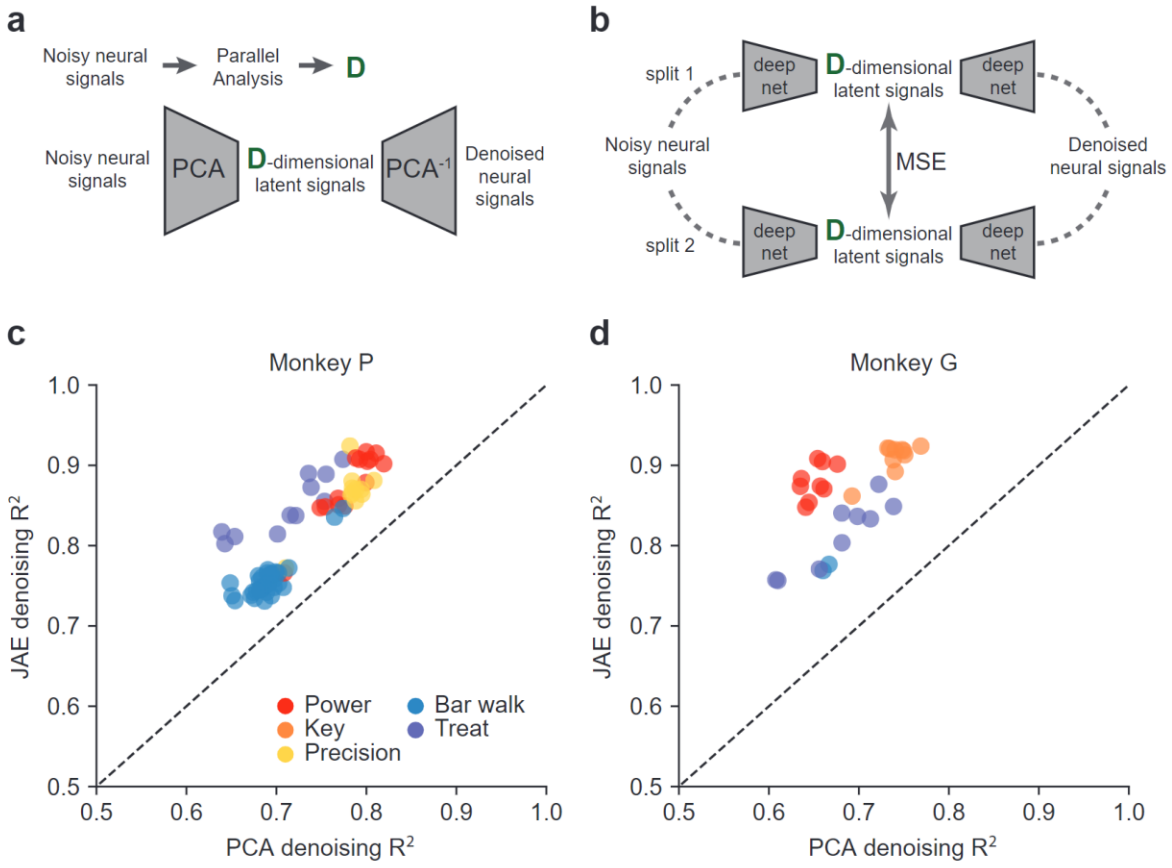

Figure 3-2: Denoising the neural signals. Two methods were employed to denoise neural signals: a PCA based approach and a deep neural network-based approach called JAE. a) In the PCA based denoising approach, noisy neural signals were projected to the leading D principal components, with D determined by Parallel Analysis. The denoised neural signals were the reconstructions from the D-dimensional bottleneck. b) For the deep neural network-based denoising approach we used the joint autoencoder (JAE). Noisy neural signals were randomly divided into two subsets. Each subset was nonlinearly projected to a D-dimensional latent space using two parallel deep networks. D was again determined using Parallel Analysis. The mean squared error (MSE) loss used to train the network forced the D-dimensional latent representations across the parallel networks to match. Denoised neural signals were obtained from the reconstructions based on the activity in the D-dimensional bottlenecks. c) Comparison of reconstruction accuracies between PCA and JAE denoising approaches for Monkey P. Each circle

represents the R^2 performance of each denoising method for a single dataset. Warm colors indicate the constrained laboratory tasks. Cold colors indicate the unconstrained cage tasks. d) Same as in panel c but for Monkey G.

For every data set for both monkeys, denoising the data using JAE yielded a higher reconstruction accuracy (R^2) than denoising the data using PCA **(Figure 3-2c and 2d**). For monkey P and for the power and precision grasp tasks, PCA yielded a reconstruction accuracy of 0.78 ± 0.03 (mean ± standard deviation) and 0.78 ± 0.02, whereas JAE yielded 0.88 ± 0.04 and 0.87 ± 0.04, respectively. The reconstruction accuracies for the bar walk and treat grasp tasks were 0.69 ± 0.02 and 0.71 ± 0.05 with PCA, and 0.76 ± 0.02 and 0.85 ± 0.04 with JAE. We observed a similar trend for Monkey G. For the power and key grasp tasks, reconstruction accuracy using PCA was 0.65 ± 0.01 and 0.74 ± 0.02, whereas the reconstruction accuracy using JAE was higher at 0.88 ± 0.02 and 0.91 ± 0.02. For the unconstrained tasks, reconstruction accuracies using PCA and JAE were 0.66 ± 0.01 and 0.68 ± 0.05, and 0.77 ± 0.01 and 0.81 ± 0.04, respectively. For the reminder of the analyses, we denoised the neural signals using the JAE approach due to its consistently superior performance over PCA.

Dimensionality estimates of denoised neural signals

Our next goal was to compute the dimensionality of the denoised signals. We applied two dimensionality estimation algorithms: PA and TNN. We used PA to estimate the embedding dimensionality and TNN to estimate the intrinsic dimensionality of the neural signals. These methods were chosen based on their assumption about linearity and superior performance to alternatives (Altan et al. 2021).

Both embedding and intrinsic dimensionality estimates were slightly higher for unconstrained behaviors compared to constrained behaviors for both Monkey P and Monkey G **(Figure 3-3**). Dimensionality estimates from PA were almost always higher than those from TNN. For Monkey

P, the PA dimensionality across the constrained laboratory tasks increased from 7.28 ± 1.67 to 10.40 ± 2.54 across the unconstrained tasks **(Figure 3-3a**). Similarly for Monkey G, the PA dimensionality across constrained and unconstrained tasks was 5.84 ± 0.78 and 9.27 ± 1.62, respectively **(Figure 3-3b**). We observed the same trend in TNN dimensionality across the two task settings. TNN dimensionality estimates for Monkey P increased from 5.01 ± 0.39 to 6.5 ± 0.58 from the constrained to unconstrained tasks **(Figure 3-3c**). Similarly for Monkey G, TNN estimates went from 4.72 ± 0.18 to 5.70 ± 0.27 **(Figure 3-3d**). In summary, neural manifolds for all tasks were consistently much smaller than the total number of sampled neurons, which defines the dimensionality of the empirical neural space that contains these task-specific manifolds. We note that both the intrinsic and embedding dimensionalities were slightly higher for unconstrained tasks.

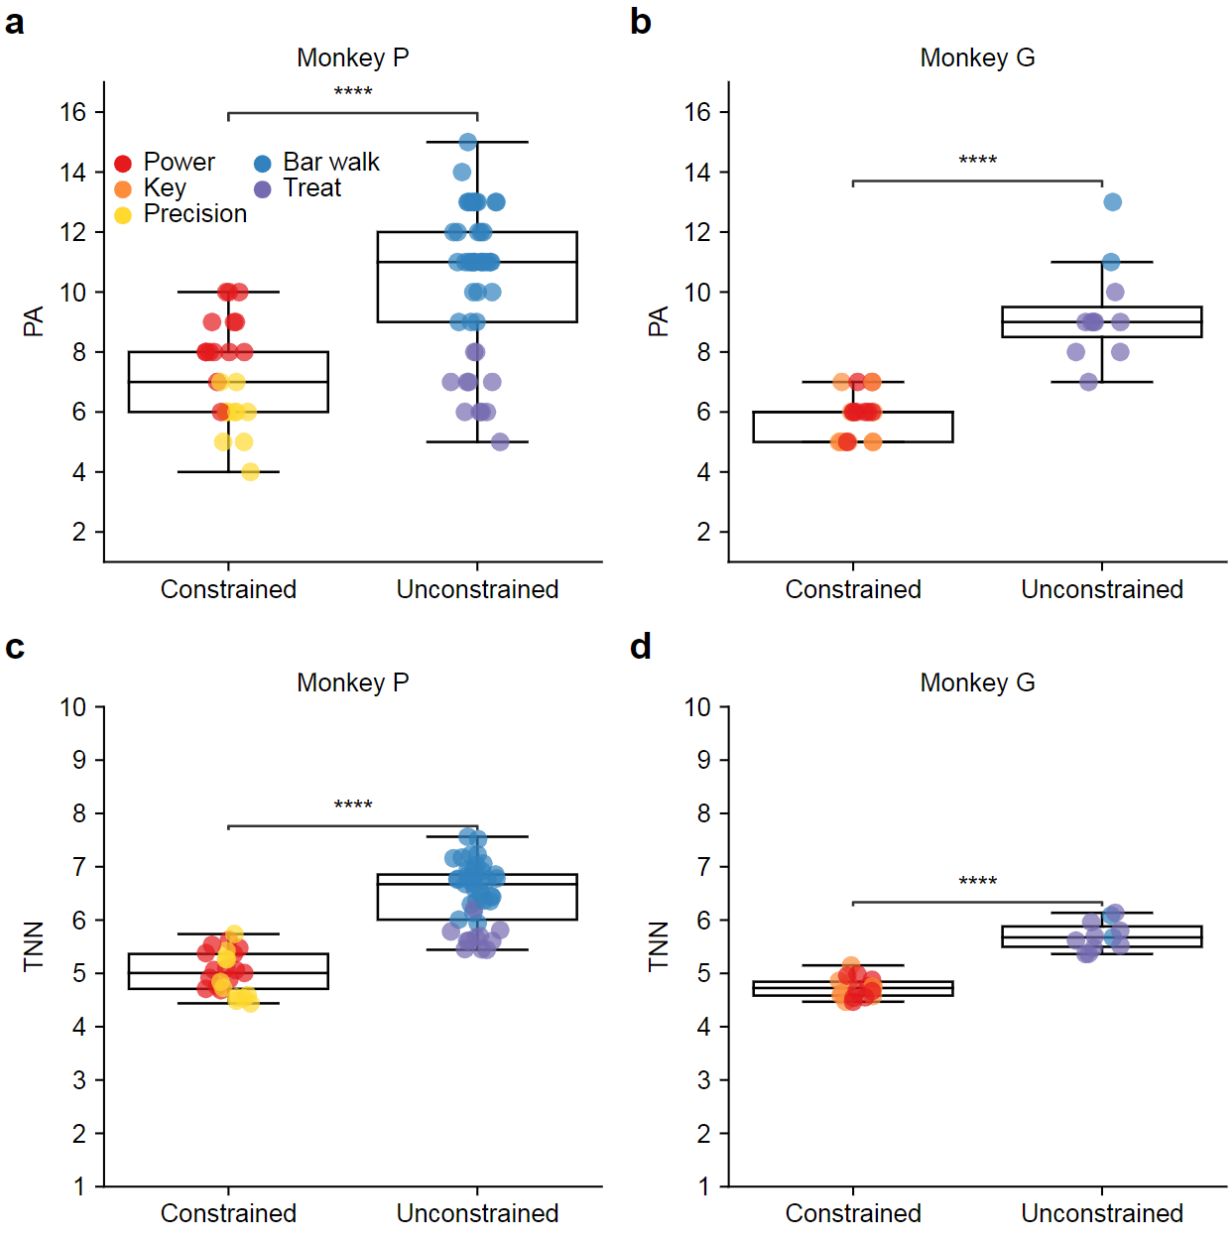

Figure 3-3: Estimating the embedding and intrinsic dimensionalities of neural manifolds.

We applied Parallel Analysis (PA) and Two-Nearest Neighbors (TNN) to estimate the embedding and intrinsic dimensionality, respectively. a) Comparison of the PA estimates for constrained and unconstrained behaviors for Monkey P ($p \approx 0$). Each circle represents the dimensionality estimate from a single dataset. Warm colors indicate constrained behaviors in the laboratory. Cold colors indicate unconstrained behaviors in the cage. b) Same as in panel a but for Monkey G (p

≈ 0). c) Comparison of the TNN estimates across constrained and unconstrained behaviors for Monkey P ($p \approx 0$). d) Same as in panel c but for Monkey G ($p \approx 0$).

We asked if the small but significant increase in the dimensionality of neural manifolds associated with unconstrained tasks is a reflection of increased task complexity (P. Gao and Ganguli 2015; P. Gao et al. 2017), and used the intrinsic dimensionality of EMG signals as a proxy for task complexity (**Supplementary Figure 3-3**). The increase in EMG dimensionality when comparing unconstrained to constrained tasks was also small but significant, and particularly noticeable in some instances of unconstrained tasks for monkey G (**Supplementary Figure 3-4**).

Investigating the extent of nonlinearity in the neural manifolds

To quantify the degree of nonlinearity in the geometry of neural manifolds, we computed the *manifold nonlinearity index* (see Methods). On average, the manifold nonlinearity index was slightly higher for the unconstrained tasks than for the constrained tasks, but rarely exceeded two **(Figure 3-4a and 4b**). On average, the manifold nonlinearity index was 1.45 ± 0.29 (Monkey P) and 1.24 ± 0.16 (Monkey G) for the constrained tasks, and 1.58 ± 0.28 (Monkey P) and 1.63 ± 0.26 (Monkey G) for unconstrained tasks, respectively. The larger manifold nonlinearity index for Monkey P was not statistically significant.

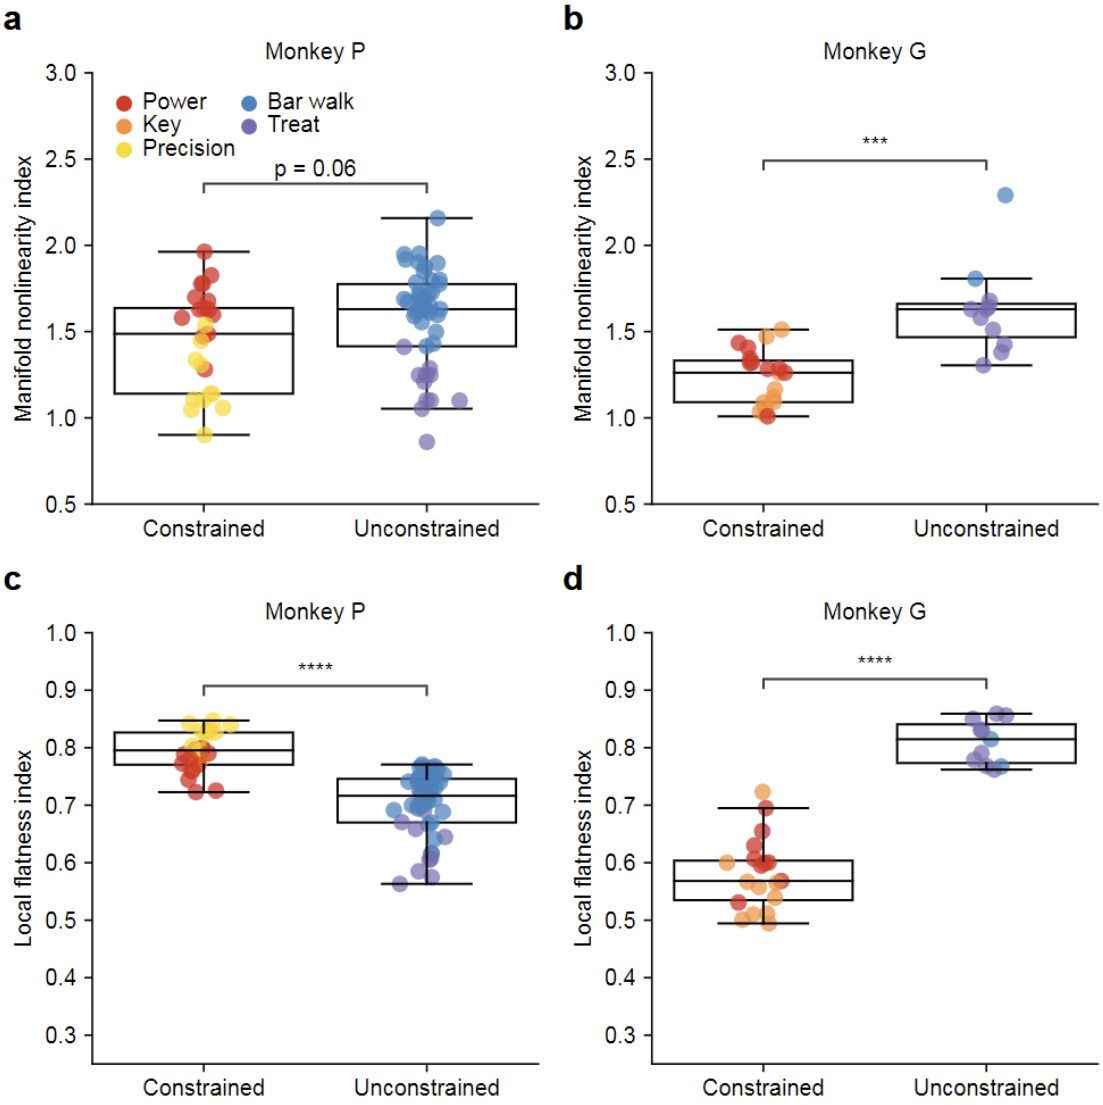

Figure 3-4: Manifold nonlinearity index and local flatness index of neural manifolds. a) Manifold nonlinearity index, the ratio of the embedding (PA) to intrinsic (TNN) dimensionality, for all datasets for Monkey P. Each circle represents the manifold nonlinearity index for a single dataset. Warm colors indicate constrained behaviors in the laboratory. Cold colors indicate unconstrained behaviors in the cage ($p = 0.06$). b) Same as in panel a but for Monkey G ($p = 0.0005$). c) Local flatness index, measured by the fractional overlap between the distributions of

Euclidean and geodesic distances (see Methods, Supplementary Figure 2), for Monkey P ($p \approx 0$). d) Same as in panel c but for Monkey G ($p \approx 0$).

The manifold nonlinearity index only partially elucidates the nature of nonlinearity of the neural manifold, because it does not reveal how neural activity samples the manifold. To complement it, we computed the local flatness index, which quantifies the extent to which the population activity samples linear regions within the manifold (see Methods). The approach is based on comparing all pairwise Euclidean and geodesic distances for each dataset. **(Figure 3-4c and 4d**). The overlap between the distributions of geodesic and Euclidean distances highlights the degree of linearity in the explored regions of the manifold: an overlap greater than 0.50 indicates that most of the data lie in regions of the manifold that are well approximated by the tangent linear subspace (**Supplementary Figure 3-2**).

For Monkey P, the constrained tasks had a local flatness index of 0.79 ± 0.04 and unconstrained tasks had a local flatness index of 0.70 ± 0.06. For Monkey G, local flatness was 0.58 ± 0.06 for constrained and 0.81 ± 0.04 for the unconstrained tasks **(Figure 3-4c and 4d**). The differences for both monkeys were statistically significant, despite being in opposite directions. We notice the low value of the local flatness index for Monkey G when executing constrained tasks, and hypothesize that this effect is likely associated with the adoption of distinct task execution strategies by Monkey G in the laboratory, as illustrated in **Supplementary Figure 3-3b**. For this hypothesis to be confirmed, we would need to analyze and compare the fine details of task execution, an analysis beyond the scope of this work. To summarize these findings the aggregate results in **Figure 3-4c and 4d** show that the local flatness index never fell below 0.5 in any constrained or unconstrained scenario; this indicates a substantial overlap between the distributions of geodesic and Euclidean distances for both monkeys. Results for the local flatness index establish that for the tasks investigated here the neural population dynamics visited states mostly confined to nearly linear regions within slightly nonlinear neural manifolds.

Number of neurons required for stable dimensionality estimates

One critical consideration for interpreting dimensionality is the amount of data needed for dimensionality estimates; previous studies show that the amount of data can affect the accuracy of the estimates (Cunningham and Yu 2014; Altan et al. 2021). We determined that datasets roughly one minute-long were sufficient for stable dimensionality estimates **(Supplementary Figure 3-1).**

Another factor contributing to the reliability of dimensionality estimates is the number of recorded neurons. How robust are the dimensionality estimates with respect to the number of neurons used in analysis? To answer this question, we assessed whether the dimensionality estimates reached an asymptote as the number of neurons used for analysis increased. An asymptotic saturation of the estimated dimensionality would signal reliable dimensionality estimates.

We found that for both monkeys the nonlinear TNN method saturated as the number of neurons increased **(Figure 3-5a and c**). The TNN dimensionality estimates saturated at roughly 20 to 35 sampled neurons for all tasks and monkeys. In contrast, the linear method PA yielded dimensionality estimates that continued to increase as the number of neurons used to estimate manifold dimensionality increased **(Figure 3-5b and d**).

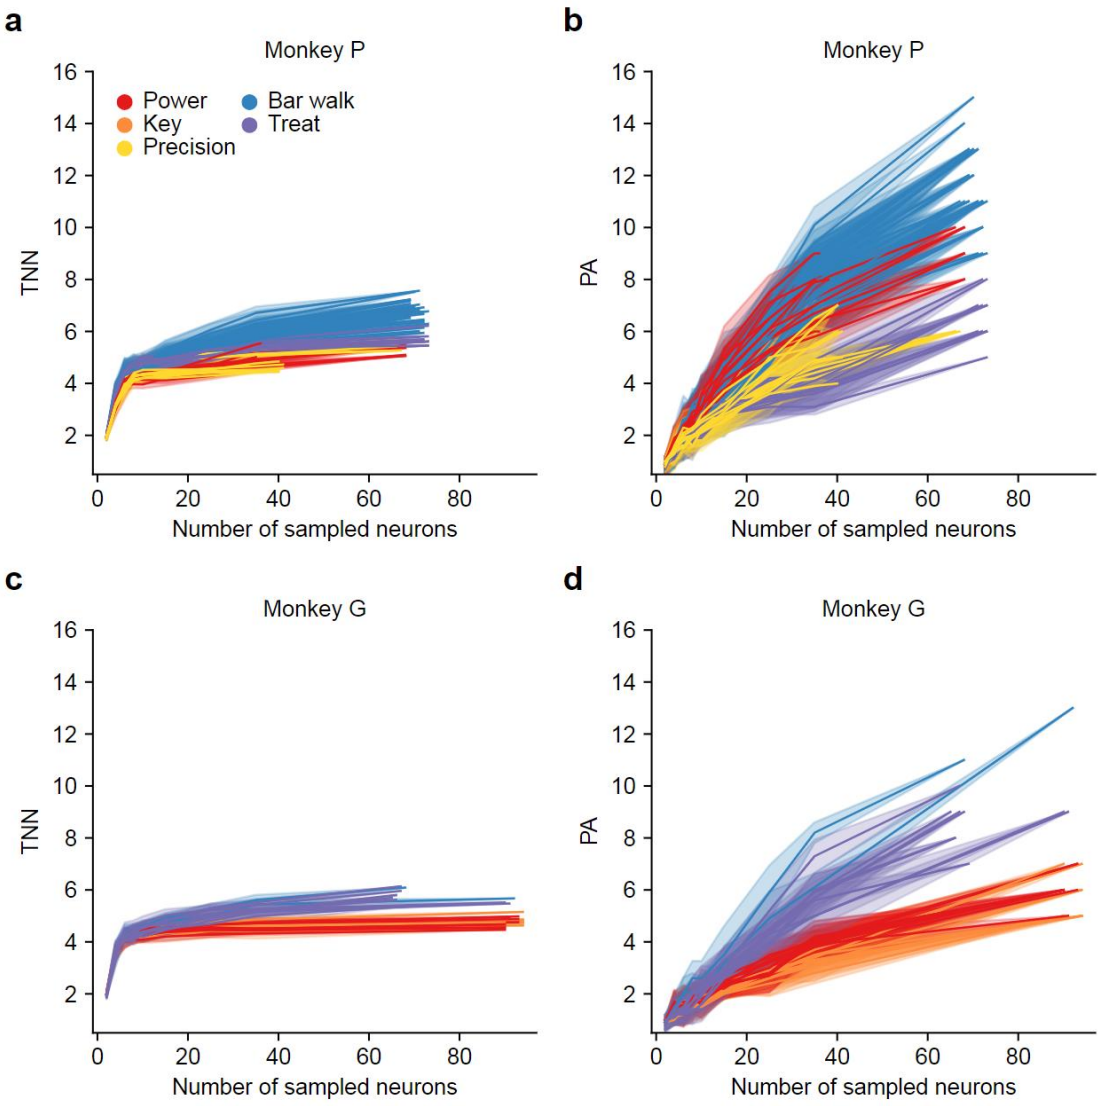

Figure 3-5: Number of neurons required for stable dimensionality estimates. We progressively increased the number of neurons included in the dimensionality estimation. We randomly subsampled, with ten repetitions, a given number of neurons from the total number of neurons available in each dataset. Solid lines indicate the average dimensionality estimate. Shaded regions indicate the standard deviation. a) TNN estimates from for Monkey P. b) PA estimates for Monkey P. c) Same as in panel a, but for Monkey G. d) Same as in panel b, but for Monkey G.

Decoding EMGs from low-dimensional latent activity

Finally, we compared how well the neural manifolds represented behavior by decoding EMGs from the activity within the low-dimensional neural manifolds obtained with both PCA and a nonlinear autoencoder **(Figure 3-6**). We compared the accuracy of these decoders to that of decoders based on the activity of all recorded neurons. In both scenarios, we reported the D_{TNN} estimates of the intrinsic dimensionality of the EMG signals and used this dimensionality as the number of leading latent variables to be used as inputs to a neural-to-EMG decoder. This approach allowed us to directly quantify the extent to which muscle-related information lives in a linear hyperplane that approximates a slightly nonlinear neural manifold of D_{TNN} dimensions. Examples of actual and decoded EMG signals based on all recorded neurons, linear latent variables, and nonlinear latent variables are shown in **Supplementary Figure 3-5** for five different muscles. We also showed the EMG decoding performance from progressively increasing neural manifold dimensionality in **Supplementary Figure 3-6**.

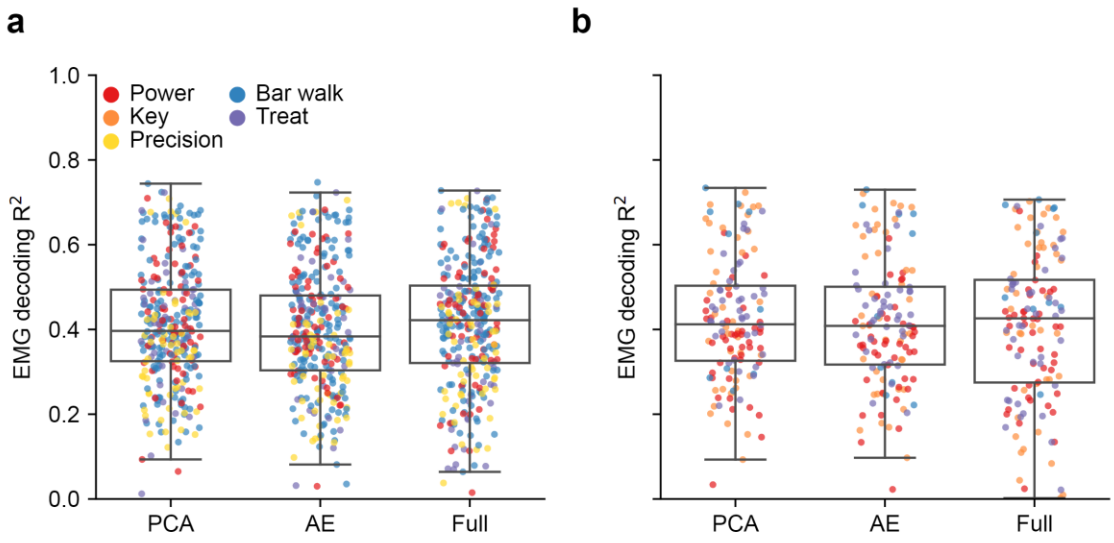

Figure 3-6: Decoding EMGs from low dimensional embeddings and all available neurons. We used linear regression to decode EMGs from denoised neural signals (see Methods for denoising). We either used PCA embeddings (PCA), AE embeddings (AE), or the full-dimensional neural signals (Full) as inputs to linear neural-to-EMG decoders. PCA and AE provided linear and

nonlinear latent variables, respectively. The dimensionality of the latent representations followed from the TNN estimate for each dataset. Each dot represents the decoding accuracy from one of the five test folds of decoding. Different colors represent different tasks. Repeated measures ANOVA was not significant between the three conditions for two monkeys (p = 0.59). a) EMG decoding accuracies for Monkey P. b) Same as in panel a, but for Monkey G.

In all datasets, the accuracies obtained when decoding EMGs from all recorded neural signals and from the low-dimensional latent variables were not significantly different. Importantly, linear decoding from the activity of PCA and autoencoder latent variables restricted to D_{TNN} dimensions yielded similar accuracy–ANOVA with repeated measures was not significant, p = 0.59. For Monkey P, the average test-fold EMG decoding accuracy using PCA, autoencoder, or all recorded neural signals was (mean ± standard deviation) 0.41 ± 0.14, 0.40 ± 0.15, and 0.41 ± 0.15 **(Figure 3-6a**). The corresponding results for Monkey G were 0.42 ± 0.15, 0.42 ± 0.15 and 0.41 ± 0.17**(Figure 3-6b**). These results indicate that there was no difference in EMG decoding from the activity of all recorded neurons or from the latent variables that characterize the low-dimensional neural manifolds. Additionally, the similar EMG decoding accuracy from D_{TNN} latent variables obtained from PCA and nonlinear autoencoder embeddings indicates that the information encoded in the motor cortex relevant to activating muscles exists in a mostly linear subspace of dimension D_{TNN} (**Supplementary Figure 3-5 and 6**).

Discussion

Neural dimensionality has emerged as an important concept for understanding the underlying dynamics and computation abilities of populations of neurons. The dimensionality of the primary motor cortex (M1) manifolds associated with specific tasks has been found to be low in numerous studies limited to constrained laboratory behaviors (Sadtler et al. 2014; Gallego et al. 2017; Russo et al. 2018; Perich, Gallego, and Miller 2018; Altan et al. 2021). In this work, we

computed the intrinsic and embedding dimensionalities of M1 recordings as monkeys engaged in unconstrained tasks in their home cage. Our primary finding was that both the intrinsic and embedding dimensionality of M1 signals were slightly higher than for tasks performed with a single arm, while seated in a primate chair. Although the manifolds associated with unconstrained tasks were found to have slightly higher dimensionalities, these dimensionalities were still very low. In addition, although we found signatures of nonlinearity in the M1 neural manifolds, most of the activity was confined to nearly linear evolved regions within the neural manifolds. Finally, the accuracy of linearly decoded EMGs from the low-dimensional latent variables that characterize the manifolds matched closely the accuracy of EMGs decoded from all recorded neurons, both in cage and in lab.

Linking neural dimensionality to neural computation and processing

There are different definitions and interpretations of dimensionality. We adopted those recently put forward by Jazayeri and Ostojic, who distinguish between intrinsic and embedding dimensionalities (Jazayeri and Ostojic 2021). While the intrinsic dimensionality quantifies the number of independent latent variables encoded by a neural population, the embedding dimensionality plays a role in understanding how the latent variables are processed and transmitted. Together, the analysis of both intrinsic and embedding dimensionalities illuminates how neural systems encode and process information.

Intrinsic dimensionality refers to the actual dimensionality of a nonlinear manifold to which the population neural activity is mostly confined; it may include different classes of information, such as sensory inputs, motor outputs, and other latent variables that correspond to learned experiences and expectations (Jazayeri and Ostojic 2021). Since different brain regions selectively represent these classes, the intrinsic dimensionality of a representation in a particular brain area may be strongly associated with one class and weakly with another. For example, the intrinsic dimensionality of latent signals in early sensory areas like the olfactory

cortex and primary visual cortex is associated with the representation of incoming stimuli such as chemical odors (Pashkovski et al. 2020) or visual gratings (Stringer et al. 2021), respectively. On the other hand, latent signals from the primary motor cortex are more closely associated with motor output (Churchland, Cunningham, et al. 2010; Churchland et al. 2012; Kato et al. 2015; Gallego et al. 2017) and less with future expectations (Zimnik and Churchland 2021). The intrinsic dimensionality of manifolds in the primary motor cortex would therefore be expected to be closely associated with motor output variables such as kinematics or muscle activation.

Embedding dimensionality refers to the dimensionality of the minimal Euclidean space sufficient to fully contain the nonlinear manifold. Although the recurrent dynamics within a given brain area might lead to nonlinear manifolds best characterized by their intrinsic dimension, it is their embedding dimension that best characterizes the signals to be communicated through linear readouts. This observation has led to the view that the embedding dimension reflects how the latent variables are processed for communication to other neural areas (Jazayeri and Ostojic 2021). However, the computational principles associated with the embedding dimensionality are not yet fully understood. What determines whether the embedding dimensionality is high or low? We have argued that a higher degree of manifold nonlinearity will require a larger embedding dimension. But how does this concept connect to the question of communication across areas? One hypothesis is that the neural code is linearly relayed to other brain areas or to the periphery (Fusi, Miller, and Rigotti 2016; Badre et al. 2021; Maass 2016; Semedo et al. 2019; Kohn et al. 2020). In this view, different linear decoders can selectively relay different task relevant information to different downstream areas without interference, a property called “mixed selectivity” (Rigotti et al. 2013; Kaufman et al. 2022). A brain area that relays distinct sets of information to several downstream areas would require the latent variables to be embedded into a higher-dimensional Euclidean subspace to facilitate mixed selectivity. Mixed selectivity thus necessitates a high embedding dimensionality (Badre et al. 2021). As an example of this

correspondence, prefrontal and primary visual cortical neurons have been reported to have a very high embedding dimensionality (Rigotti et al. 2013; Fusi, Miller, and Rigotti 2016; Stringer, Pachitariu, Steinmetz, Carandini, et al. 2019); these observations fit well with the mixed selectivity hypothesis. In contrast, a brain area that does not require extensive mixed selectivity, such as the primary motor cortex, would generate latent variables that only require a relatively low embedding dimensionality, an observation consistent with many reports on M1 manifolds (Sadtler et al. 2014; Gallego et al. 2017; Perich, Gallego, and Miller 2018; Gallego et al. 2018; Pandarinath, O'Shea, et al. 2018; Pandarinath, Ames, et al. 2018; Altan et al. 2021).

Low dimensionality of M1 transcends task constraints

Thus far, the low dimensionality of M1 manifolds has only been investigated in the context of constrained laboratory settings. It is unclear whether the low dimensionalities observed in M1 truly reflect some intrinsic computational property of M1 or whether they are a byproduct of the constraints in stereotyped, repeated laboratory tasks. To begin to elucidate this question, our approach was to obtain and analyze a rich collection of datasets corresponding to two distinct settings: laboratory and cage. These two settings captured different levels of constraint. In the laboratory setting, we placed the monkeys in a primate chair with restraints such that they were only able to move the hand contralateral to their neural implant in specific, well-instructed trial segments. In the cage setting, the monkeys had more freedom to move around the perch bars and could grab small treats from the experimenters as they pleased. We did not provide any instructions on how or when to perform the tasks in the cage; the monkeys could take as much time as they needed and execute these tasks in a non-stereotyped manner. Our results show that the low dimensionality of the primary motor cortex is largely independent of task setting and constraints. Both the intrinsic and embedding dimensionalities of M1 were low in both constrained laboratory and unconstrained cage settings, with only slightly higher dimensionality estimates for the unconstrained tasks **(Figure 3**). Thus, we hypothesize that the low

dimensionality of M1 may not be a byproduct of constrained movements, but rather reflects its computational strategy. Our results, obtained through the lens of dimensionality, fit in with recent evidence of context independence in M1 (Russo et al. 2020). In contrast, recent evidence from mouse cerebellar parallel fiber recordings showed a roughly four-fold increase in embedding dimensionality from constrained limb-actuated lever tasks to spontaneous standing, running, or whisking tasks (Lanore et al. 2021). The difference between M1 and cerebellar results, though from a different species, highlights that different brain areas differentially process task constraints and context.

M1 representations are nonlinear, but only slightly

Task-specific M1 representations were slightly nonlinear for all tasks and in both settings. For a given bottleneck dimension D, the nonlinear JAE denoising algorithm consistently resulted in a better reconstruction of neural signals than PCA **(Figure 1**). The manifold nonlinearity index was between one and two for all datasets **(Figure 4**). Finally, the local flatness index was, on average, around 0.7 when aggregated across all datasets. Although these results describe neural manifolds that were monkey and task specific, a general trend emerges: despite the slight degree of nonlinearity in the geometry of the neural manifolds, task-specific neural dynamics sampled mostly linear regions within the respective nonlinear manifolds **(Figure 4**). In other words, there was evidence of only mild nonlinearity in M1 manifolds associated with the motor behaviors analyzed here; in addition, the population activity was mostly confined to linear regions within these manifolds.

Interpreting the low dimensionality and mild nonlinearity of M1

How should we interpret the low dimensionality and mild nonlinearity of M1 manifolds? One explanation is related to the computational tradeoff between generalizability and expressivity that has been studied in both artificial and biological neural networks (Musslick et al. 2017;

Musslick and Cohen 2019; Pryluk et al. 2019; Sagiv et al. 2020). Let's first consider an extreme case: if the task-specific information encoded in a brain area must be selectively communicated to several different brain areas, latent signals ought to be confined to distinct, linearly independent subspaces, and would therefore require a high embedding dimension. Such an organization would facilitate linear readouts without interference and would therefore be more expressive (Rigotti et al. 2013; Fusi, Miller, and Rigotti 2016; Jazayeri and Ostojic 2021; Badre et al. 2021). The advantages of high-dimensional embeddings for linear readouts are well understood in artificial networks and widely used in the context of kernel methods and support vector machines (Boser, Guyon, and Vapnik 1992; Maass 2016; Cohen et al. 2020).

While there is evidence of some degree of expressivity in M1 (Churchland et al. 2012; Kaufman et al. 2014; Saxena and Cunningham 2019), our findings of low-dimensional manifolds implies that mixed selectivity is lower in M1 than in some higher-order brain areas. For example, a brain area that must selectively relay the abstract variables relevant to decision making and ultimately leading to motor output, such as the dorsolateral prefrontal cortex, requires high mixed selectivity and has been reported to have a large embedding dimensionality (Rigotti et al. 2013; Fusi, Miller, and Rigotti 2016). The low dimensionality and mild nonlinearity of M1 indicates that such complex representations are unnecessary for M1. Instead, M1 exhibits a more generalized, low-dimensional representation that encodes different inputs into a small set of common activity patterns (Fusi, Miller, and Rigotti 2016; Badre et al. 2021). Thus, from a functional perspective, the low-dimensional and generalizable computational strategy of M1 facilitates the reliable generation of muscle commands that are largely unaffected by task constraints.

Another interpretation of the low dimensionality and mild nonlinearity of M1 neural population activity refers to the strength of recurrent connectivity in M1. The dynamical systems perspective highlights the existence of recurrent connections, but we do not know much about

the strength of these connections (Vyas et al. 2020). Recent work on artificial and biological networks related the strength of recurrent connections to the embedding dimensionality of neural representations (Litwin-Kumar et al. 2017; Mastrogiuseppe and Ostojic 2018; Schuessler, Dubreuil, and Mastrogiuseppe 2020; Beiran et al. 2021; Pollock and Jazayeri 2020). In this view, networks with low-rank connectivity matrices generate low-dimensional dynamics. While the body of work that relates recurrent connectivity structure to dynamics is in its infancy and warrants future studies, the low dimensionality of M1 is consistent with weak recurrent connections (Langdon, Genkin, and Engel 2023).

Decoding EMGs from low-dimensional manifolds

Recent theoretical work directly relates the dimensionality of M1 manifolds to the accuracy of movement parameter decoding (P. Gao and Ganguli 2015; P. Gao et al. 2017). In this view, the low dimensionality of M1 permits sampling far fewer than the millions of active neurons for accurate decoding. While this theoretical work was verified in laboratory settings, the relationship between low dimensionality and decoding accuracy had not been shown in more natural settings. We showed that decoding EMGs from latent variables works reasonably well in either context **(Figure 6, Supplementary Figures 5 and 6**), indicating that the proposed theoretical ideas also apply in unconstrained settings. Our results indicate that current hardware recording technology that samples only hundreds to thousands of neurons allows for the identification and description of neural manifolds associated with unconstrained tasks. Generalizable representations in M1 are low dimensional and in turn decodable even when our recording devices vastly undersample the active neurons in the motor cortex. These findings contrast with the recent findings about dimensionality and decoding accuracy in the prefrontal cortex (PFC). The high embedding dimensionality of PFC representations make it difficult to decode behavioral parameters form neural data acquired with current recording technology,

despite moment-by-moment changes in neural activity during behavior (Bhandari, Gagne, and Badre 2018); in some cases, decoding behavior from PFC was barely above chance levels.

Limitations

When the cage experiments were designed, we anticipated the unconstrained tasks to be more complex, and hoped to quantify this expectation through a higher intrinsic dimensionality of EMG signals. However, we failed to see the anticipated effect. For Monkey P, the EMG reconstruction accuracies with progressively increasing latent dimensionality of linear and nonlinear EMG manifolds were slightly lower for the treat grasp task than for the other behaviors, signaling higher task complexity (**Supplementary Figure 4a**). The reconstruction curves for bar walk overlapped with those of the constrained tasks. For Monkey G, the EMG reconstruction curves associated with both bar walk and treat grasp tasks appeared to be below those for constrained tasks (**Supplementary Figure 4b**). However, for all tasks and both monkeys, roughly 4 to 5 latent dimensions were sufficient to explain over 0.70 of the variance in EMG signals. Importantly, the TNN dimensionality of EMGs corresponding to constrained and unconstrained settings were similar, and between 2 and 6 for both monkeys (**Supplementary Figure 3**). These results indicate there was no significant difference in task complexity between constrained and unconstrained tasks, in contrast to our original expectations. Therefore, one limitation of our study was that all tasks that we tested were relatively simple, even the unconstrained tasks in the more naturalistic cage setting. A definitive comparison between neural manifold associated with simple versus complex tasks is still needed.

Interpreting the dimensionality and complexity of behaviors is difficult (Bialek 2020). Earlier studies showed that the embedding dimensionality of hand kinematics did not exceed 8 even when humans individually moved the joints on their hands (Todorov and Ghahramani 2004). However, a more recent study showed evidence that even low-variance Principal Components contain some degree of task-relevant information, leading to the proposal that the distribution of

eigenvalues of the covariance matrix of joint angles should not be truncated, and that the embedding dimensionality of everyday manual behaviors could be as high as 30 (Yan et al. 2020). Future studies should investigate how task complexity affects the dimensionality of the EMG signals associated with from very simple to highly dexterous hand gestures; such an investigation would help elucidate the current ambiguity in quantifying and interpreting behavioral complexity.

The use of linear models for decoding EMG signals might also be considered a limitation. Linear decoding is based on a weighted sum of neural signals; it is thus the projection of collective neural activity along a specific direction in the neural state space. Although this widely used approach (Carmena et al. 2003b; Pandarinath, Ames, et al. 2018; Glaser et al. 2020) has been superseded by nonlinear and recurrent alternatives (Glaser et al. 2019, 2020; Perkins et al. 2023; Deo et al. 2023), linear decoders are simple, interpretable, and effective, and provide a useful tool for hypothesis testing. In this study, we were interested not in the absolute EMG decoding accuracy achieved when all neurons are used as inputs, but in the relative accuracy achieved when only the latent variables associated with the low-dimensional manifolds are used as inputs. Linear decoders were amply sufficient to quantify this comparison. Future studies, such as those involving neural prostheses, might benefit from the use of nonlinear and recurrent models of decoding for improved EMG decoding accuracy (Deo et al. 2023), although the degree of improvement is likely to depend on task complexity and its effect on manifold nonlinearity.

Conclusion

Our study demonstrates that the low dimensionality of task-specific manifolds in primary motor cortex is not limited to constrained laboratory tasks but is also present in unconstrained tasks. This finding suggests that the low dimensionality of M1 manifolds reflects an intrinsic computational property of M1, rather than being a byproduct of constrained movements. Even if

the unconstrained tasks were relatively simple, to have established the existence of low-dimensional neural manifolds in M1 beyond our study.

Our study also revealed signatures of nonlinearity in the geometry of the M1 neural manifolds, indicating that the neural population activity can represent complex relationships between the latent variables. However, the population neural dynamics mostly explored nearly linear regions of the neural manifolds. This finding is consistent with our results demonstrating the ability to linearly decode EMG signals from relatively few linear latent variables, and provides a manifold-based understanding of the surprising effectiveness of linear methods in decoding the motor output in brain-machine interfaces (Carmena et al. 2003b; Pandarinath, O'Shea, et al. 2018; Glaser et al. 2020; Perkins et al. 2023).

Additionally, our study emphasizes the importance of distinguishing between intrinsic and embedding dimensionalities when analyzing neural population activity. Our results suggest that the low embedding dimensionality of M1 manifolds reflects that the processing and communication requirements of M1 are relatively simple in comparison to those of other brain areas. For example, areas such as prefrontal cortex, parallel fibers in the cerebellum, and the primary visual cortex exhibit quite higher embedding dimensionalities. Future studies should explore the relationship between neural dimensionality, manifold nonlinearity, and information processing, both across different brain regions and across behavioral contexts of distinct complexity.

Overall, our study provides new insights into the computational properties of primary motor cortex and highlights the potential of low-dimensional representations for decoding motor output. The insights gained from this research have implications for extending the applicability of neural prosthetics and brain-machine interfaces to natural environments, while contributing to a broader understanding of how the motor cortex represents and processes movement-related information.

Supplementary Figures

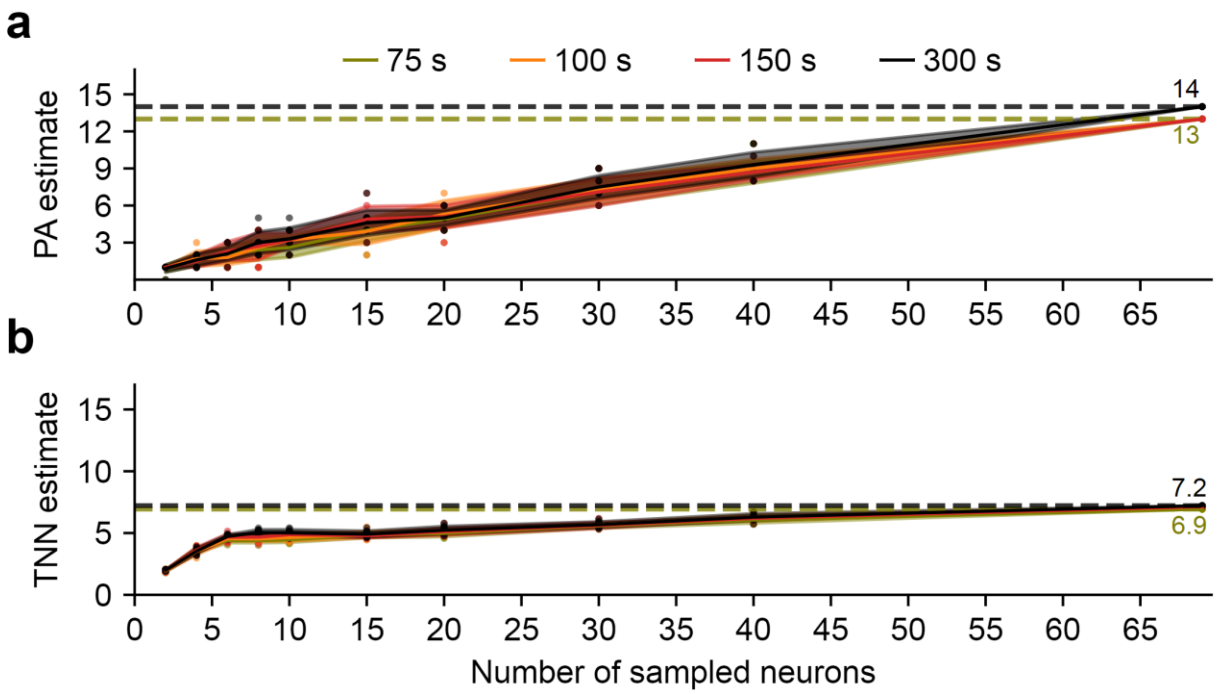

Supplementary Figure 3-1: The effect of temporal length of data and number of neurons on estimating embedding and intrinsic dimensionalities. a) Parallel Analysis (PA) estimates for the embedding dimensionality with increasing lengths of temporal data and number of neurons for Monkey P doing the bar walk. Colors represent the different amounts of temporal data used for the dimensionality estimate, ranging from 1500 bins of neural data corresponding to 75 seconds (binned at 50 ms) to 6000 bins of neural data corresponding to 300 seconds. The analysis was repeated 10 times, sampling a subset of neurons ranging from 2 to 69 (maximum neurons available for this session). b) Same as in panel a, but Two Nearest-Neighbors (TNN) estimates for the intrinsic dimensionality instead of PA estimates for the embedding dimensionality.

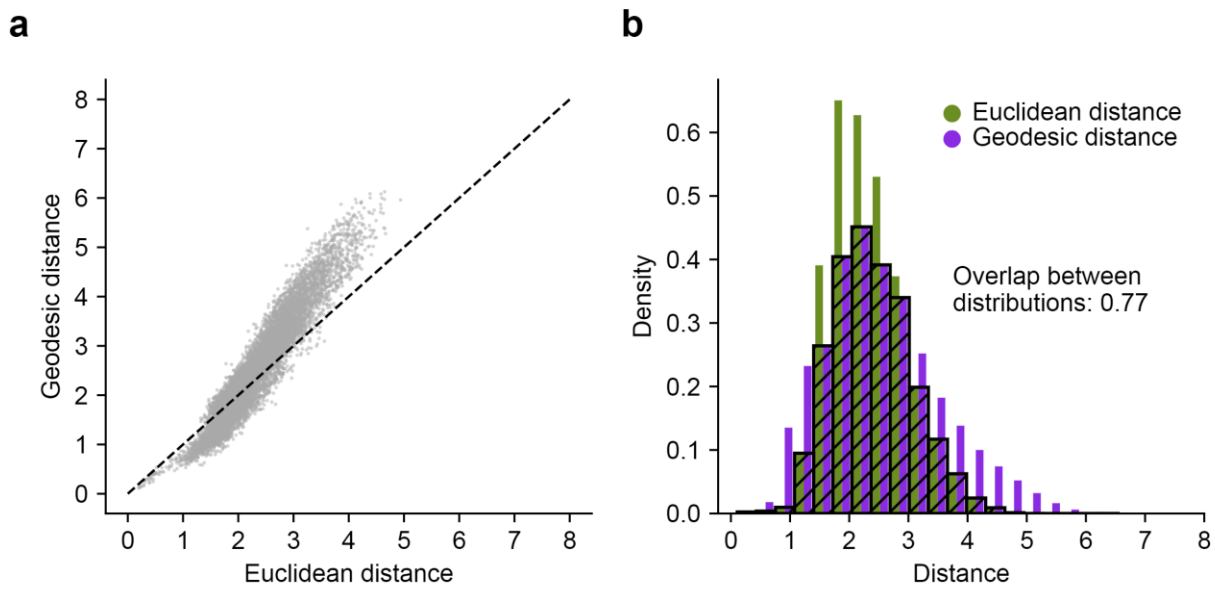

Supplementary Figure 3-2: Computation of the local flatness index. a) Comparison between the Euclidean and geodesic distances for all pairs of neural population data points recorded for Monkey P while performing the bar walk task. b) The distributions of Euclidean (olive) and geodesic (violet) distances for the same task. The distributions were normalized to represent probability densities. The sum of the overlap between the two distributions (shown as hatched black bars) is the local flatness index, equal to 0.77 for this dataset.

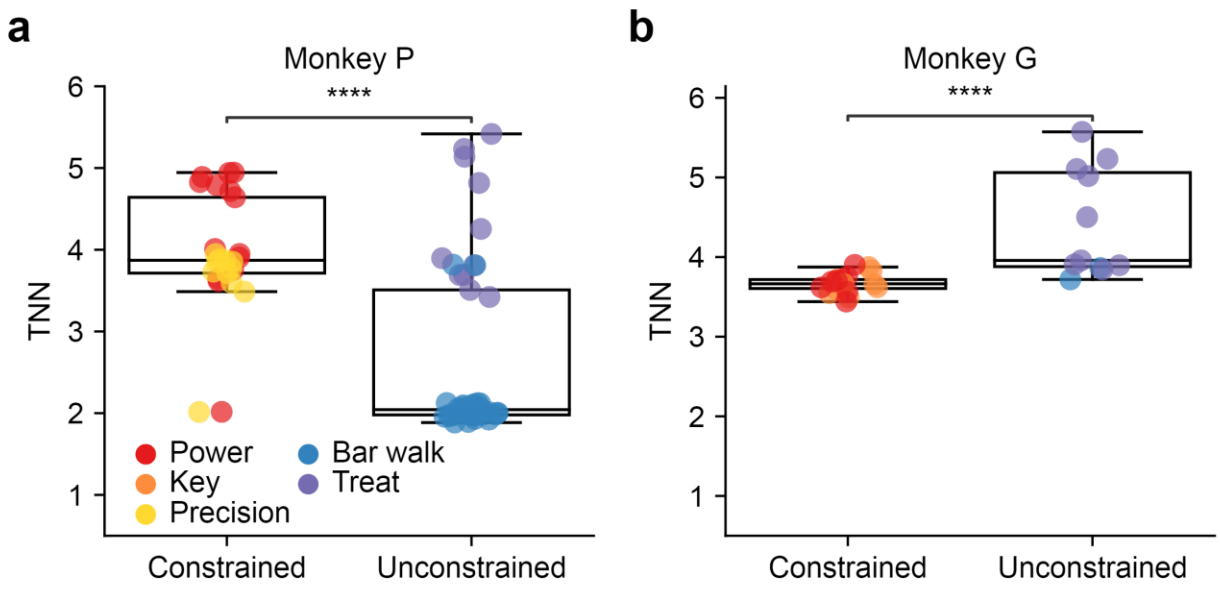

Supplementary Figure 3-3. Intrinsic dimensionality of EMG signals. The intrinsic dimensionality of the EMG signals was estimated using Two Nearest-Neighbors (TNN). Different colors represent different tasks. Warm colors indicate data collected in the constrained laboratory setting and cold colors indicate data collected in the unconstrained cage setting. a) Intrinsic dimensionality of EMG for Monkey P ($p \approx 0$). b) Same as in panel a, but for Monkey G ($p \approx 0$).

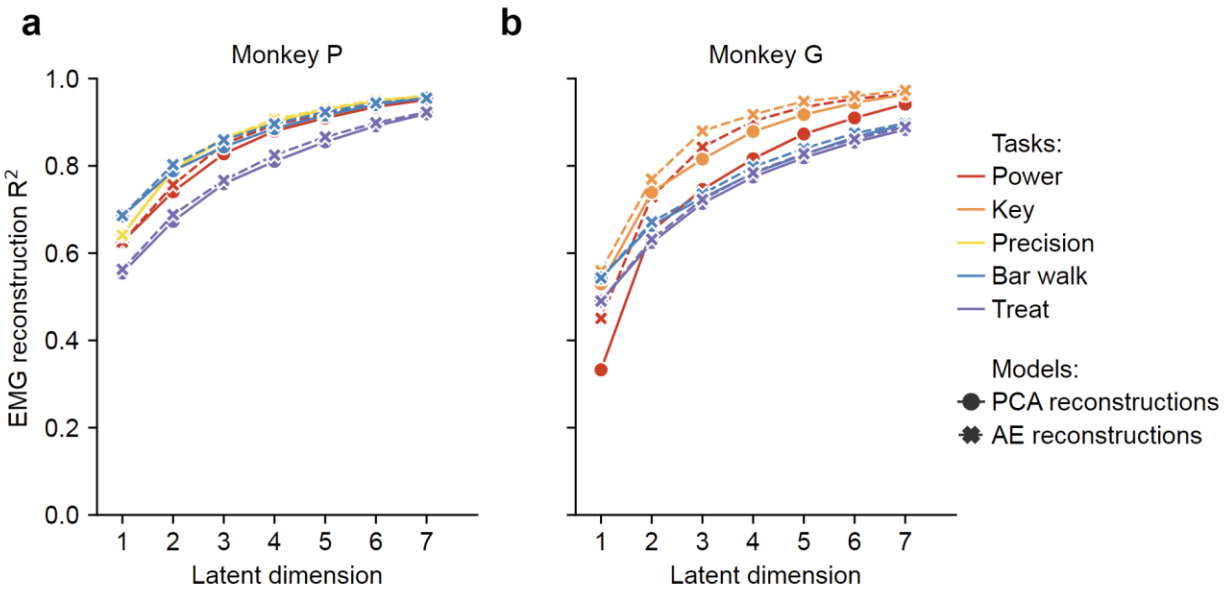

Supplementary Figure 3-4: EMG reconstruction accuracies with progressively increasing EMG manifold dimensionality. a) We progressively increased the latent dimensionality of the linear (Principal Component Analysis, PCA) and nonlinear (autoencoder, AE) EMG manifolds of Monkey P performing a variety of tasks. Each color corresponds to a different task. All results for a given task were averaged. Circle and cross symbols indicate the reconstruction accuracy of the EMG data from the low dimensional EMG manifolds of varying latent dimensionality using PCA and autoencoder (AE), respectively. b) Same as in panel a but for Monkey G.

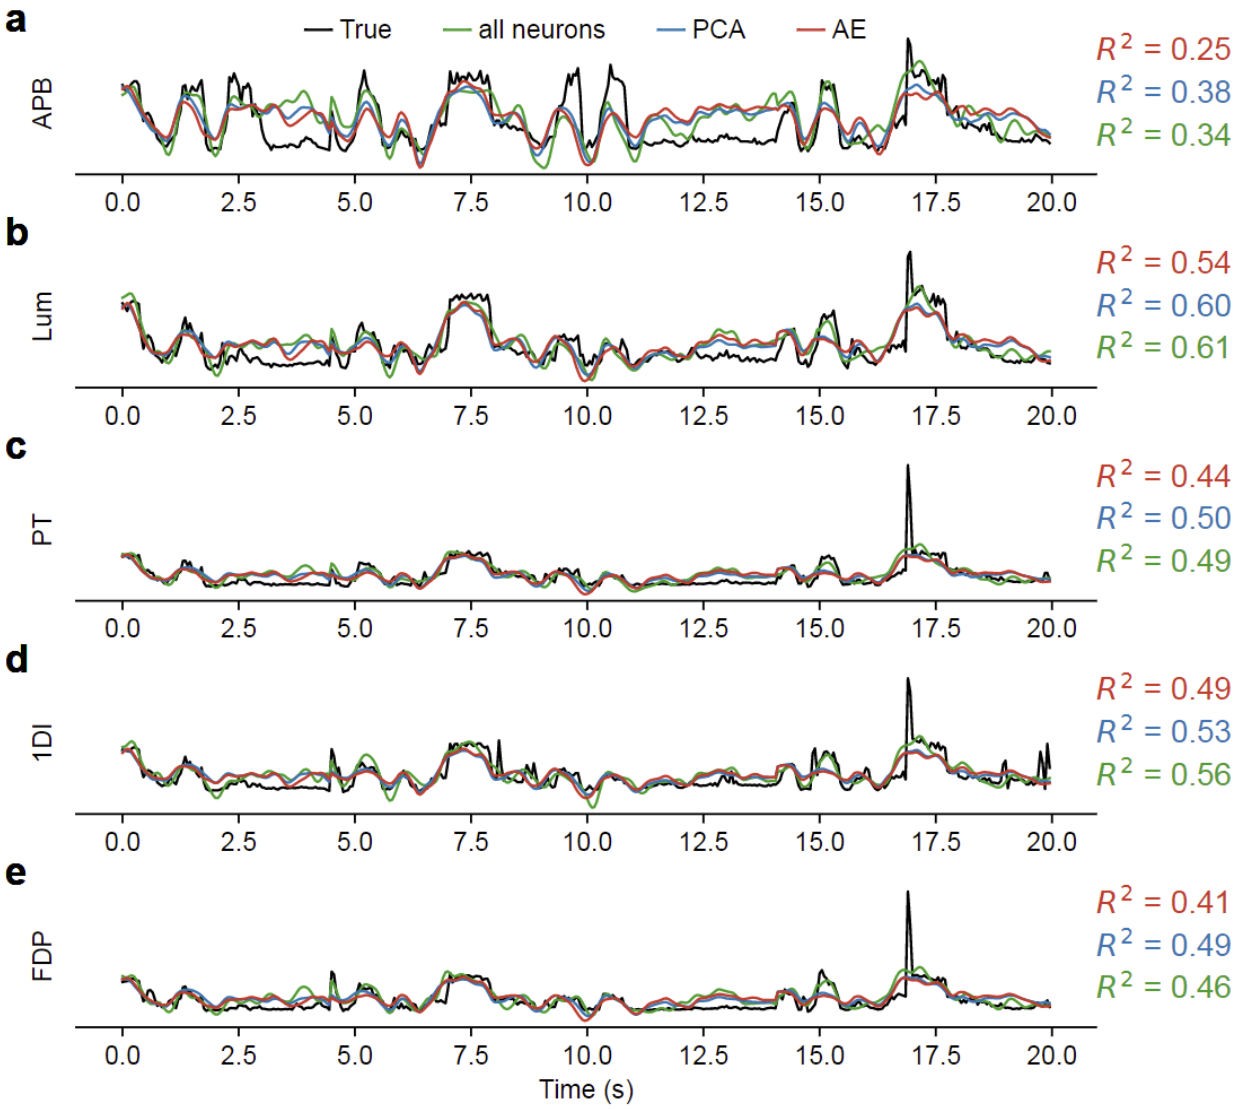

Supplementary Figure 3-5: EMG prediction examples for Monkey P performing the bar walk task in the cage. Each panel corresponds to the predictions of intramuscular EMG recorded from a different muscle. Black traces represent the normalized EMG activity. Green traces represent EMG activity predictions using all recorded neurons. Blue and red traces represent EMG predictions from low-dimensional linear (PCA) and nonlinear (autoencoder) latent variables, respectively. Corresponding R^2 values are denoted on the right of each panel. EMG predictions for a) Abductor Pollicis Brevis (APB), b) Lumbricals (Lum), c) Pronator teres (PT), d) First dorsal interossei (1DI), e) Flexor Digitorum Profundus (FDP).

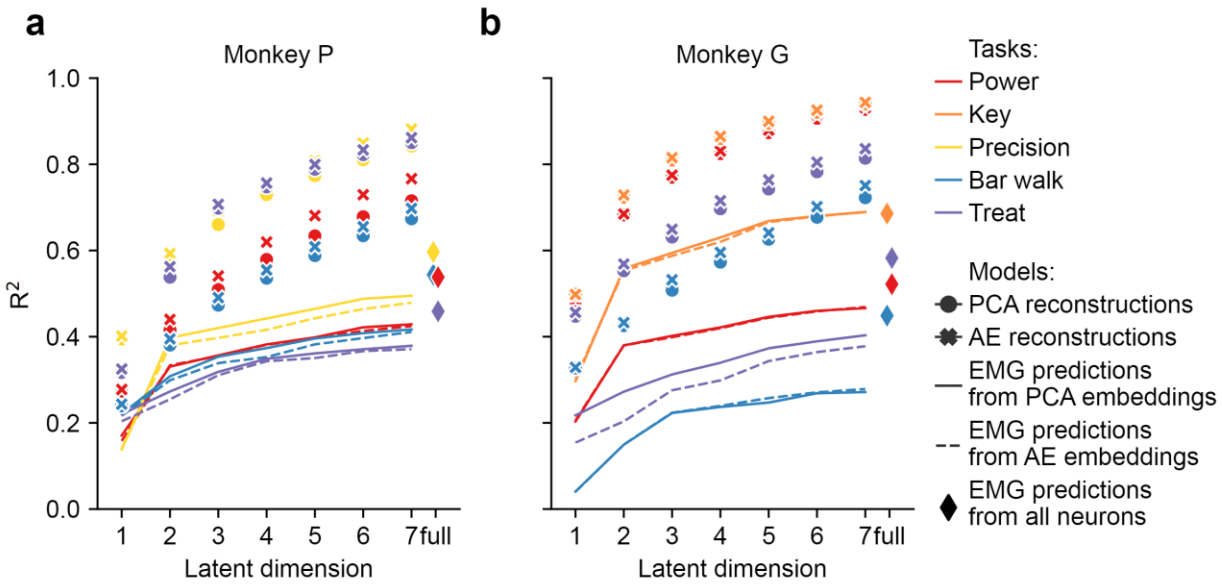

Supplementary Figure 3-6. EMG predictions and neural reconstruction accuracies with progressively increasing neural manifold dimensionality. a) We progressively increased the latent dimensionality of the linear (Principal Component Analysis, PCA) and nonlinear (autoencoder, AE) neural manifolds of Monkey P performing a variety of tasks. Each color corresponds to a different task. All results for a given task were averaged. Solid lines indicate average EMG predictions from linear neural manifolds of varying latent dimensionality. Dashed lines indicate average EMG predictions from nonlinear neural manifolds of varying latent dimensionality. Diamond symbols on the right indicate the average EMG predictions using all available neurons. All EMG predictions shown in the figure are the averages of the test folds of five-fold cross validation. Circle and cross symbols indicate the reconstruction accuracy of the neural data from the low-dimensional neural manifolds of varying latent dimensionality using PCA and autoencoder (AE), respectively. b) Same as in panel a but for Monkey G.

CHAPTER 4: FROM MONKEYS TO HUMANS: OBSERVATION-BASED EMG BRAIN-COMPUTER INTERFACE DECODERS FOR HUMANS WITH PARALYSIS

Fabio Rizzoglio*[,1,+], Ege Altan*[,1,2], Xuan Ma[1], Kevin L. Bodkin[1], Brian M. Dekleva[3], Sara A. Solla[1,4], Ann Kennedy[1], Lee E. Miller[1,2,5,6]

[1] Department of Neuroscience, Northwestern University, Chicago, IL, United States of America

[2] Department of Biomedical Engineering, Northwestern University, Evanston, IL, United States of America

[3] Rehab Neural Engineering Labs, Department of Physical Medicine and Rehabilitation, University of Pittsburgh, Pittsburgh, PA, United States of America

[4] Department of Physics and Astronomy, Northwestern University, Evanston, Illinois, United States of America

[5] Shirley Ryan AbilityLab, Chicago, IL, United States of America

[6] Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL, United States of America

* Equal contributions

Foreword

An earlier version of this work has been submitted to *Nature Biomedical Engineering* on November 2, 2022.

Abstract

Intracortical brain-computer interfaces (iBCIs) rely on “observation-based” decoders to enable individuals with paralysis to control the movement of virtual limbs and robotic arms. A large fraction of persons with high-level spinal injury indicate they would undergo brain surgery if it would restore some use of their own hands. However, the current reliance on an observed target motion for decoder development precludes its application to the prediction of unobservable motor output like muscle activity. We tested two possible solutions to this problem, each using data from a human iBCI user and a monkey who were performing similar motor actions. In one approach, we simply mapped neural signals recorded from the human to electromyographic (EMG) activity of the monkey. We used a second approach, based on the

hypothesis that the low-dimensional “latent” neural representations of motor behavior, known to be preserved across time for a given behavior, might also be preserved across individuals. We “transferred” an EMG decoder, trained solely on monkey data to the human iBCI user after using Canonical Correlation Analysis to align the human latent signals to those of the monkey. We found that both direct and transfer decoding approaches allowed accurate EMG predictions both between pairs of monkeys and from a monkey to a human. Our findings suggest that these latent representations of behavior are consistent across animals and even species. These methods are an important initial step in the development of iBCI decoders that generate EMG predictions for individuals with paralysis. These predictions could serve as signals for a biomimetic decoder controlling motion and impedance of a prosthetic arm or even muscle force directly through Functional Electrical Stimulation.

Introduction

Intracortical brain-computer interfaces (iBCIs) promise to restore voluntary movement to persons with paralyzed limbs. A kinematic iBCI uses a “decoder” to transform neural activity into signals that can be used to control a cursor or a robotic limb (Carmena et al. 2003a; Collinger et al. 2013; Musallam et al. 2004; Serruya et al. 2002; Taylor, Esnouf, and Hobby 2002; Wodlinger et al. 2014). In a proof of principle, EMG predictions were used to reanimate the temporarily paralyzed muscles of monkeys’ hands by using them to control the intensity of Functional Electrical Stimulation (FES) (Ethier et al. 2012). FES is used in a variety of clinical applications for both lower (Kralj, Bajd, and Turk 1988; Luo et al. 2020; M. R. Popovic et al. 2001; Triolo et al. 1996) and upper extremities (Eraifej et al. 2017; Marquez-Chin and Popovic 2020; Taylor, Esnouf, and Hobby 2002). None of the current clinical applications are brain-controlled, but two recent studies implemented rudimentary FES control of hand (Bouton et al. 2016) or arm and hand (Ajiboye et al. 2017) indirectly, driven by a simple grasp classifier or a kinematic decoder respectively.

Decoders are typically computed with supervised learning methods, tuning parameters to minimize the error between decoder-predicted and a motor-related output (Collinger et al. 2013; Hochberg et al. 2006; Willett et al. 2021). For iBCI applications with intact monkeys, the decoder can be computed between M1 firing rates and actual movement (or muscle activity). For application to a paralyzed human with no motor output, this is not possible. In this case, the standard approach is to use an “observation-based” decoder, computed to map M1 activity into an observed (and presumably attempted) kinematic trajectory (Hochberg et al. 2006; Willett et al. 2019). This approach is not feasible for a decoder that maps neural activity into EMGs, which cannot be observed.

There are two potential solutions to this problem, both of which involve a “target” human iBCI user with a paralyzed arm attempting the same type of movement performed by another “source” human, or even monkey. One approach is to create a direct mapping between the iBCI user’s recorded neural and the recorded “source” EMG signals. This approach is similar to an observation-based decoder, except the neural signals are mapped to the recorded EMG signals instead of observed endpoint kinematics. An advantage of this “direct decoding” approach is its simplicity. It essentially mimics the well-established observation approach, for movement-related signals that can’t actually be observed. A disadvantage is that it necessarily requires a close match between the movements (or motor actions, more generally) attempted by the source and target individuals.

An alternative solution might be to use a single, fixed decoder trained on a source monkey, where both neural and EMG activity can be recorded, and then inputting to that decoder, M1 signals recorded from the target iBCI user. This would obviously not be possible using single-neuron firing rates as inputs, since they differ completely for the source and target individuals. A possible solution might be found through the use of the latent signals that exist within a low-dimensional “neural manifold” computed from neural firing rates (Churchland et al. 2012;

Cunningham and Yu 2014; P. Gao and Ganguli 2015; Gallego et al. 2017). Several groups have shown that, despite being embedded in different coordinate systems, the latent signals computed from neural recordings made in different experimental sessions can be “aligned” to one another. By this method, a single, fixed decoder using aligned latent signals as inputs can remain stable, despite changing neurons, for many months (Degenhart et al. 2020a; Farshchian et al. 2018; Gallego et al. 2020; Karpowicz et al. 2022; Ma et al. 2022). Similar alignment methods might allow the latent signals from human and monkey to be aligned to each other, enabling what we call “transfer decoding”.

We tested these two approaches using data collected from three monkeys performing a version of the standard center-out task requiring production of forces in eight directions about the wrist. After developing and testing the alignment methods across monkeys, we applied them to a source monkey and a target human with a paralyzed arm. Both direct and transfer decoding produced accurate EMG predictions, though the direct method was consistently higher. Critically, the accuracy of EMG predictions was tied closely to the similarity between latent neural signals of the source and target. We consider these findings to be supportive of future clinical implementations enabling the use of more ‘biomimetic’ iBCI decoders that could either directly restore voluntary arm movements through FES or allow the control of an anthropomorphic prosthetic arm’s motion and impedance.

Methods

Monkey task and recordings

Neural and muscle activity data were collected from three adult male rhesus macaque monkeys. All surgical and experimental procedures were approved by the Institutional Animal Care and Use Committee (IACUC) of Northwestern University.

We considered data from three different monkeys on four different recording sessions (within a span of 47 days), for a total of 12 sessions. In each session, the monkeys were seated in a primate chair that faced a computer monitor, with their left hand secured within a small box instrumented with a 6 degree of freedom load cell (JR3 Inc., CA) to measure the forces exerted (**Figure 4-1a**). The load cell measurement axes were aligned with those of the wrist so that flexion/extension forces moved a cursor on the monitor right/left, while radial/ulnar deviation forces moved it up/down. The monkeys were required to move the cursor from a central target towards one of eight peripheral targets uniformly distributed on a circle. A trial started when the monkeys held the cursor in the central target for a random time between 0.2 s and 1.0 s. Then, one of the eight peripheral targets was selected randomly and presented together with an auditory go cue. Monkeys had to reach the outer target within 2.0 s and maintain that force for 0.8 s to receive a liquid reward. In this study, we used data within a window starting 0.5 s before onset of cursor movement and ending 1.0 s after movement onset (see **Supplemtary Figure 4-1** for examples of cursor trajectories).

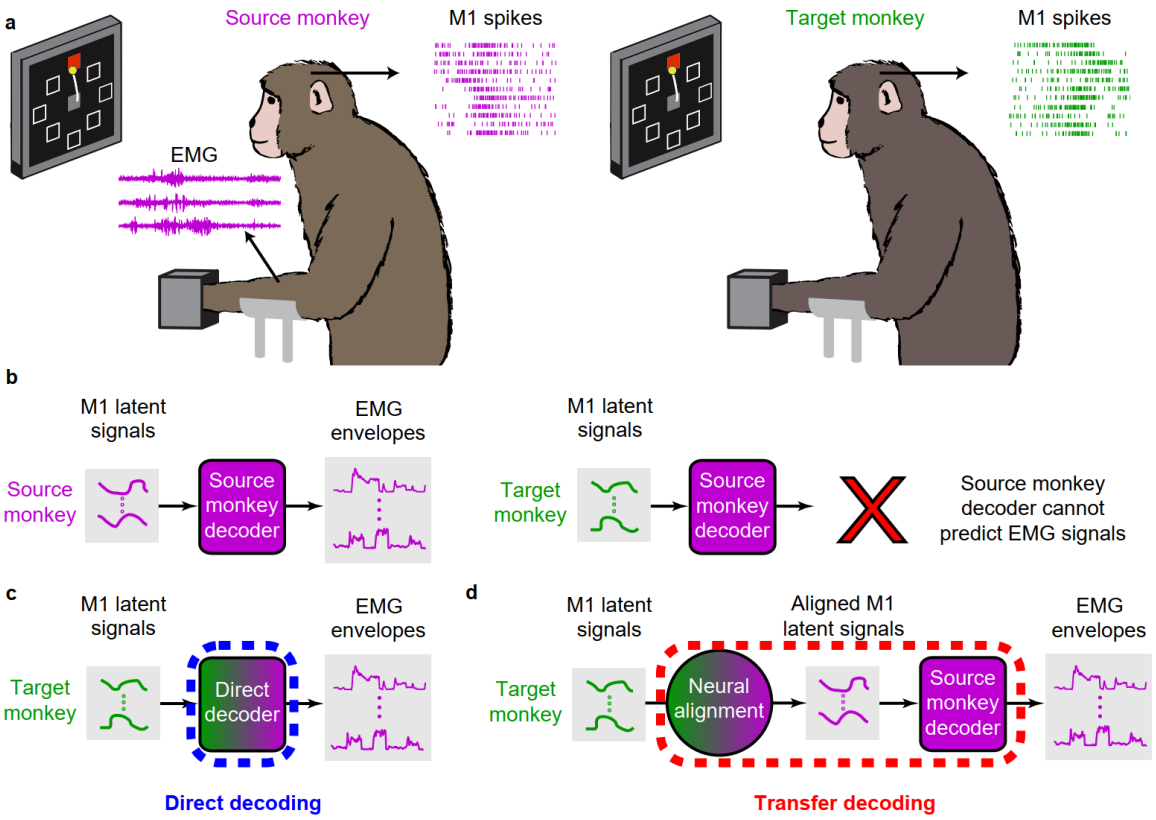

Figure 4-1: Cross-user decoding of EMG. a) We recorded neural firing rates from M1 and EMGs from forearm and wrist muscles of a "source" monkey trained to perform the isometric wrist task. We also recorded M1 data from a "target" monkey. b) The firing rates of the source monkey were projected in a low-dimensional neural manifold and the resulting latent signals used as input to train a source-monkey iBCI decoder. c) In a first approach, we obtained source monkey EMG predictions by training a decoder between the target monkey M1 signals and the source monkey EMG signals. We call this approach "direct decoding." d) In a second approach, we computed a decoder soley from the source monkey data, then preprocessed the target M1 latent signals to align them to those of the source monkey. We used the aligned latent signals as input to the source-monkey decoder to obtain predicted EMG. We call this approach "transfer decoding."

Each monkey was implanted with a 96-channel Utah electrode array (Blackrock Neurotech, Inc.) in the hand area of the right M1, contralateral to the hand used for the task, using standard surgical procedures. In a separate procedure, monkeys were also implanted with intramuscular leads in the forearm and hand muscles of the left arm. The location of each electrode was verified during the surgery by observing the muscle contraction evoked when it was stimulated. Five muscles were implanted in all three monkeys: three major wrist muscles: extensor carpi radialis longus (ECRl), flexor carpi radialis (FCR), and flexor carpi ulnaris (FCU), and two extrinsic finger muscles: extensor digitorum communis radialis (EDCr) and flexor digitorum profundus (FDP). Beside these muscles that were in common across monkeys, we implanted additional muscles in each monkey. These included extensor carpi radialis brevis (ECRb) and extensor carpi ulnaris (ECU) in monkey J, flexor digitorum indicis (FDI) and opponens digiti minimi (MD) and FDP (2 pairs) in monkey S, and ECU, extensor digitorum communis (EDC), flexor digitorum superficialis (FDS), pronator teres (PT), and supinator (SUP) in monkey K.

We recorded M1 activity using a Cerebus system (Blackrock Neurotech, Inc.). The signals on each channel were sampled at 30 kHz, digitally bandpass filtered (250 ~ 5000 Hz) and converted to spike times based on threshold crossings. The threshold on each channel was set with respect to its root-mean square (RMS) amplitude (monkey J: -5.5 x RMS; monkey S: -6.25 x RMS; monkey K: -6.0 x RMS). To extract the smoothed firing rate used in the analysis, we applied a Gaussian kernel (standard deviation of 100 ms) to the spike counts in 20 ms, non-overlapping bins for each channel.

The recorded EMG signals were amplified, bandpass filtered (4-pole, 50 ~ 500 Hz), and sampled at 2 kHz. To extract the envelopes used during the analysis, we subsequently rectified and lowpass filtered (4-pole Butterworth, 10 Hz) each EMG channel digitally and subsampled it to 50 Hz to correspond to the bin size of the M1 signals. EMGs were clipped to avoid data points larger than the mean plus 6 times the standard deviation of each channel. We removed

the baseline of each EMG channel by subtracting the 2nd percentile of its amplitude and then normalized its activity to its 90th percentile.

Human participant task and recordings

The participant, male, 28 years old at the time of the implant, was part of a multi-site clinical trial (NCT01894802) and provided informed consent prior to the experimental procedure. He presented with a C5 motor/C6 sensory ASIA B spinal cord injury that occurred 10 years prior. He had no spared control of the intrinsic or extrinsic muscles of the right hand but had limited control of wrist flexion and extension. Proximal limb control at the shoulder was intact, as was elbow flexion. However, he had no voluntary control of elbow extension.

We secured a rigid wrist brace to the participant's right wrist and hand, oriented in a neutral posture on his lap. On each trial, the participant attempted to produce isometric wrist forces in response to the movement of a cursor to eight different radial targets displayed on a screen in front of him. Upward movement corresponded to radial deviation, rightward movements to wrist extension, downward movements to ulnar deviation, and leftward movements to wrist flexion. Diagonal targets corresponded to the appropriate combinations of these gestures. Each trial began with the presentation of the upcoming target. One second after target appearance, a go cue occurred, followed by movement of the cursor to the target (Move: 0.2 s), static hold at the target (Hold: 2.0 s), and return to center (Return: 0.2 s). We instructed the participant to produce step-like force profiles in response to the movement of the cursor. To match the time scale of the monkey data, here we only used data from a 1.5 s window, starting 0.76 s before and ending 0.74 s after the go cue (see **Figure 4-6b**).

The participant was implanted with four NeuroPort microelectrode arrays (Blackrock Neurotech, Inc.) in the left hemisphere. Two 96-channel arrays were implanted in the hand and arm areas of M1, while the other two were implanted in somatosensory cortex. In this study, we only used

the arrays placed in M1. The signals on each channel were recorded at 30 kHz and high-pass filtered at 750 Hz. Whenever the signal crossed a threshold (-4.25 x RMS), a spiking event was recorded. We used multiunit threshold crossing and considered spike counts in 20 ms non-overlapping bins. We applied a Gaussian kernel (standard deviation of 100 ms) to obtain the smoothed firing rate used in the analysis. For all data analyses, we mirrored the target labels for the human data as his cortical implant is in the opposite hemisphere as that of the monkeys, who used their left hands to perform the task.

Computation and evaluation of EMG decoders

In this study, we tested two cross-user decoding approaches: direct decoding and transfer decoding. Both types of decoders were trained to predict the source monkey's EMG signals, and used as inputs latent signals computed within a low-dimensional manifold of M1. The direct decoder was trained on target M1 signals (monkey or human; **Figure 4-1c**), whereas the transfer decoding was trained on M1 data from the source monkey (**Figure 4-1d**).

To find the latent signals, we performed dimensionality reduction by applying PCA to the M1 firing rates. We set the dimensionality of the latent space to 13, the largest estimate of linear dimensionality across all monkeys and sessions using Parallel Analysis (Altan et al. 2021). We computed Wiener filter decoders that implemented linear regression to predict the EMG at the current time bin from the latent neural signals stretching to five time bins (100 ms) into the past. The output of the linear filter was then rectified to more nearly match the statistics of the EMG envelopes.

We compared the resulting EMG predictions with the actual EMG recordings of the source monkey. Ideally, the accuracy of the cross-monkey and cross-species decoder as measured by R^2 would approach the performance of the corresponding within-monkey decoder. As the

decoded EMGs are multi-dimensional, we computed an R^2 for each EMG and then reported the average across all muscles, weighted by the variance of each recorded EMG.

CCA alignment of low-dimensional neural signals

The transfer decoder requires that we "align" the latent signals of the target monkey (or human) to those of the source monkey that was used to compute the decoder (**Figure 4-1d**). Similarly, in direct decoding, we tested alignment of the latent signals over time, as a means to address the steady decline in performance of a fixed decoder over time due to the inherent instabilities of the recorded neural signals (Downey et al. 2018; Perge et al. 2013; Sussillo et al. 2016). For both applications, we used Canonical Correlation Analysis (CCA). We labelled the latent neural spaces that we wished to align as $\boldsymbol{S}$ and $\boldsymbol{T}$. Both were *M x p* matrices, where *M* is the number of samples and *p* the latent dimensionality (*p=13*). The goal of neural alignment is to make $\boldsymbol{T}$ as similar as possible to $\boldsymbol{S}$. When we tested neural alignment with direct decoding, $\boldsymbol{S}$ and $\boldsymbol{T}$ were the neural signals associated with the same user but across sessions. In the transfer decoding case, $\boldsymbol{S}$ and $\boldsymbol{T}$ were associated with source and target users, respectively.

CCA linearly transforms both $\boldsymbol{S}$ and $\boldsymbol{T}$ such that the correlation between $\boldsymbol{C_S}(S)$ and $\boldsymbol{C_T}(T)$ is maximal, where $\boldsymbol{C_S}$ and $\boldsymbol{C_T}$ are the matrices that implement the linear transformations of the latent signals (**Figure 4-2a**). Since CCA requires a one-to-one correspondence in time between data points, we extracted single trials from the continuous data. We then ordered all trials by target direction and concatenated them prior to applying CCA. Finally, because CCA transforms both sets of signals, we further multiplied $\boldsymbol{C_T}(T)$ by the inverse transformation $\boldsymbol{C_S^{-1}}$ to represent $\boldsymbol{T}$ in the coordinate frame of $\boldsymbol{S}$. We refer to the resulting aligned latent signals $\widetilde{\boldsymbol{T}} = \boldsymbol{C_S^{-1}}(\boldsymbol{C_T}(T))$.

Cross-validation of direct and transfer decoding approaches

For both direct and transfer decoding approaches, we implemented cross-validation at the input level (i.e., the target user neural data) to avoid overfitting with what would typically be limited

data. However, it is important to note that this was implemented differently in the two methods. While direct decoding uses target M1 data during decoder computation (**Figure 4-1d**), transfer decoding employs it during the CCA alignment step (**Figure 4-1c**). Therefore, when training the direct decoder across monkeys, we performed 4-fold cross-validation, which resulted in a training set of 96 trials (12 trials and 18 seconds of data per target direction) and a test set of 32 trials (4 trials and 6 seconds of data per target direction) for each fold. For the more limited human data, we used only three folds, each consisting of three trials for each of the eight targets. Likewise, to test transfer-decoding alignment, we used four folds for the monkey-to-monkey case and three folds for human-to-monkey alignment. As noted above, we used all the available data to train the source monkey transfer decoder itself (128 trials in total, 16 trials and 24 seconds of data per target direction).

Results

Direct mapping between target M1 and source EMG signals enables cross-monkey decoding

Direct decoding offers a simple solution to the problem of mapping neural activity to EMGs in iBCI applications for individuals with a paralyzed arm and hand. Direct decoding establishes a mapping between the neural signals recorded from the iBCI user (the “target”) to EMG signals from a “source” individual as they each perform (or attempt to perform) the same movement (**Figure 4-1c**). We began our investigation of direct decoding using monkeys as both source and target.

We determined the EMG decoding accuracy by computing the coefficient of determination (R^2) between the EMG predictions and the actual EMG signals recorded from the source monkey. **Figure 4-2a** (blue lines) shows example EMG predictions from a single source-target monkey pair (source monkey: J, first session, target monkey: S, second session). The EMG decoding

accuracy ranged between R^2 = 0.48 and 0.77 for individual muscles. The average EMG decoding accuracy across all muscles, weighted by their individual variances, was 0.64. We compared these to the accuracy of decoders built from the same monkey and session, which ranged between 0.69 and 0.86 for the same muscles, with an average R^2 of 0.77 (grey lines).

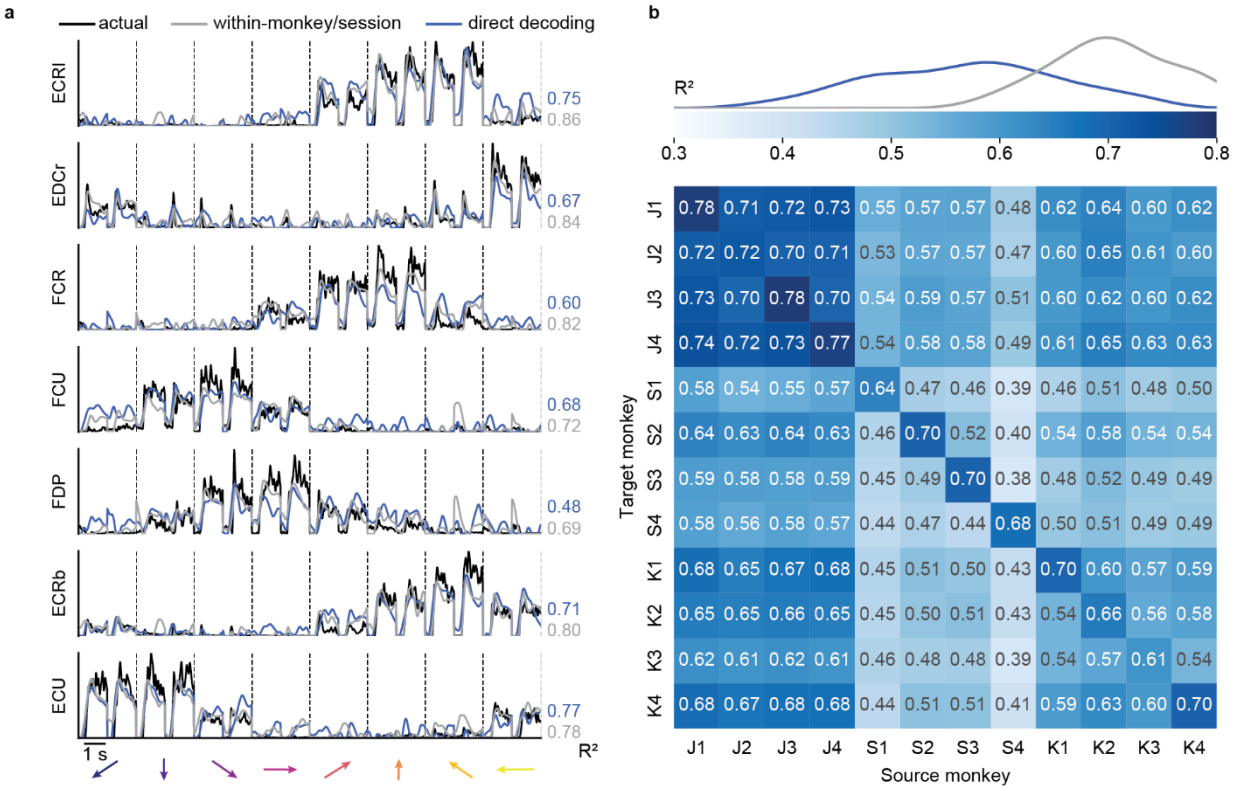

Figure 4-2: Direct decoding of EMG across monkeys. a) Cross-monkey EMG predictions obtained with target-monkey latent signals via direct decoder (blue lines). Data for a representative pair of monkeys for two trials in each target direction (directions separated by vertical dashed lines). These predictions are almost as good as those obtained by using the source-monkey latent signals as input to the decoder (grey lines). The R^2 for both within- (grey) and cross- (blue) monkey decoding are shown for each muscle, computed relative to actual EMG recordings of the source monkey (black lines). b), Overall cross-monkey decoding accuracy (R^2) with direct decoding for all pairs of monkeys. A kernel density estimate plot at the top shows the distribution of cross-monkey (off-diagonal) R^2 (blue) compared to within-monkey (on-diagonal) R^2 (grey).

Next, we extended the cross-monkey EMG decoding to all source-target monkey pairs (**Figure 4-2b**). The diagonal entries of the matrices in **Figure 4-2b** show the weighted averages across all muscles of the within-monkey decoding accuracy for each session. The off-diagonal elements represent either cross-monkey or within-monkey/cross-session accuracy. We included the within monkey/cross-session decoding accuracy as a more realistic comparison with the cross-monkey (and implicitly cross-session) analyses. The curves at the top of the panel indicate the kernel density estimate of the off-diagonal and on-diagonal accuracies. On average, cross-monkey direct decoding retained 80% of the accuracy of within-monkey decoding and 97% of the within-monkey/cross-session accuracy.

Latent signals across monkeys become similar after neural alignment

An alternative approach to obtain EMG predictions for individuals with paralysis is to use a fixed decoder trained on a source monkey from which both neural and EMG activity can be recorded, and transferring it for use with the M1 signals recorded from the target iBCI user. However, a decoder using neural firing rates as inputs cannot be used directly for another individual. Instead, we compute decoders based on inputs from low-dimensional latent signals computed from the firing rates. An essential step is to align the latent signals of the target iBCI user to those of the source monkey (**Figure 4-1d**) (Degenhart et al. 2020a; Farshchian et al. 2018; Gallego et al. 2020; Ma et al. 2022).

First, we used data from three monkeys to test the extent to which we could align the latent neural signals collected from a source monkey to those of a target monkey. **Figure 4-3a**, left column shows an example of the latent signals obtained by projecting the neural activity of the source and target monkeys into their corresponding neural manifolds (source monkey: J, first session, target monkey: S, second session). The unaligned latent signals of the two monkeys are similar, in that the trajectories corresponding to each of the eight targets are well-separated and traverse roughly parallel paths through their respective principal component spaces.

However, despite their similar shapes, the two sets of latent signals differ in scale and orientation.

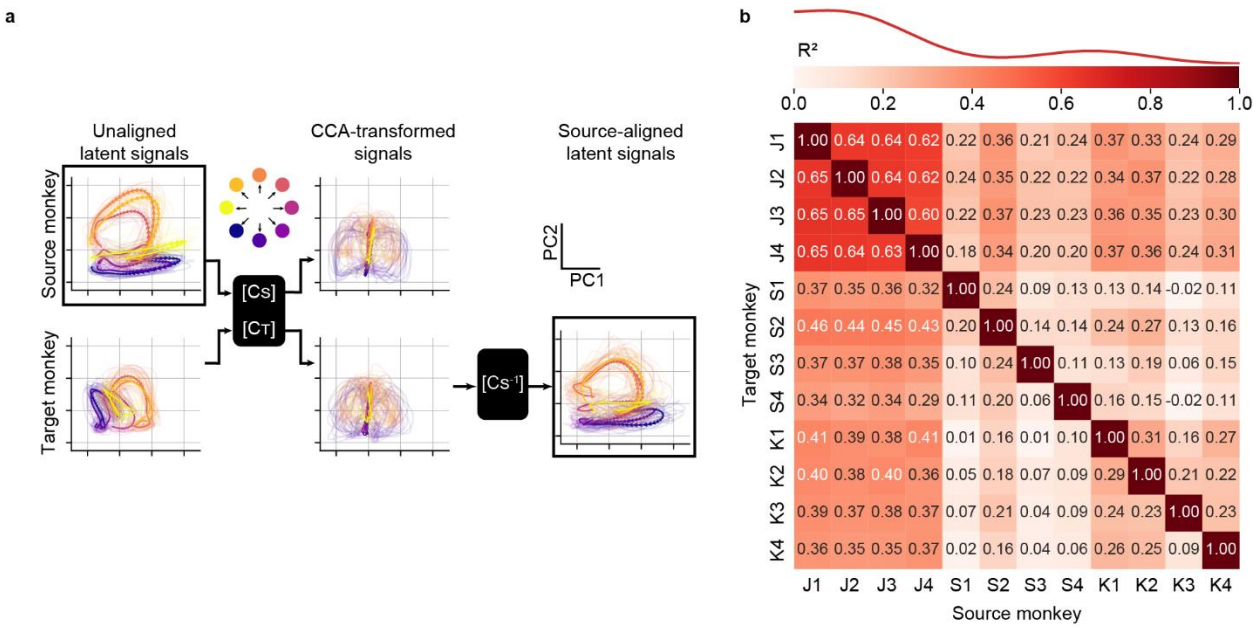

Figure 4-3: Latent neural signals becomes more similar across monkeys. a) Representative latent signals described by the first two principal components for a source (top left) and target (bottom left) monkeys. We used CCA to transform the latent signals such that they were maximally correlated (center). The CCA-transformed signals are more similar, but neither looks like the original source-monkey trajectories. Consequently, we further multiplied the target-monkey latent signals in the CCA-transformed space by the inverse of the source-monkey CCA weights (C_S^{-1}) to rotate them back to the original source-monkey latent coordinates (bottom right). Data were averaged across all trials for each target direction; single trial trajectories are shown as lighter traces. Arrows indicate the temporal evolution of the trajectories (from 0.5 s before to 1 s after cursor movement onset). b) Overall cross-monkey latent signals similarity (R^2) after CCA alignment for all pairs of monkeys. A kernel density estimate plot shows the distribution of cross-monkey latent signals R^2 similarity.

We next used Canonical Correlation Analysis to linearly transform both sets of low-dimensional signals such that they were maximally correlated (**Figure 4-3a**, middle column). The transformed sets of latent signals traverse very similar regions of the new (transformed) latent space, but neither looks like the original source monkey trajectories, as both were transformed by CCA. Consequently, we further processed the transformed target monkey latent signals by using the inverse of the canonical correlation transformation from the source monkey. This two-step process allowed us to align the target monkey signals to those of the source monkey (**Figure 4-3a**, right column). The aligned trajectories for the two monkeys (**Figure 4-3a**, upper left, lower right) are a remarkably close match.

We computed the coefficient of determination (R^2) between source monkey neural signals and CCA-aligned target monkey neural signals for the entire dataset comprising four sessions for each of three monkeys (**Figure 4-3b**). CCA alignment significantly improved the similarity of all source/target monkey latent signal pairs, both within-monkey/cross-session (R^2 before alignment: -0.37 ± 0.11; after alignment: 0.34 ± 0.04 (mean ± s.e.); P ~ 0, Wilcoxon's signed rank test) and cross-monkey (R^2 before alignment: -1.08 ± 0.08; after alignment: 0.25 ± 0.01 (mean ± s.e.); P ~ 0, Wilcoxon's signed rank test).

Neural alignment allows to transfer EMG decoder across monkeys

Next, we evaluated the performance of the transfer decoding approach using the CCA aligned input signals. Just as we did for direct decoding, we quantified EMG decoding accuracy by computing the coefficient of determination (R^2) between the EMG predictions and the actual EMG traces of the source monkey.

Transfer decoding results are summarized in **Figure 4-4** for all source-target monkey pairs. The average cross-monkey EMG decoding accuracy was 0.53, virtually identical to the average within-monkey/cross-session accuracy of 0.55. Overall, cross-monkey transfer decoding

retained 95% of the within-monkey/cross-session accuracy and 75% even of the within-monkey/within-session decoding accuracy.

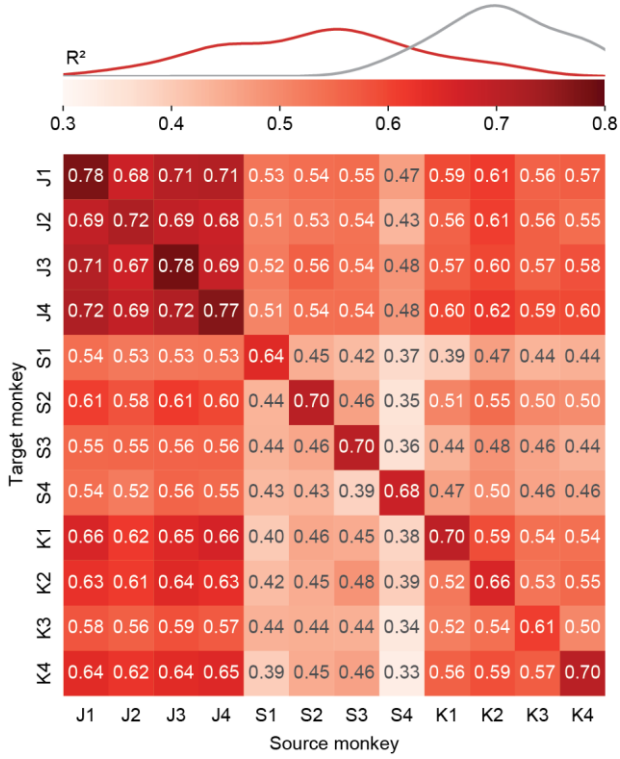

Figure 4-4: Transfer decoding of EMG across monkeys. Overall cross-monkey decoding accuracy (R^2) with transfer decoding for all pairs of monkeys. A kernel density estimate plot shows the distribution of cross-monkey R^2 (red) compared to within-monkey R^2 (grey) as in Figure 2.

Figure 4-5a shows example EMG predictions for both the direct and transfer approaches, as well as the actual EMGs of the source monkey for the same source-target monkey pair shown in **Figure 4-2 and Figure 4-3**. Although the accuracy of the two methods was quite similar, that of direct decoding was slightly higher.

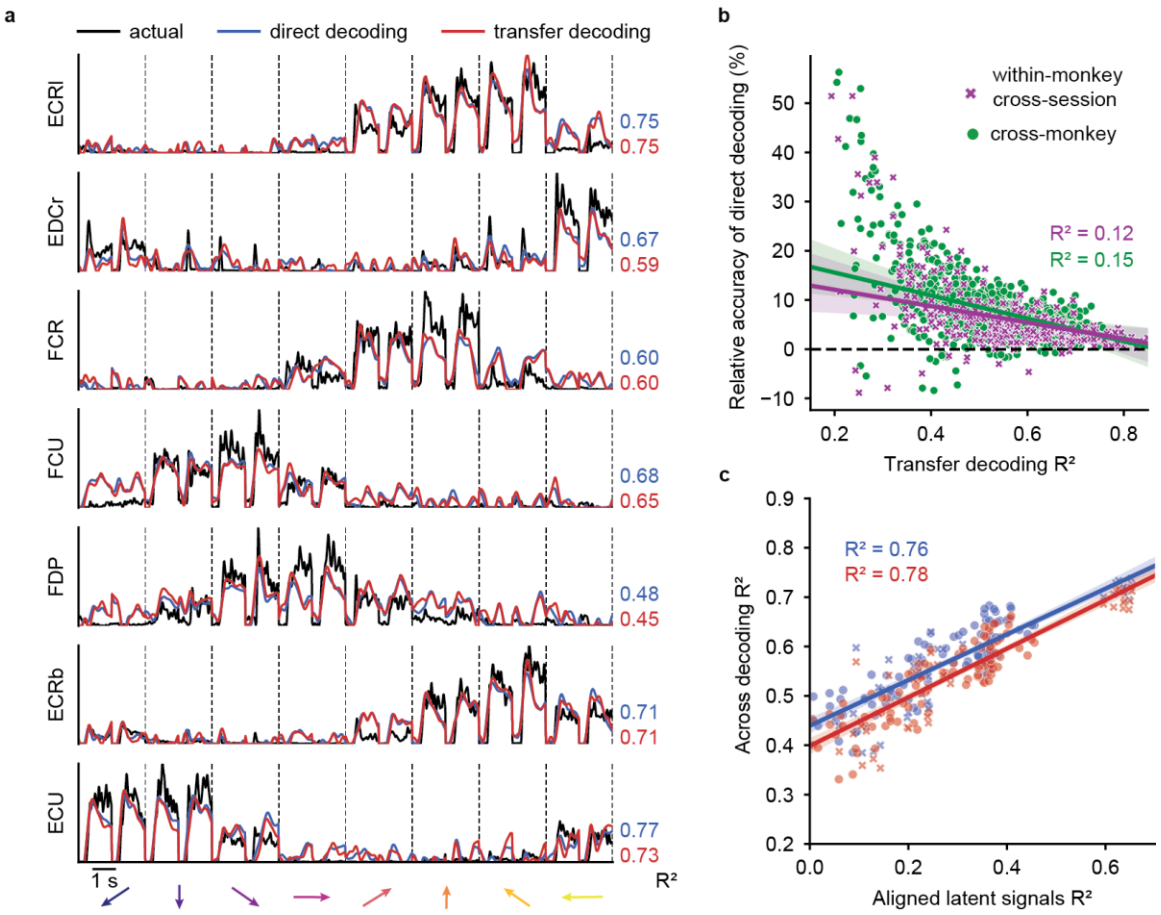

Figure 4-5: Comparing direct decoding and transfer decoding. a) Representative cross-monkey EMG predictions obtained with direct (blue lines) and transfer decoding (red lines) compared to source monkey ground truth (black lines) as in Figure 4-2a. b) Element-by-element scatter plot comparing direct and transfer decoding performance. Each point represents the within-monkey/cross-session (purple) or cross-monkey (green) accuracy of transfer decoding (x-axis) versus the relative accuracy of direct decoding (y-axis) of a single muscle. Direct decoding generally yielded higher decoding accuracy for both cross-monkey and cross-session predictions, especially when decoding accuracy was low. c) Element-by-element scatter plot for the matrices in Figure 4-4 and Figure 4-3b. Latent signals of monkey pairs with higher similarity after CCA alignment yielded higher decoding accuracy for both cross-monkey (purple) and cross-session (green) predictions.

We quantified these results for all monkey pairs in **Figure 4-5**. Each point in the scatter plot shows the R^2 for an individual muscle for both methods. For both the cross-monkey (**Figure 4-5b**, green symbols) and cross-session (purple symbol) cases, the direct decoding yielded a small but consistently higher prediction accuracy. Interestingly, the difference in decoding accuracy between direct and transfer decoding becomes smaller with increasing R^2.

We hypothesized that cross-monkey decoding accuracy with either approach would depend on the ability to achieve good neural alignment with CCA. The similarity between the EMG decoding accuracy matrices in **Figure 4-2a and 4-4a** and the neural alignment accuracy matrix in **Figure 4-3b** suggests that accuracy of decoder performance is dependent on successful alignment. We quantified the relationship by constructing an element-by-element scatter plot of the decoding and alignment matrices (**Figure 4-5c**). Indeed, both the cross-monkey (green) and cross-session (purple) decoding accuracy was well predicted by the success of the CCA alignment, with an R^2 of 0.76 for direct and 0.78 for transfer decoding across all datasets.

Monkey to human EMG decoding

Having demonstrated that cross-monkey EMG decoding is possible, we asked if the same two approaches could be implemented for cross-species decoding from a monkey to a human. We thus tried to either train a direct decoder or transfer a monkey decoder to predict EMG signals from the neural activity of a human with paralysis. For this analysis, we chose to use the first session from monkey J (dataset J1) as the source monkey, as it had the highest within-session decoder performance among all datasets.

We performed these experiments with one participant who had a partial spinal cord injury at C5/C6, who was implanted with two 96-channel electrode arrays, one each in the arm and hand areas of the left primary motor cortex (M1). We recorded multi-unit spiking activity as the participant, who had some remaining ability to produce wrist extension forces, attempted to

produce isometric wrist forces in eight radial directions. Cursor control and task design were essentially as in the monkey experiments, with the following exceptions. First, the participant attempted to use their right arm and hand (contralateral to the arrays), thereby reversing in world coordinates, the flexion and extension muscle activity relative to those of the monkey. Consequently, for direct decoder and CCA alignment computation, we mirrored the target directions for the human data about the vertical axis. Second, cursor movement was not under the control of the participant, but occurred automatically as in the standard approach to "observation" decoder training (Hochberg et al. 2006; Willett et al. 2019). One second after target appearance, a go cue occurred and the cursor moved to the target in 0.2 s, where it remained for 2.0 s before returning to the center in 0.2 s. We instructed the participant to attempt to produce the forces necessary to control the position of the cursor.

Figure 4-6a shows the time course of the neural activity of the hand M1 array projected onto the first principal component during trials corresponding to two oppositely directed targets. The human data was more variable across trials than that of the monkey in both timing and magnitude, likely due to the lack of actual force and force-related feedback. To minimize the effect of this variability on the cross-species decoding process, we computed the EMG decoding R^2 as a function of the time window we used for direct decoding and CCA alignment computation, for the direct decoder and transfer decoding approaches respectively. We achieved maximal decoding accuracy when aligning the human trials 0.84 s and 0.78 s prior to the go cue (**Figure 4-6b**) for direct and transfer decoding respectively.

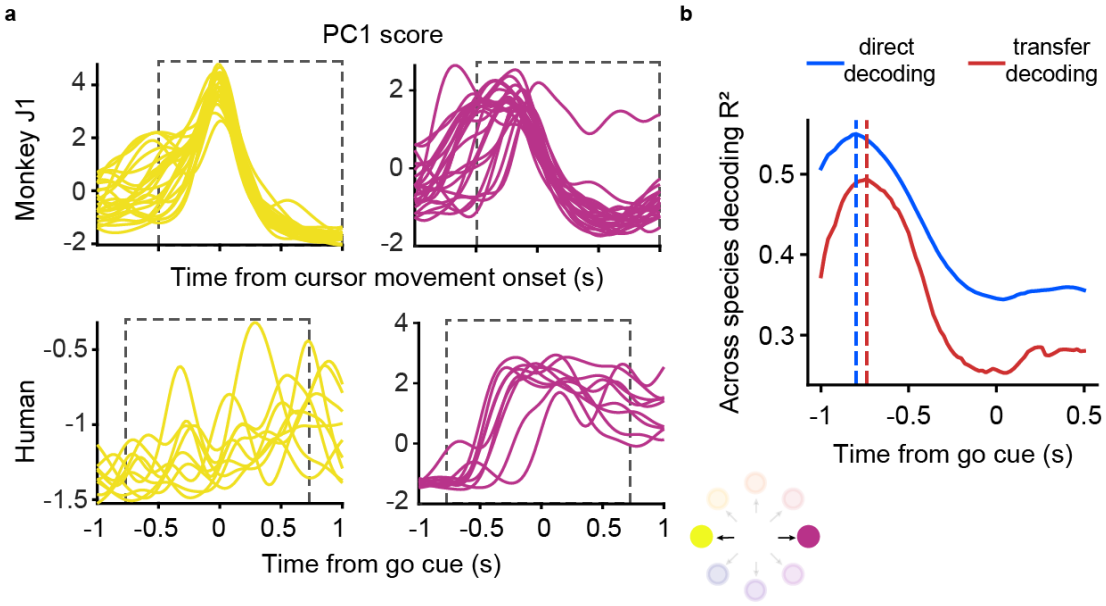

Figure 4-6: Accuracy of monkey-to-human EMG decoding depends on the latency relative to the go cue used for human trial segmentation. a) Single-trial M1 data from the source monkey (left, first session from monkey J) and the human participant (right) projected on the first principal component (computed in each case using the entire corresponding dataset) for a pair of oppositely directed targets. The neural responses recorded during the human's attempted task had greater trial-by-trial variation in timing and magnitude compared to those of the monkey. b) EMG decoding accuracy with direct (red) and transfer decoding (blue) as a function of the time index relative to the go cue used to segment the human trials. The vertical dashed line indicates the time yielding greatest EMG decoding accuracy.

Given optimal time-alignment, the average R^2 prediction accuracy with direct decoding approach was 0.55 (**Figure 4-7a**, blue lines), 98% as accurate as the corresponding cross-monkey predictions, and 78% as accurate even as the average within-monkey/with session decoding. The transfer decoding approach achieved similar accuracy, with an average R^2 of 0.49 (**Figure 4-7a**, red lines).

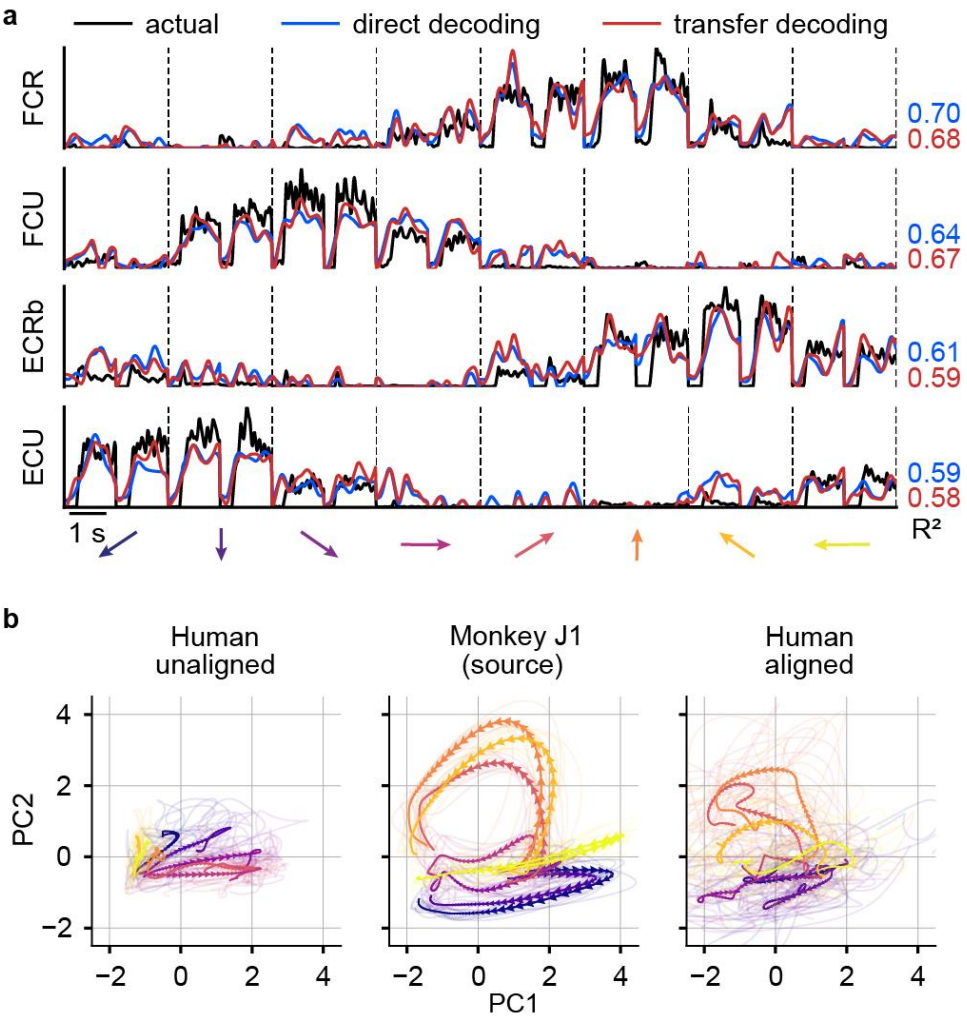

Figure 4-7: EMG decoding from a human with tetraplegia. a) EMG predictions for the four major wrist muscles obtained via direct (blue lines) and transfer decoding (red lines). Actual EMG recordings of the source monkey (black) provide a ground truth. b) Latent M1 trajectories described by the first two principal components of the M1 data (left) recorded as the participant attempted to perform a wrist isometric task, and those for the source monkey (center; first session from monkey J). The human neural recordings had a stereotypical low-dimensional structure for each of the eight target directions, albeit with increased inter-trial variability (shown by the lighter traces). Despite this increased variability, CCA alignment recovered a shape similar to that of the source monkey's latent signals (right).

Despite the greater inter-trial variability, the latent signals of the human data traversed well-separated, stereotypical trajectories for each of the eight targets (**Figure 4-7b**, left panel). CCA alignment allowed us to match the two sets of latent signals, if not quite as closely as for typical pairs of monkeys (**Figure 4-7b**, middle and right panels; R^2 increased from -0.56 to 0.10).

Decoders need to be stabilized in the face of changes in the neurons that are recorded over time

The performance of any fixed iBCI decoder tends to decline over time due to the inherent instabilities of the recorded neural signals (Downey et al. 2018; Perge et al. 2013; Sussillo et al. 2016). One simple solution to combat this effect is to recompute the iBCI decoder at each session as we have done with the direct decoding approach. Alternatively, various groups have proposed the use of neural aligners that match the statistics of neural recordings from a later day ("day-k") to those from the first day when the decoder was calibrated ("day-0") (Degenhart et al. 2020a; Farshchian et al. 2018; Gallego et al. 2020; Karpowicz et al. 2022; Ma et al. 2022). These neural alignment approaches were motivated to avoid the scenario in which an iBCI user would need to relearn the dynamics of each recomputed decoder. In the transfer decoding approach, CCA alignment works implicitly across time as well as across users, thereby automatically addressing this problem. However, this is not the case for direct decoding. Its maintained accuracy was only possible because we trained a new decoder for every session. If instead the direct decoder is fixed, its performance rapidly declines with time, as expected (**Figure 4-8**, blue squares). Therefore, we investigated the effectiveness of using CCA to align the day-k and day-0 latent signals using monkey data (**Figure 4-8**, blue dots and line). In this scenario, the greater direct decoding accuracy over the transfer decoding approach disappears (red and blue circles).

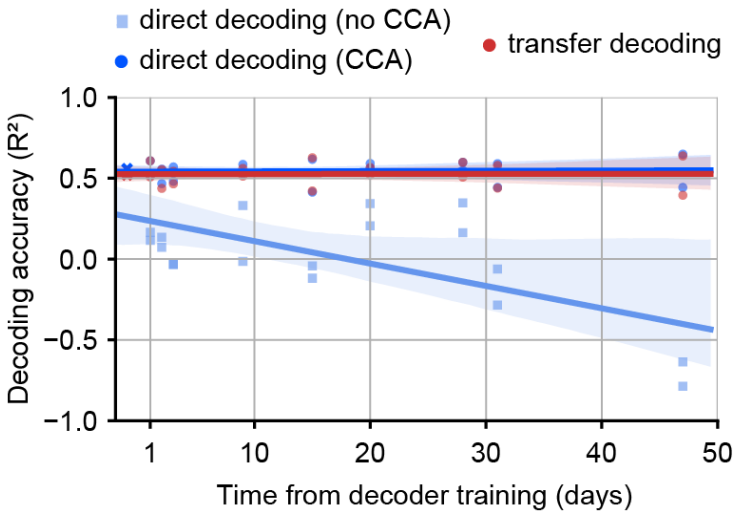

Figure 4-8: A fixed direct decoder needs to be aligned across time. EMG prediction accuracy over time using a fixed 'day-0' direct decoder trained on the first target monkey session. We compared the performance of the fixed direct decoder before (turquoise) and after (blue) within-monkey/across-time CCA alignment between the day-k and day-0 target neural data. CCA alignment stabilized the performance of the fixed direct decoder over time, such that it achieved performance similar to that of the transfer decoding approach (red).

Task generalization of cross-individual EMG decoding

Beyond the question of how well a given decoder performs when trained and tested on similar movements, is the question of how well they extrapolate to different movements. This question may be particularly important for the cross-individual decoders we are developing, as they each depend on the two individuals performing similar movements.

Using exclusively the monkey datasets, we tested the task generalization performance for both cross-individual decoding approaches. We first investigated their performance on test targets that interpolated the four cardinal directions used for trianing (**Figure 4-9a**). We also tested an extrapolation condition in which we trained on the four lower targets and tested on the upper targets (**Figure 4-9b**). For the three conditions (full training set, interpolated training data, extrapolated training data), cross-monkey EMG decoding R^2 was 0.56, 0.42, and -0.45,

respectively, for the direct decoding approach (**Figure 4-9b**, red). In contrast, the cross-monkey EMG decoding R^2 was 0.53, 0.40, and -0.23, respectively, for the transfer decoding approach (**Figure 4-9b**, blue). Direct decoding significantly outperformed transfer decoding when using all training set and in the interpolation case (P ~ 0, Wilcoxon's signed rank test). Neither approach worked well for the extrapolation case as the average R^2 was negative for both.

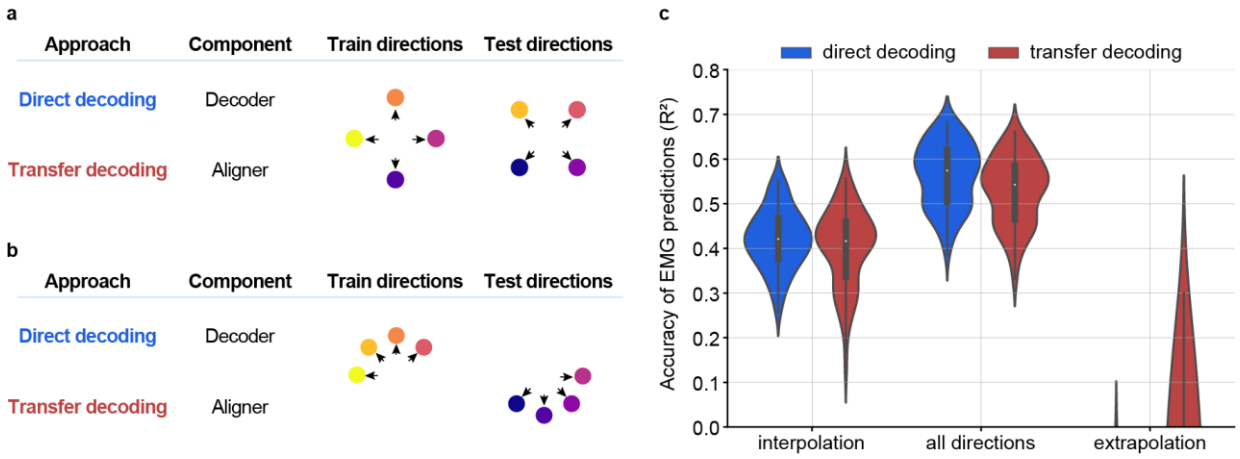

Figure 4-9: Task generalization of cross-monkey EMG decoding. Generalizability of the direct (blue) and transfer (red) decoding was assessed by training either the direct decoder or the cross-monkey CCA alignment using only a subset of all eight movement directions. a) Task generalization when *interpolating* (i.e., training on cardinal directions and testing on diagonal directions); b) Task generalization when *extrapolating* (i.e., training on adjacent lower directions and testing on upper directions). c) Violin plots for the overall cross-monkey decoding accuracy for all pairs of monkeys when testing on all target directions (center), when interpolating (left), and when extrapolating (right). When interpolating, cross-monkey decoding is still possible with both direct and transfer decoding, albeit with lower accuracy. When extrapolating, the cross-monkey decoding generally failed, as indicated by the negative $\boldsymbol{R^2}$ values.

Discussion

When surveyed, nearly 80% of the individuals with high-level spinal cord injury would elect brain surgery to recover some control of their own hands (Blabe et al. 2015). However, the “observation based” decoders currently used to map motor cortical activity to movement kinematics are not directly applicable to the prediction and control of an unobservable motor output like muscle activity. The focus of this study is to test two possible solutions to this problem. The first is the "direct decoding" approach where we mapped the neural signals recorded from an iBCI user to muscle activity of another individual (or monkey) performing the same movements. In the second approach, which we call “transfer decoding”, we used a decoder trained on both M1 and EMG signals from a “source” monkey and then transferred it to a human iBCI user. The transfer decoding approach relies on aligning the low-dimensional latent neural signals from that user to those of the source monkey for which the decoder was computed. We investigated the performance of both approaches both between pairs of monkeys and from a monkey to a human with quadriplegia.

The two approaches yielded similar EMG prediction accuracy, although direct decoding was on average 5% more accurate, provided the decoder was retrained at each session. However, when we held the direct decoder fixed and fed it with latent signal inputs that were aligned between sessions with CCA, there was no difference between the two approaches. Finally, we showed that while both approaches were able to generalize to a set of interpolated targets, neither could extrapolate to test targets in a completely different part of the workspace.

Overall, our work indicates that cross-user EMG decoding is possible even between monkeys and humans. This would allow the use of predicted EMG as a control signal for functional electrical stimulation of muscles as a means to restore voluntary arm movement or movement of an anthropomorphic limb actuated with control properties designed to mimic those of the musculoskeletal system.

Clinical applications of a monkey-to-human biomimetic decoder

Existing iBCIs have allowed paralyzed individuals to gain control of a computer cursor with only a few minutes of practice (Brandman et al. 2018; Hochberg et al. 2006) and three dimensions of a robot end-effector (Hochberg et al. 2012), even within the first few hours of training (Collinger et al. 2013). Higher-dimensional control is more difficult for the user to learn, requiring weeks and even months of practice to achieve control of 7 (Collinger et al. 2013) or 10 degrees of freedom (Wodlinger et al. 2014). However, these kinematic iBCIs allow no direct control of applied grasp forces. Attempts to combine position and joint torque control for 2D planar reaching have met with limited success (Chhatbar and Francis 2013; Fagg et al. 2009). These limitations are present despite the well documented evidence that force and muscle-like information is encoded in neurons in the primary motor cortex (Cheney and Fetz 1980; Evarts 1968; Hepp-Reymond et al. 1994; Holdefer and Miller 2002; Kalaska and Hyde 1985; Lemon, Johansson, and Westling 1995; Maier et al. 1993; Morrow, Jordan, and Miller 2007; Oby, Ethier, and Miller 2013; Sergio and Kalaska 2003).

The mammalian neuromuscular system controls the motion of the arm and digits, the stiffness of joints, and exerted forces, all through the modulation of muscle activity. One approach to a more biomimetic form of robotic limb control allowing more intuitive control of many degrees of freedom might be to feed real-time predictions of muscle activity to a musculoskeletal model that would compute muscle forces and the evolution of limb state; these signals could be used to control the motion and impedance of an anthropomorphic prosthetic arm (Blana et al. 2017, 2020; McFarland et al. 2023).

Alternatively, the predicted EMG signals could be used to control muscle force directly through the electrical stimulation of the muscles or peripheral nerves. This technique, known as Functional Electrical Stimulation (FES), is used to improve stance and walking (Daly et al. 2011; Granat et al. 1993; Thrasher, Flett, and Popovic 2006) as well as grasping (Peckham et al.

2001; Peckham, Marsolais, and Mortimer 1980; M. B. Popovic et al. 2002; Snoek et al. 2000) following stroke or spinal cord injury. Patients with C5-C6 spinal cord injury (SCI) were able to use their preserved voluntary shoulder movements to trigger preprogrammed stimuli and regain some ability to grasp objects (Taylor, Esnouf, and Hobby 2002). This approach becomes much more limited in patients with higher level cervical SCI, as those needing the greatest restored movement have the least available peripheral control signals. iBCIs offer an alternative solution by providing the means to obtain more natural, higher-dimensional control signals than those derived from residual movements.

Our group previously designed an iBCI-controlled FES system that enabled monkeys with temporary paralysis of the hand muscles induced by a peripheral nerve block to perform hand movements: decoded EMGs were used to modulate stimulation of five electrodes implanted in different compartments of three hand flexor muscles (Ethier et al. 2012).

Two other groups have developed brain-controlled FES systems in which human participants with cervical spinal cord injury were able to control simple elbow, wrist, or hand movements. One approach used six parallel decoders, each predicting one of six wrist or finger movements. The decoder with the highest output triggered a muscle stimulation pattern consisting of at most, three intensity levels, designed to approximate the decoded movement (Bouton et al. 2016). In another approach, real-time velocities of the elbow, wrist, hand, and shoulder were decoded from M1, and controlled in a feedback system using a combination of muscle and nerve stimulation (Ajiboye et al. 2017).

While promising as additional proofs of concept, these approaches achieved only limited control of a small number of dimensions. The former, in particular, seems unlikely to be able to scale to more complex movements. In our study, we aimed to achieve the means to provide more intuitive control of more degrees of freedom, including nonkinematic aspects of movement, by directly inferring muscle activity from M1. We showed two possible approaches to map motor

cortical activity of a patient with paralysis into intramuscular EMG signals that resembled those of a monkey performing the same isometric force-generation task.

The redundancy of the hand musculature allows most actions to be produced with a variety of different patterns of muscle activity (Bernstein 1966; d'Avella, Saltiel, and Bizzi 2003; Santello and Soechting 2000). In our monkeys, EMGs were consistent across time for a given monkey (**Supplementary Figure 4-2**), but less so across monkeys (**Supplementary Figure 4-3**). As the muscle activity patterns became more dissimilar across monkeys, so did their latent neural representations (**Supplementary Figure 4-5d**). Even more so, the human participant may have intended to use a somewhat different pattern of EMG from that of the source monkey. However, to the extent that the biomechanics of the human hand resembles that of the monkey, the resultant EMG predictions should be biomechanically appropriate to accomplish the task as the source monkey. The discrepancy would impose an interesting motor-adaptation problem as the user learns to interact with the decoder.

Neural representations of motor intent are similar across monkeys and humans

The neural population activity in many brain areas is constrained to a low-dimensional neural manifold. Furthermore, there is increasing evidence that the dynamics of specific patterns of activity within the manifold, the latent signals, underlie the computations required for planning and executing movements (Cunningham and Yu 2014; Elsayed and Cunningham 2017; Gallego et al. 2017; P. Gao and Ganguli 2015; Mante et al. 2013; Mazor and Laurent 2005; Williamson et al. 2019). In the primary motor cortex, recorded latent signals that differ across days can be mathematically transformed ("aligned") to be more similar to each other. These aligned latent signals maintain a remarkably stable relation to behavior over months and even years (Gallego et al. 2020). As a consequence, a fixed decoder that uses these aligned signals as inputs remains accurate across long periods without supervised recalibration (Degenhart et al. 2020a; Gallego et al. 2020; Karpowicz et al. 2022; Ma et al. 2022).

Manifolds transcend the analysis of individual neurons, and make the comparisons of population activity across individuals more readily interpretable (Dabagia, Kording, and Dyer 2022). Recently, Safaie et al. used similar alignment techniques to show that latent signals in M1 were preserved across monkeys as they performed a center-out reaching task using a planar manipulandum; latent signals were also preserved in the dorsolateral striatum of mice that grasped and pulled a joystick (Safaie et al. 2022). Other recent studies revealed that low-dimensional neural structure within the hippocampus and the sensorimotor cortex is preserved across rats for a variety of behaviors, ranging from locomotion along a linear track inside a maze to unconstrained movement in an arena (H.-T. Chen, Manning, and van der Meer 2021; Melbaum et al. 2022; Nieh et al. 2021; Rubin et al. 2019).

Not only was the neural representation of the isometric wrist task similar across monkeys, but also the similarity extended even to a paralyzed human attempting to perform the same task. The preserved motor intent signals between individuals enabled successful monkey-to-human EMG decoding, even though the decoding accuracy was not as high as when using M1 data from a target monkey. The lower cross-user decoding accuracy from the human M1 data could be attributed to the greater inter-trial variability of the corresponding latent trajectories compared to those of the monkeys (**Figure 4-6a**). The increased trial-to-trial variability partly resulted from the absence of any actual force or force-related feedback. Moreover, the variability was influenced by the participant's tendency to anticipate the go cue during the attempted movements (**Figure 4-6a**). As a result, the accuracy of both direct decoding and transfer decoding peaked for time-alignment *prior to* the go cue (**Figure 4-6b**). Remarkably, under these limitations, the EMG predictions from either cross-user decoding approach with human M1 data still resembled the actual EMGs of the source monkey.

Comparison of direct and transfer decoding

Direct decoding was simple and consistently outperformed transfer decoding by a small amount, as long as the decoder was retrained each session. Unlike transfer decoding, which requires that both M1 and EMG data be collected from the source monkey, direct decoding has the advantage that it would be possible to use EMGs collected from able-bodied humans instead of monkeys, which would reduce any concern about the similarity of monkey and human biomechanics.

The task we studied here was highly stereotypic, such that individual trials could be aligned in time, a requirement for both direct decoding and the CCA alignment used in transfer decoding. This limitation is fundamental for direct decoding; however, CCA could, in principle, be replaced with an unsupervised method that aligns the statistics of clouds of points independently of the time course of the associated signals. Such approaches have been used to align neural signals from a given monkey across time, (Degenhart et al. 2020a; Farshchian et al. 2018; Karpowicz et al. 2022; Ma et al. 2022) but they have not been successfully applied across monkeys. These methods should be further investigated as a means to allow transfer decoding to be applied to a broader range of unstructured tasks.

In addition, because it separates the decoder computation and alignment phases, transfer decoding could take advantage of large quantities of more readily available data from monkeys. With more extensive data, it may be possible to train sophisticated decoders that can operate across various behaviors, while relying on the more limited data from the human user for neural alignment. Although their effectiveness needs to be validated in future studies, deep learning-based decoders could be used for this purpose, as they have been shown to be more robust to variability in recording conditions and have higher decoding accuracy compared to linear alternatives (Glaser et al. 2020; Sussillo et al. 2016; Willsey et al. 2022).

Summary

To our knowledge, this study is the first that demonstrates the feasibility of cross-user EMG decoding, both a “direct” method and one that relies on “transfer” a decoder computed from a monkey to a human user. We validated both approaches on multiple pairs of monkeys, and then tested them between a human with paralysis and a monkey. Future studies should focus on developing more sophisticated decoders and neural aligners to further expand the applicability of these approaches to varied and unstructured movements typical of daily living. While their feasibility would have to be tested online to address the problems associated with motor adaptation, the current study serves as an initial and encouraging proof of concept.

Acknowledgments

We thank Eric J. Perreault for valuable discussions. We also thank current and former members of the Miller Limb Lab, including Stephanie Naufel, Matthew Perich, and Christian Ethier, for their contributions to data collection. The work was supported in part by grants to L.E.M. (R01 NS053603, R01 NS074044).

Supplementary Figures

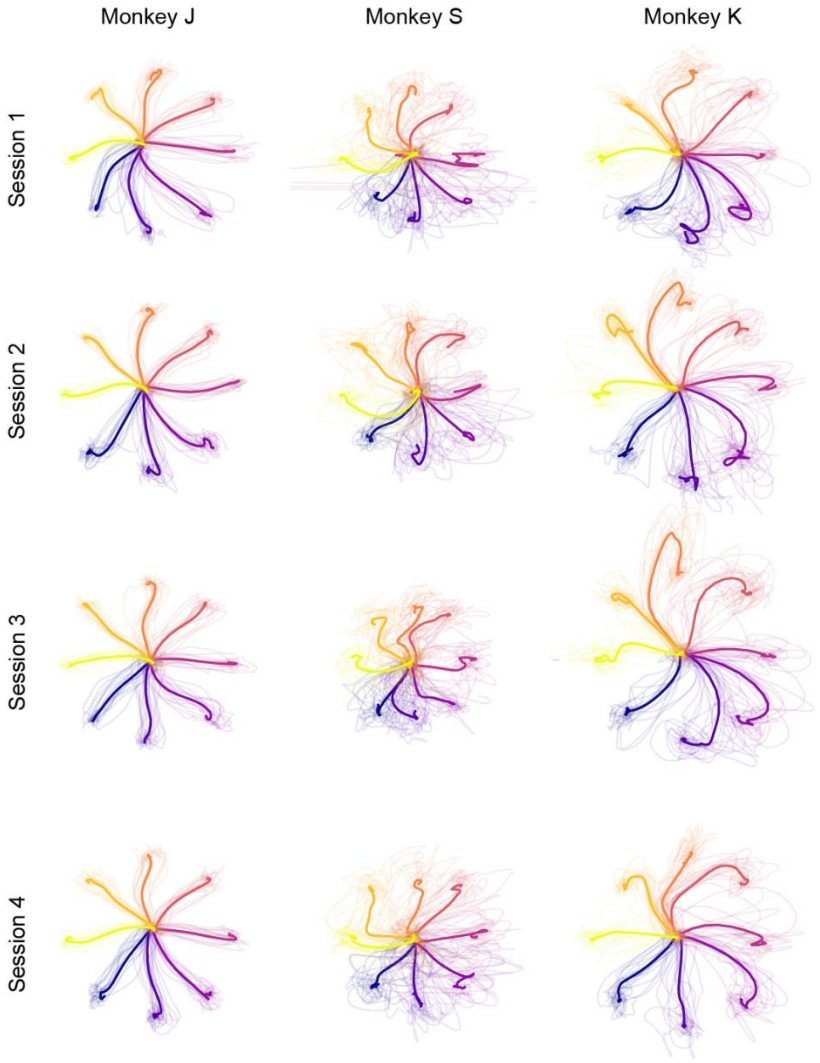

Supplementary Figure 4-1: Monkeys have slightly different cursor trajectories. Cursor trajectories of the three monkeys during the four sessions analyzed in this study. Data is averaged across all trials for each target direction (single trial trajectories are shown as shaded curves). Among the three monkeys, Monkey J has the straightest trajectories. Monkey K and (especially) S show the highest trial-by-trial variations.

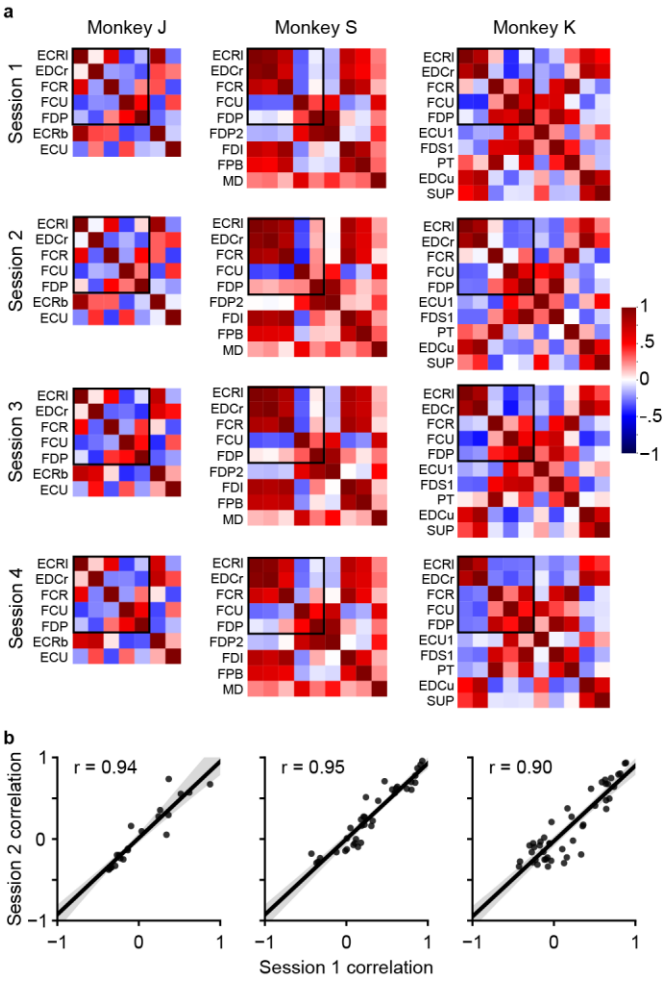

Supplementary Figure 4-2: Muscle correlation patterns are preserved across time. a) Matrices of EMG correlations in each monkey during the four recorded sessions. Each matrix is ordered to cluster muscles that have been recorded in common across all three monkeys (indicated by the black square). b) Element-by-element scatter plots of the matrices in the top (x-axis) and bottom (y-axis) row of a. Correlation of EMGs in the same monkey are very consistent across time.

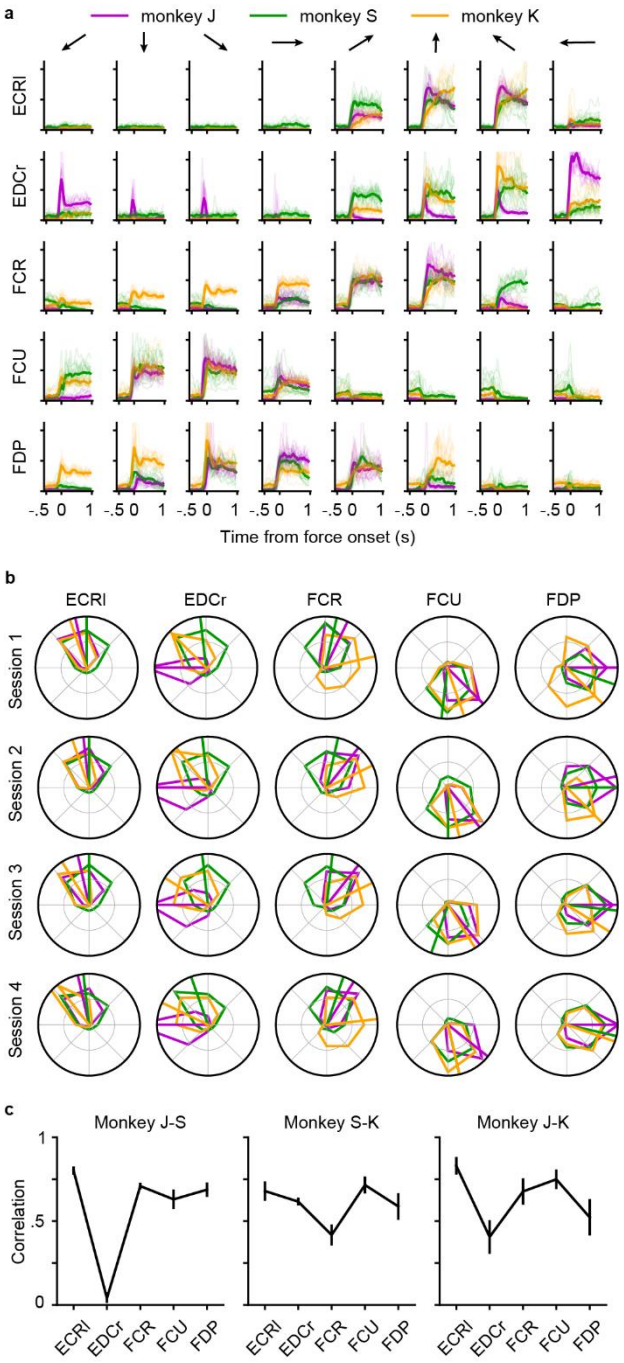

Supplementary Figure 4-3: Monkeys perform the same task using slightly different muscle strategies. a) EMG traces recorded during the isometric box task for the common muscles of the first session of monkey J (purple), monkey S (green), and monkey K (orange). Columns show the average EMG activity for the eight target directions (shaded lines are single trials). The EMG traces are centered around force onset. b) Muscle tuning curves expressing the level of activity as a

function of the direction of the isometric force. The radial vectors indicate the preferred directions. Each row shows the tuning curves and PDs of monkey J (purple), monkey S (green), and monkey K (orange) recorded during the four sessions. c) Correlation of EMGs that have been recorded in common between Monkey J and S (left), Monkey S and K (center) and Monkey J and K (right). Mean and standard deviation across all sessions is reported.

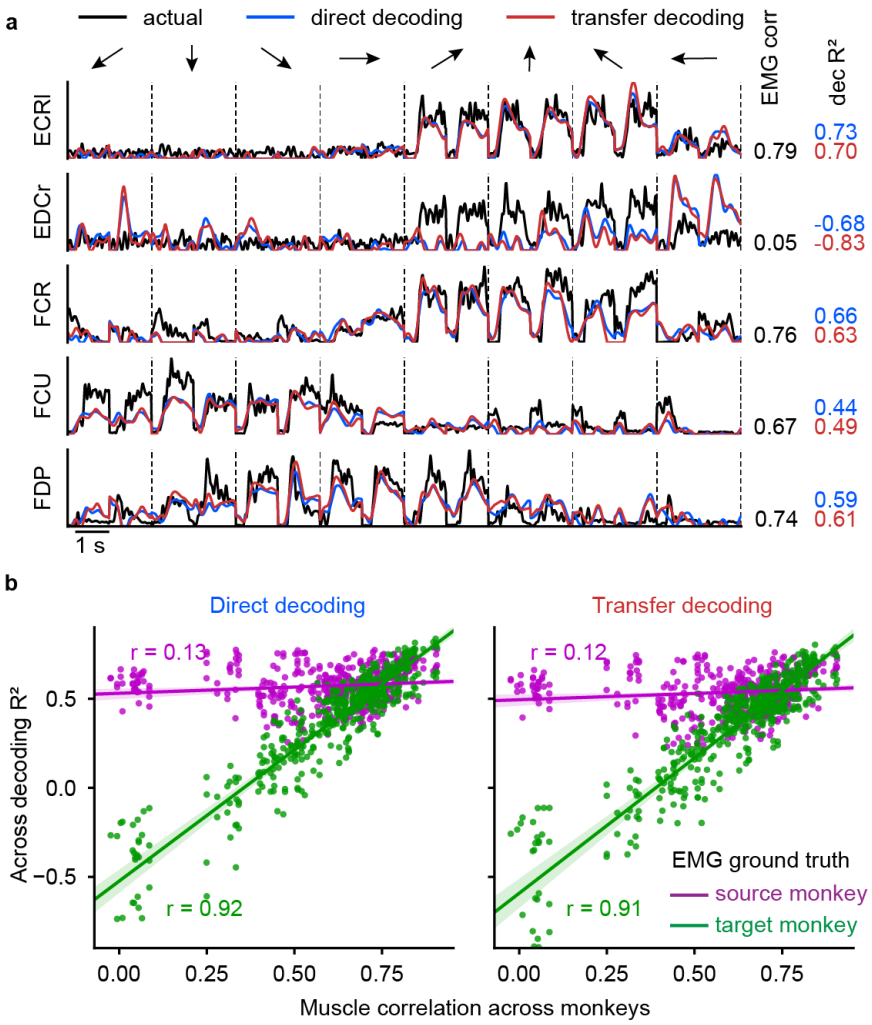

Supplementary Figure 4-4: Using target monkey EMGs as ground truth for cross-monkey decoding. a) Cross-monkey EMG predictions obtained with target-monkey latent signals via direct decoder (blue lines) and with aligned target-monkey latent signals via transfer decoding (red lines). Actual EMG recordings of the target monkey (black lines) provide a ground truth for measuring decoding accuracy. Note that we cannot compare the cross-monkey predictions with the entire set of recorded muscles of the target monkey, but only with the set of EMGs that have been recorded in common across monkeys, as the fixed decoders have been trained to predict the actual EMGs of the source monkey. For each muscle, the cross-monkey predicted R^2 is reported with the correlation between the actual EMG recordings of the source and the target monkey. b) Muscle correlation between source and target monkeys is plotted against the single muscle cross-monkey

R^2 with direct (left) and transfer (right) decoding when using source (purple) and target (green) monkey EMGs as ground truth. With both methods, we can decode target monkey EMGs that have high correlation with those of the source monkey. Each dot refers to a single muscle for a given source/target monkey pair.

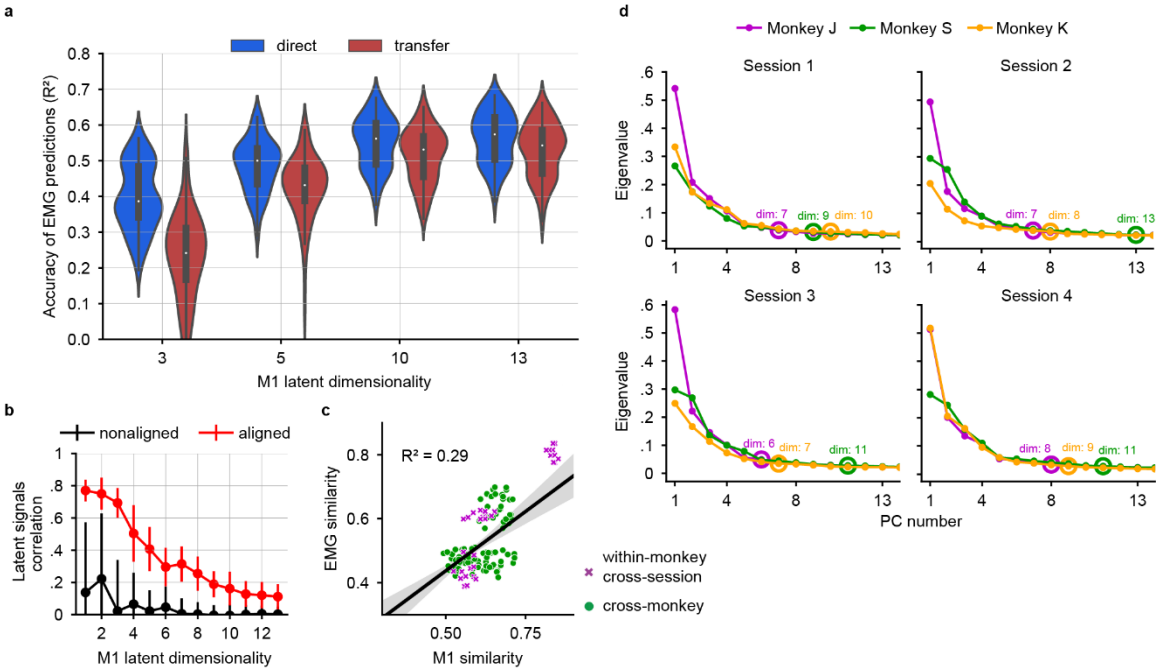

Supplementary Figure 4-5: Dimensionality analysis. a) Violin plot showing the overall cross-monkey EMG decoding R^2 with direct (red) and transfer (blue) decoding for all pairs of monkeys when using different latent space dimensionality. b) Correlation between the latent signals of source and target monkeys on each of the first 13 latent dimensions before (black) and after (red) CCA alignment. The increased correlation magnitude indicates increased similarity across monkeys as a result of the CCA alignment. The mean and standard deviation across all pairs of monkeys is shown. c) M1 similarity is plotted against the EMG similarity across monkeys (green dots) and within monkey/across sessions (purple x symbols). The similarity is computed as the correlation between the CCA-aligned M1/EMG latent signals averaged across the first five latent dimensions. For the EMG similarity, we applied PCA on the EMG signals that were recorded in common across monkeys (see Table S1) and then compute CCA alignment on the corresponding latent signals. d) Scree plot of the first 13 principal components for the four sessions of monkey J (purple), monkey S (green) and monkey K (orange). The Parallel Analysis estimate of linear dimensionality for each monkey's latent signal is indicated as a circle of the corresponding color.

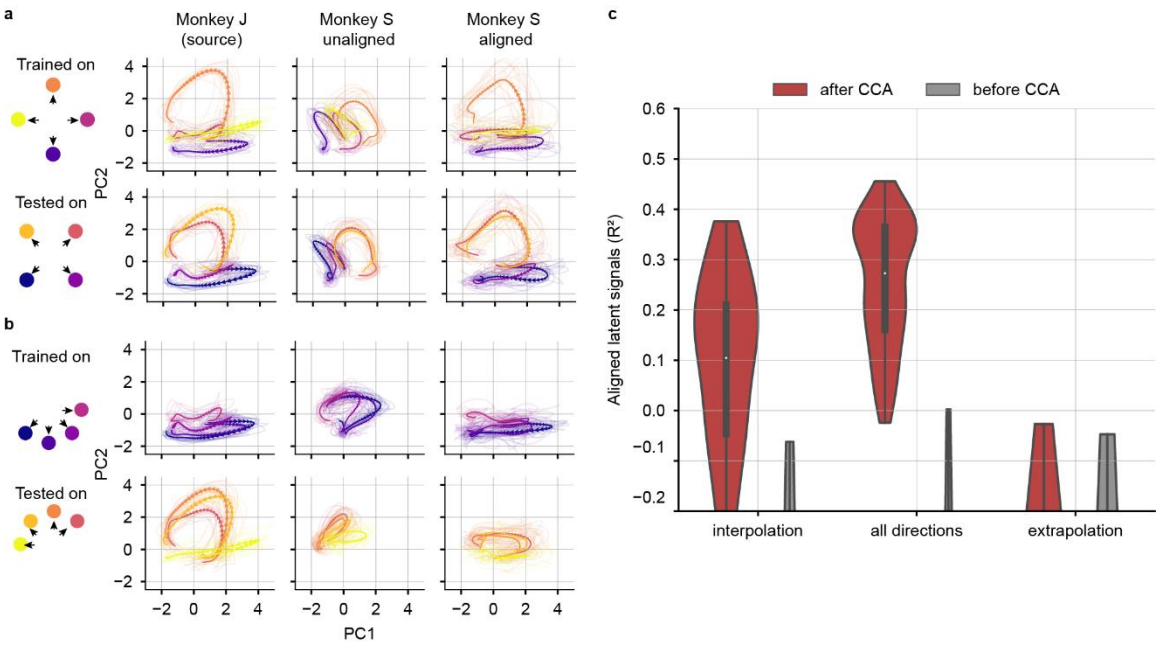

Supplementary Figure 4-6: CCA alignment generalizes when interpolating, but not when extrapolating. a-b, Representative latent signals described by the first two principal components of a single monkey pair (source monkey: J, first session; target monkey: S, second session) before and after CCA alignment. a) Generalizability of the CCA alignment when interpolating (i.e., trained on cardinal directions and tested on diagonal directions); b) Generalizability of the CCA alignment when extrapolating (i.e., trained on adjacent lower directions and tested on upper directions). When interpolating, the aligned neural trajectories of the target monkey resembled those of the source monkey when tested on the held-out directions (a, bottom row). When extrapolating, CCA accurately aligned the training data (b, top row), but not the trajectories of the held-out directions (b, bottom row). c) Violin plot showing the similarity, as measured by R^2, between the latent signals of all pairs of source and target monkeys before (grey) and after (red) CCA alignment when data from all (center), only cardinal (left), and only lower (right) directions were available. When using all directions and when interpolating, the similarity significantly increased after CCA alignment. When extrapolating, the R^2 values are still negative even after alignment.

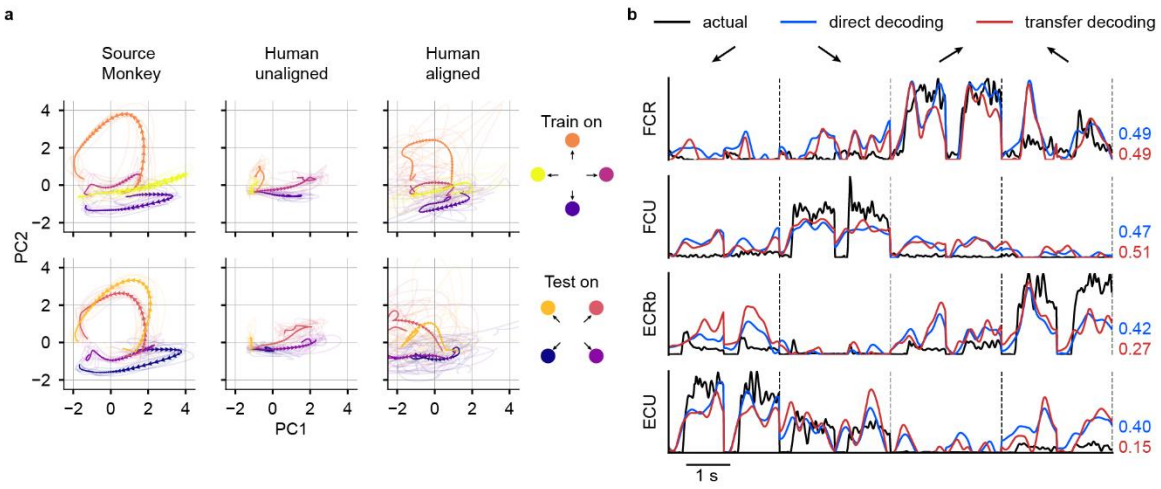

Supplementary Figure 4-7: Cross-species decoding generalizes when interpolating. a) Latent M1 trajectories described by the first two principal components of the source monkey (left) and the paralyzed patient attempting to perform a wrist isometric task before (center) and after (right) CCA alignment. The first row shows the latent signals of the directions used to obtain the CCA transform, while the bottom row shows the latent signals of the held-out directions. In both cases, CCA made the neural trajectories of the human more similar to those of the source monkey. b) Corresponding cross-monkey EMG predictions obtained with human latent signals via direct decoder (blue lines) and with aligned human latent signals via transfer decoding (red lines). Actual EMG recordings of the source monkey (black lines) provide a ground truth for measuring decoding accuracy. The $\boldsymbol{R^2}$ for both direct (red) and transfer (blue) decoding are shown on the right for each muscle. Vertical dashed lines separate muscle traces for the eight target directions.

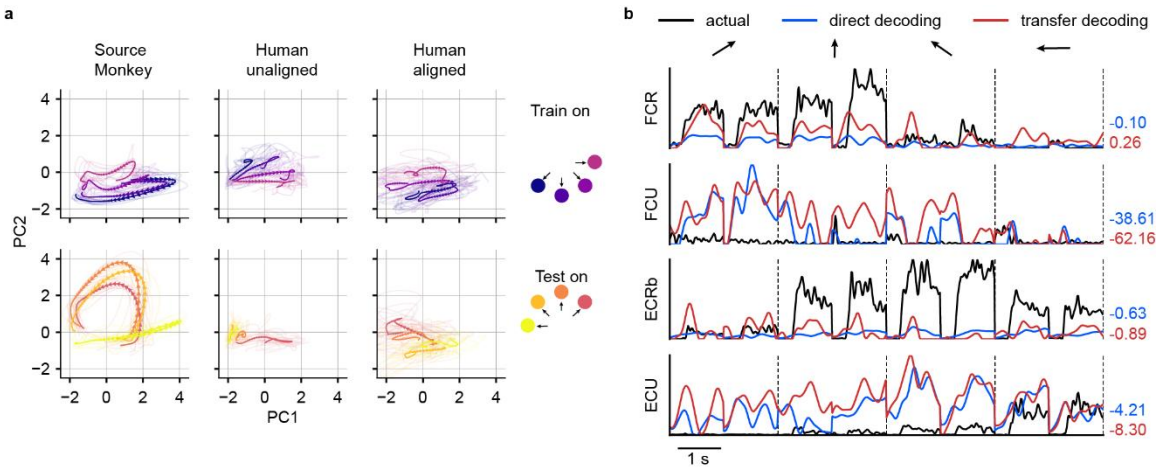

Supplementary Figure 4-8: Cross-species decoding does not generalize when extrapolating.

a) Latent M1 trajectories described by the first two principal components of the source monkey (left) and the paralyzed patient attempting to perform a wrist isometric task before (center) and after (right) CCA alignment. The first row shows the latent signals of the directions used to obtain the CCA transform, while the bottom row shows the latent signals of the held-out directions. While CCA accurately aligned the latent signals of the training directions, the latent signals of the human bear no resemblance to those of the source monkey in the held-out targets. b) Corresponding cross-monkey EMG predictions obtained with human latent signals via direct decoder (blue lines) and with aligned human latent signals via transfer decoding (red lines). Actual EMG recordings of the source monkey (black lines) provide a ground truth for measuring decoding accuracy. The $\boldsymbol{R^2}$ for both direct (red) and transfer (blue) decoding are shown on the right for each muscle. Vertical dashed lines separate muscle traces for the eight target directions.

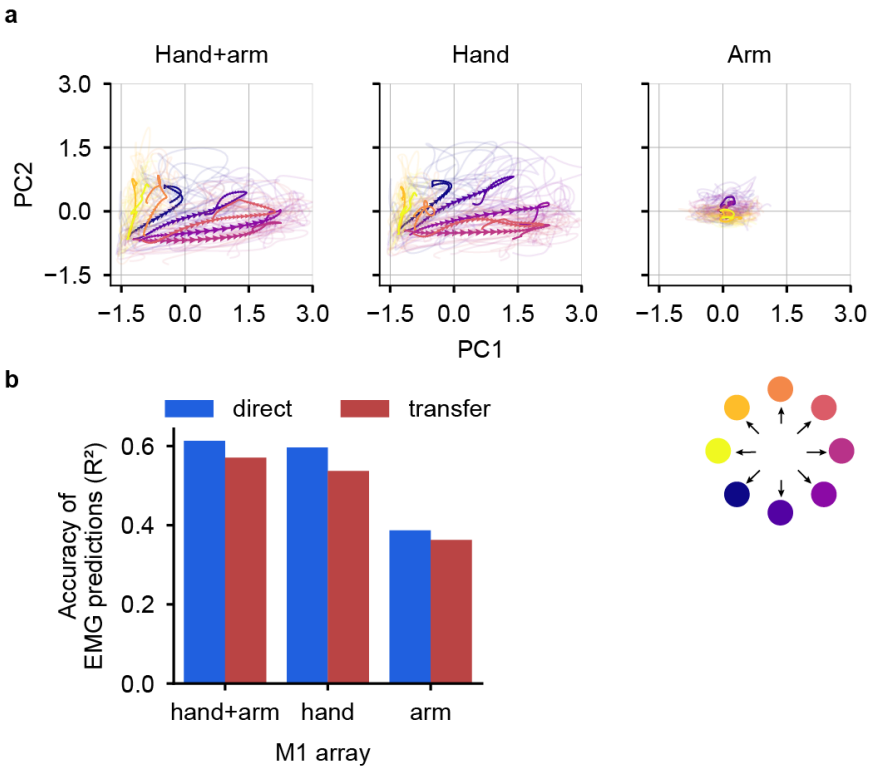

Supplementary Figure 4-9: Additional cross-species decoding analysis. a) Latent signals described by the first two principal components of the human recording using channels from both hand and arm (left), only hand (center) and only arm (right) area of M1. Data is averaged across all trials for each target direction (single trial trajectories are shown as shaded curves). Arrows indicate the temporal evolution of the trajectories. b) Cross-species R^2 for direct (blue) and transfer (red) decoding when using the human latent signals from the hand, arm, and the combined areas of the M1 array. For the cross-species decoding, we only considered the four main wrist muscles.

CHAPTER 5: DISCUSSION

Summary of findings

My first objective was to refine the methodology for accurately estimating the geometry, namely the dimensionality and nonlinearity, of neural manifolds. The second objective was to understand whether the observed low dimensionality of M1 manifolds is a general computational strategy or a trivial result of constraints associated with motor behaviors performed in the laboratory. The final objective, again using the concept of the neural manifold, was to develop innovative strategies that address the challenges faced by observation-based iBCI decoders, potentially enabling individuals with paralysis to control their muscles using their neural activity. I addressed these three main objectives in Chapters 2, 3, and 4, respectively.

In Chapter 2, I laid the groundwork for the subsequent chapters by evaluating techniques for estimating the intrinsic and embedding dimensionality of neural recordings and devised a pipeline to do so in a principled manner. Through the analysis of linear and nonlinear algorithms, I found that none of the tested algorithms work for all possible scenarios. However, I identified important insights regarding the conditions under which estimates of intrinsic dimensionality are likely to be valid or not. These conditions mostly depended on the extent of nonlinearity in the underlying manifold's geometry, the amount of available data, and the levels of noise. In light of these findings, I devised an analysis pipeline that will aid in computing the neural manifolds and in quantifying the complexity of information encoded by neurons.

Chapter 3 extended this investigation by exploring the intrinsic and embedding dimensionalities of M1 manifolds in monkeys engaging in unconstrained tasks. I discovered that both the intrinsic and embedding dimensionality of M1 manifolds were slightly higher in unconstrained, natural motor behaviors compared to those in the constrained laboratory tasks that I analyzed and that had been previously reported. Although signatures of nonlinearity were present in the M1

manifolds, most of the neural activity was restricted to nearly linear regions. Furthermore, the accuracy of linearly decoded electromyograms (EMGs) from low-dimensional latent variables closely matched the accuracy of EMGs decoded from all recorded neurons, implying that most of the information related to the activation of muscles in M1 lived in a linear subspace within the slightly nonlinear neural manifold.

Finally, in Chapter 4, I focused on predicting and controlling motor outputs that cannot be observed, such as the activity of muscles that produce movement, from M1 activity for intracortical brain-computer interfaces (iBCIs). To achieve this, I utilized the latent signals on the low-dimensional neural manifolds computed according to methods described in Chapters 2 and 3 to perform cross-user EMG decoding between a source and a target user. The source user I tested was always a monkey, from which the neural and EMG data were relatively easy to collect. The target user's neural data came from either another monkey or a human with paralysis attempting to perform the same movements as the source monkey. I compared two approaches. The first approach, "direct decoding," maps neural signals from the target user to the muscle activity of the source individual performing the same movements. Alternatively, "transfer decoding" transfers a decoder trained on M1 and EMG signals from a source monkey to a target monkey or human iBCI user. Both approaches showed similar EMG prediction accuracy, and cross-user EMG decoding was possible even between monkeys and humans. This finding suggests potential applications in restoring voluntary arm movement or controlling anthropomorphic limbs with properties designed to mimic the musculoskeletal system.

Importance of nomenclature on neural dimensionality

The concept of dimensionality of neural signals has emerged as a fundamental tool in understanding computational processing carried out by a population of neurons. This interest is driven by the ability of dimensionality to quantify the redundancy and to pinpoint shared

components of collective dynamics of interconnected networks of neurons with respect to task variables. However, the term “dimensionality” has different meanings for different groups. An important challenge has been the lack of a unified approach and shared terminology surrounding the term despite its perceived importance in the context of analyzing the activity of neural populations.

Different studies investigating dimensionality often leaned on their individual interpretations and definitions of the concept, leading to a fragmented understanding. The lack of a common framework, or a universal language, can result in confounding interpretations and assessments. To illustrate this, let us consider recordings from the primary visual cortex when a mouse is presented with a series of images. While the interaction of a substantially large number of neurons drives how V1 may process these images, the relevant variables for the task are likely to be bounded by the statistics of the stimuli (Simoncelli and Olshausen 2001).

One group has recently found that natural visual stimuli evoke high-dimensional geometry in V1 manifolds from mice (Stringer, Pachitariu, Steinmetz, Carandini, et al. 2019). They interpret this finding to be consistent with the “efficient coding” hypothesis (Atick and Redlich 1990; Simoncelli and Olshausen 2001), which means that neural code is expressive, as it optimizes the transfer of information by reducing correlations present in natural images. According to this hypothesis, such neural codes facilitate the read out of the complex features by downstream networks.

Is V1 truly high-dimensional as they claim, or are the representations, in fact, lower-dimensional but highly nonlinear? Without a proper definition of what we mean by dimensionality or accurate methods to measure it, the interpretations of neural redundancy and the underlying computational principles could be varied. This ambiguity could hinder our understanding of how population activity carries out and represents task-relevant variables. Moreover, without common terminology, we may overlook parallels across investigations that span multiple brain

areas, potentially impeding our overall progress in understanding how neural populations process and relay information relevant to behaviors between one another.

Fortunately, the pressing need for a common terminology in studying neural dimensionality has been formalized by Gao and Ganguli (P. Gao and Ganguli 2015; P. Gao et al. 2017) and was finally addressed by Jazayeri and Ostojic just over a year and a half ago (Jazayeri and Ostojic 2021). In the context of neuroscience, they divide the notion of dimensionality into three distinct terms. The first, and the least important of the three regarding neural computations, is *ambient* dimensionality, which describes the total number of recorded neurons. For example, in the multielectrode array data that I analyzed in Chapters 2-4, ambient dimensionality would equal the roughly 100 neurons that were sampled in M1. The remaining two are *intrinsic* and *embedding* dimensionality, which help characterize the geometry of neural manifolds and assist us in interpreting the computations that occur. Intrinsic dimensionality offers a mirror into the information that is encoded, like how the visual stimuli presented to the mouse in the earlier example are encoded by V1. On the other hand, embedding dimensionality is related to how neural circuits process and manage this information.

The work I did in Chapter 2 preceded these definitions and the associated theory regarding neural computation. However, the tools I developed and sharpened using simulated scenarios with known ground truth allowed for testing the theory with experimental data in Chapter 3. For example, Parallel Analysis, the most accurate linear dimensionality estimator I tested, serves as an accurate embedding dimensionality estimator. Similarly, Two Nearest Neighbors and Levina-Bickel Maximum Likelihood Estimation are two accurate methods to estimate the intrinsic dimensionality. These methods would work well to test theories regarding dimensionality and neural computations as long as the data were properly denoised.

In Chapter 3, I applied these tools and definitions directly to M1. The primary aim here was to uncover the underlying computational principles of M1 during unconstrained, natural tasks. By

computing both the intrinsic and embedding dimensionalities of M1 recordings during these tasks, I was able to explore how M1 represents and manages task-relevant variables in different contexts.

My findings revealed that the intrinsic and embedding dimensionalities of M1 signals were slightly higher during complex, unconstrained tasks than during simple tasks performed with a single arm. This suggests a small increase in computational complexity when the motor cortex is driving more freely executed tasks. Yet, both embedding and intrinsic dimensionalities remained much lower than the millions of neurons active in M1. Furthermore, I observed that neural activity that was predictive of EMGs was mostly confined to nearly linear regions within the neural manifolds. This insight is instrumental for future studies that aim to investigate the dynamics of the signals within M1 manifolds and their relation to behavior in different contexts.

Linking the low intrinsic dimensionality, low embedding dimensionality, and mild nonlinearity of M1 to functional organization and recurrent connectivity

What conclusions can we draw from the low dimensionality and modest nonlinearity of M1 regarding its function and connectivity? A potential explanation, as detailed in Chapter 3, is anchored in the computational balancing act between generalizability and expressivity, a concept that has been examined in both artificial and biological neural networks (Musslick et al. 2017; Musslick and Cohen 2019; Pryluk et al. 2019; Sagiv et al. 2020). For example, if task-specific information encoded in a brain area needs to be selectively transmitted to various other brain areas, it would be beneficial to confine the components of this information to distinct, linearly independent subspaces. This form of arrangement would support readouts without interference, exhibiting increased expressivity. Such an organizational structure would necessitate a high embedding dimensionality. Artificial neural network studies have demonstrated the merits of high embedding dimensionality, which has been applied in the

context of kernel methods and support vector machines (Boser, Guyon, and Vapnik 1992; Maass 2016; Cohen et al. 2020).

My findings in Chapter 3 showed that we do not observe this scenario in M1. The low estimates of intrinsic and embedding dimensionality in both constrained and unconstrained tasks suggest that this expressivity is less pronounced in M1 than in some higher-order brain areas. For instance, a brain area required to selectively convey latent variables crucial to decision-making and motor output, such as the dorsolateral prefrontal cortex, would demand high mixed selectivity. Consistent with the theory of dimensionality, these areas have been reported to have a large embedding dimensionality even in simple decision-making tasks (Mante et al. 2013; Rigotti et al. 2013; Fusi, Miller, and Rigotti 2016). The low dimensionality and mild nonlinearity of M1 suggest that such intricate representations are not necessary for M1, even when the laboratory constraints are removed. Instead, M1 presents more generalized representations, which may encode different inputs into a limited set of common activity patterns or “neural modes” (Gallego et al. 2017; Bhandari, Gagne, and Badre 2018). From a functional standpoint, M1’s low-dimensional and generalizable computational strategy could ensure the reliable generation of movement commands that remain largely unaffected by task constraints and contexts.

Another interpretation of the low dimensionality and minor nonlinearity in M1 neural population activity could be related to the strength of recurrent connectivity within M1. The dynamical systems perspective (Shenoy, Sahani, and Churchland 2013) emphasizes the existence of recurrent connections, yet our knowledge of the strength of these connections is limited (Vyas et al. 2020). Although still in its infancy, recent work on artificial and biological networks has tied the strength of recurrent connections to the dimensionality of neural representations (Litwin-Kumar et al. 2017; Mastrogiuseppe and Ostojic 2018; Schuessler, Dubreuil, and Mastrogiuseppe 2020; Beiran et al. 2021; Pollock and Jazayeri 2020). According to this view,

networks with low-rank connectivity matrices give rise to low-dimensional dynamics. While the research connecting recurrent connectivity structure to dynamics is still emerging and calls for further investigation, the low embedding and intrinsic dimensionality of M1 are in line with the presence of relatively weak recurrent connections (Langdon, Genkin, and Engel 2023).

The methods developed in Chapter 2 and tested in Chapter 3 present an enriched understanding of the neural computations within the primary motor cortex. Through the development of precise tools for estimating intrinsic and embedding dimensionality and combining these tools with theories on neural computation, we can interpret the computational implications of the geometry of the neural manifolds within neural circuits. By applying these tools to the motor cortex, we begin to uncover the underlying computational principles, which appear to favor generalizability over expressivity, even in unconstrained and natural settings, while hinting at the potential influence of recurrent connectivity strength on the nature of neural representations.

Muscle-related information lives in a low-dimensional, linear subspace within M1

Considerable research has leveraged linear dimensionality reduction methods like PCA to derive the latent signals that define the M1 manifolds (Churchland et al. 2012; Cunningham and Yu 2014; Sadtler et al. 2014; Kaufman et al. 2014; P. Gao and Ganguli 2015; Gallego et al. 2017; Williams et al. 2018; Gallego et al. 2018, 2020). The advantages of linear models, such as computational efficiency and straightforward interpretability, have rendered them an invaluable tool in this context. Importantly, M1 manifolds computed with PCA contained movement-related variables, allowing insights into how M1 generates movement.

The significance of this finding extends beyond scientific interest—it has practical implications for the development of iBCI decoders. These decoders can be constructed using low-

dimensional, linear M1 manifolds, which have been shown to harness temporally stable dynamics even under neural turnover. The temporal stability of M1 manifolds is an important feature in the context of iBCIs, as it may allow for the creation of reliable brain-to-behavior decoders that perform consistently over time (Gallego et al. 2020).

While these linear approaches yielded significant insight into how M1 generates movement and facilitated iBCI advancements in the laboratory, their effectiveness in unconstrained, real-world environments was unknown. In these settings, one possibility is that information about muscle activation might be embedded nonlinearly in M1. If this were the case, the utility of linear methods like PCA could be severely compromised, potentially limiting their practical and real-world applications.

In Chapter 4, I examined the EMG decoding accuracies by computing neural manifolds using the intrinsic dimensionality estimate and decoding EMG signals. I adopted both linear (PCA) and nonlinear (feedforward autoencoder) dimensionality reduction methods to compute latent signals associated with linear and nonlinear manifolds. My findings showed that the performance of EMG decoding from both linear and nonlinear M1 manifolds was comparable. Moreover, the decoding performance from these manifolds was on par with that derived from all available neurons (**Figure 3-6**).

These observations imply that muscle-related information in M1 appears to reside in a low-dimensional linear subspace within M1. This is a critical discovery for two reasons. First, the linearity confirms the validity and robustness of linear approximations to M1 manifolds in both constrained settings as well as in more natural, unconstrained settings. Second, the low dimensionality paves the way for translating iBCIs outside the laboratory even when we are vastly undersampling the many active M1 neurons.

Cross-user decoding of muscle activity using linear M1 manifolds

In Chapter 4, I put into practice the concept of linear approximations to M1 manifolds in the iBCI context, using their low dimensionality and muscle-like information content. Building on the insights from Chapter 3, I determined PCA to be a sufficient tool for approximating M1 manifolds. I incorporated the embedding dimensionality estimate obtained through Parallel Analysis, the efficacy of which was established in Chapter 3, to determine the dimensionality of the linear approximations to the M1 manifolds. The embedding dimensionality estimates varied between 6 and 13 across all monkey datasets analyzed in Chapter 4 (**Supplementary Figure 4-3**). This is relatively low in comparison to the roughly hundred neurons we can record and the millions of neurons that modify their activity within the motor cortex. This observation underscores the low dimensionality characteristic of these manifolds, highlighting the utility of our linear approximation approach in the context of iBCIs.

In the same chapter, I explored two methods for predicting the muscle activity of a "source" monkey iBCI user from the latent neural activity of a "target" monkey or a paralyzed human iBCI user, both undertaking a center-out task. The first, "direct decoding," directly mapped the target user's neural signals into the EMG signals of the source monkey. This method successfully predicted the EMG activities of a source monkey using M1 manifolds from another monkey and transformed latent signals on the M1 manifolds of a human with paralysis into muscle activities that resembled those of a monkey. The second method, "transfer decoding," involved the transfer of a decoder trained on neural and EMG data from a source monkey to other target users. This approach yielded similar EMG predictions without the necessity for a new decoder computation. Although direct decoding slightly outperformed transfer decoding when it was retrained across sessions (**Figure 4-5, 4-8, and 4-9**), this advantage dissipated when the direct decoder was fixed for one session and session-aligned with CCA akin to the transfer decoding approach (**Figure 4-6**).

Comparison of the potential of direct and transfer decoding for closed-loop iBCIs

Both the direct and transfer decoding methods demonstrated promising potential for predicting unobservable muscle activity based on latent neural signals. However, each method presents unique advantages and potential challenges that must be addressed for effective application in closed-loop situations.

Direct decoding was a simple yet effective approach that consistently outperformed transfer decoding when it was retrained at every session. This method, suitable only for simple behaviors with stereotypic trajectories, does not necessarily have to rely on a monkey as the source user. Instead, in future work, we could collect EMGs from able-bodied humans performing specific tasks and construct a direct decoder.

Notably, the redundancy of hand musculature permits most actions to be produced with a variety of different patterns of muscle activity (d'Avella, Saltiel, and Bizzi 2003; Bernstein 1966). In fact, the monkeys that performed the same center-out task had different patterns of EMG, even though their motor patterns were self-consistent across sessions (**Supplementary Figure 4-2**). It is reasonable to expect that the human participant intended to use a different pattern of EMG from that of the source monkey. The direct decoding method was capable of predicting EMG activities that were biomechanically appropriate for the task at hand, given the similarity between human and monkey hand biomechanics. However, this discrepancy might present an intriguing motor-adaptation problem as the user learns to generate monkey-like muscle commands.

On the other hand, transfer decoding, despite appearing slightly less effective than direct decoding, has the potential for wider applications and presents several distinct advantages for future research. While the neural alignment method utilized in this approach, Canonical Correlation Analysis (CCA), is primarily effective on tasks with well-defined onsets and offsets, it

could conceivably be substituted with alternative unsupervised techniques, such as domain adaptation methods. These methods align the statistical properties of point clouds independently from the time course of the associated signals and have been successfully used to align neural signals from the same monkey performing different behaviors over time, including tasks with minimal trial structure (Degenhart et al. 2020; Farshchian et al. 2019; Karpowicz et al. 2022; Ma et al. 2022; Sun and Saenko 2016). However, these unsupervised methods have yet to be successfully applied to match neural activity statistics collected from multiple monkeys performing varied and unstructured movements typical of daily living activities. Thus, further investigation of these and related domain adaptation methods under more realistic conditions and across different users is warranted.

Transfer decoding also has an inherent advantage in its compatibility with abundant monkey-based data. Given the relative ease of acquiring training data from monkeys and the vast existing datasets from various motor behaviors, it is feasible to train monkey decoders that are more complex than the one tested in this study. Recent studies have shown that deep-learning-based decoders that correlate neural activity with behavior outperform simple linear decoders (Sussillo et al. 2016; Willsey et al. 2022; Deo et al. 2023). These advanced decoders require more data and a longer training time than linear decoders, which may pose a challenge for human iBCI research due to the difficulty and invasiveness of collecting neural data. This limitation is not present in monkey settings. The extent to which these advanced decoders can be transferred to decode EMGs for humans is yet to be fully explored.

In conclusion, while both methods have shown promise, further investigation is required to assess the application of these advanced decoders in decoding EMGs in humans, marking an exciting direction for future research.

The low dimensionality of M1 in natural settings may help translate iBCIs outside of the laboratory

The EMG-based iBCI decoding approaches I examined in Chapter 4 were primarily focused on controlled laboratory tasks. Nonetheless, there is a considerable possibility that these EMG-decoders could be adapted for use in natural settings and for more complex tasks. This potential can be attributed to the similar low neural dimensionality of these tasks, as demonstrated in Chapter 3.

Theoretical advancements have underscored the correlation between the embedding dimensionality of M1 manifolds and the accuracy of decoding movement parameters. The inherently low embedding dimensionality of M1 enables accurate decoding from a substantially smaller subset of the millions of active neurons, assuming a sufficient amount of data is available (P. Gao and Ganguli 2015). Are natural behaviors in unconstrained settings associated with similarly low-dimensional M1 manifolds? The answer to this question has implications for translating the intracortical brain-computer interfaces (iBCIs) from controlled laboratory settings to unconfined environments, thereby enabling broader and more practical usage.

In the third chapter of my thesis, I put these theoretical assertions to the test using simultaneous M1 and EMG recordings from monkeys in the unconstrained cage setting, which was as natural and “real-world” as we could get with monitored monkeys. I successfully demonstrated that the M1 manifolds were still low-dimensional, and we could decode EMGs regardless of the experimental setting (**Figure 4-6, Supplementary Figure 4-5 and 4-6**). This finding confirms the applicability of the theoretical relationship between low embedding dimensionality and linear decodability beyond controlled laboratory conditions, demonstrating its robustness even in unconstrained settings.

My results highlight the potential of current electrophysiology hardware technology, which samples from merely hundreds to thousands of neurons, for iBCIs to operate within activities of

daily living outside the laboratory. Despite the small neural sample size, we can effectively identify the neural manifolds associated with unconstrained tasks. This illustrates that the generalizable representations in M1, though low-dimensional, can be decoded reliably even when our recording devices vastly under-sample the active neurons in the motor cortex. Given the “above-Moore’s Law” advancements in recording technologies that allow us to sample from an increasing number of neurons (Urai et al. 2022), the integration of iBCIs into the daily lives of individuals with paralysis is becoming more feasible. This progress draws us closer to the goal of transitioning iBCIs from being a promising yet laboratory-bound technology to becoming an innovative tool in real-world scenarios.

A dimensionality estimation pipeline for neural populations beyond M1

In Chapter 2 of my thesis, I explored various methods for estimating the dimensionality of neural recordings, with the motivation to answer questions regarding the primary motor cortex (M1). I used this pipeline to investigate M1 manifolds in Chapters 2 and 3. However, this pipeline can be used in other neural populations as well.

In Chapter 2, I found PCA with an arbitrary variance cutoff, the embedding dimensionality estimator that many use, to be the least accurate method among those tested. Instead, Parallel Analysis (PA) emerged as a more effective linear approach for estimating embedding dimensionality, even for mildly noisy and nonlinear datasets. It is essential to note that all the methods assessed in this study were unsupervised, making them applicable to datasets with no repeated trial structure.

In contrast to what I found in Chapter 3 about M1, there is growing evidence suggesting that neural manifolds may be nonlinear for some of the sensory and cognitive areas (Low et al. 2018; Nieh et al. 2021; Gardner et al. 2022; De and Chaudhuri 2022; Y. Li et al. 2023). In these scenarios, one would expect a larger discrepancy between the intrinsic and embedding

dimensionality than I have observed in M1. The pipeline that I used would help identify computational principles from these brain areas just as I did in M1. Below is a general recipe that would work with any set of neural recordings.

First, one should begin with Parallel Analysis for an embedding dimensionality estimate of the neural data. Parallel Analysis was not only accurate for estimating embedding dimensionality, but also computationally efficient and relatively robust to noise. Next, one should perform a denoising procedure on the signals. This step involves projecting neural signals into a subspace of dimension given by the PA estimate of the embedding dimensionality and reconstructing the signals based on these projections. There are two options: PCA-based denoising, which is straightforward and resource-efficient, and the JAE that I proposed, which can account for potential nonlinearities. The effectiveness of denoising (and the first indication of the nonlinearity in the neural data) can be evaluated through the reconstruction accuracy between the denoised (reconstructed) signals and the original signals. If the reconstruction accuracy between PCA and JAE is similar, then the underlying manifold is likely linear. In this scenario, the intrinsic dimensionality estimate would be very close to the embedding dimensionality estimate. If the JAE yields a significantly superior reconstruction performance, the underlying neural manifold is likely to be nonlinear. In this scenario, one could use Levina-Bickel Maximum Likelihood Estimation or Two Nearest Neighbors on the reconstructions of JAE for an estimate of the intrinsic dimensionality.

Limitations

In all chapters of my thesis, the behavioral tasks were simple. Particularly, in Chapter 3, I helped design cage experiments anticipating that the unconstrained tasks would be more complex. The goal was to quantify this anticipated complexity through a higher intrinsic dimensionality of EMG signals compared to constrained, lab-bound tasks. Despite this expectation, a clear trend to

support this hypothesis was absent. All tasks, even those in the more naturalistic cage setting, turned out to be relatively simple.

Simulations played a key role when I benchmarked the efficacy of the intrinsic and embedding dimensionality estimators, which I ultimately used throughout my thesis. In the simulations carried out in Chapter 2, I focused on modeling the essential features of experimental neural data but did not take into account some of their characteristics. For example, I only considered additive Gaussian isotropic noise, which may not fully reflect the noise models observed in real-world recordings. These recordings might have non-additive, non-isotropic, and non-Gaussian noise. Additionally, the JAE could have employed a variational framework, which is known to be more robust to noise than the simple feedforward approach that I took (Im et al. 2017; Kingma and Welling 2019). A deep learning based denoising approach that I referenced in the introduction of my thesis, LFADS, employs a variational framework and is the current state-of-the-art for denoising experimental data, albeit with high computational cost and significant difficulty in tuning its hyperparameters (Pandarinath, O'Shea, et al. 2018).

Another constraint in the simulations was that I normalized the firing rates of each channel to the [0,1] range, a simplification that did not mirror the range in the experimental neural firing rate. In reality, the range of firing can differ significantly, even across neurons of the same type. Furthermore, I simulated the firing rates using latent signals with the same first-order statistics as found in the actual data. While this decision overlooked some of the heterogeneity observed in the experimental neural data, the simulations proved to be vital to understand the best- and worst-case scenarios when computing the intrinsic and embedding dimensionality of real-world data.

In the context of Chapters 3 and 4, the use of linear models for decoding EMG signals could be considered a limitation. While these models offer interpretability and simplicity, they may not fully capture the complexity of the neural signals and the tasks performed. In Chapter 3, my interest

was to compare the relative performance of EMG decoding using latent signals on the neural manifolds and using all available neurons. In Chapter 4, I needed a simple decoder to test the relative performance of the two cross-user decoding approaches I tested. Linear models of EMG decoding were sufficient for my purpose in both chapters, even though they may fail to capture the potentially complex and nonlinear relationship between M1 and muscles.

Future directions

An important future direction could be toward the use of nonlinear and recurrent decoders instead of the linear decoders I used throughout my work. Given the modest nonlinearity of M1 and the current effectiveness of linear models, the proposed shift toward these complex decoders might seem counterintuitive. However, despite their utility, linear decoders fall short in accurately capturing the relationship between neural and muscle activity across a broad range of tasks.

The solution to this problem may lie in more nonlinear and recurrent decoders. With their potential to better capture data complexity, these sophisticated decoders could improve decoding accuracy and provide insights into the role of M1 in generating multiple behaviors. This insight is crucial for the real-world application of iBCIs, which need to handle diverse tasks of daily living.

Additionally, in Chapter 4, both direct and transfer decoding approaches faced limitations in generalization capabilities. Better decoders may help overcome these limitations. The ability to train complex, nonlinear decoders with large amounts of monkey data prior to augmentation with subject-specific data may be a potent approach to overcome limitations regarding task generalization.

In addition to the development of more sophisticated decoders, future work should also investigate how task complexity affects the dimensionality of EMG signals. For this purpose, we

can record EMG signals from able-bodied humans repeatedly performing tasks of varying complexity and identify a set of simple to complex tasks that we can teach monkeys to perform. For example, we could record EMGs from humans performing complex tasks that can be decomposed into simpler ones: tapping fingers in complicated patterns versus simple, individual patterns. Would the EMG dimensionality be different in these scenarios? If so, would the corresponding neural manifolds combine similarly as they do in the task space, where the complex task is composed of individual finger taps? Such investigations would illuminate the ambiguities around the current view on neural manifolds, which are thought to be task-specific (Gallego et al. 2017; Ebitz and Hayden 2021). A greater understanding of the relationship between task complexity and the corresponding neural manifold geometry could help us understand the nuances in M1's role in managing different types of motor tasks, from a combined set of simple to highly dexterous hand gestures.

In my current work, I used CCA for neural alignment. While CCA was effective, it has major limitations in terms of applicability to more complicated tasks. This constraint makes CCA inappropriate for the activities of daily living, which are likely to have minimal trial structures and consistency. To overcome this, I propose that CCA could be replaced with alternative unsupervised alignment methods, such as domain adaptation techniques (B. Sun, Feng, and Saenko 2015; B. Sun and Saenko 2016; Zhu et al. 2017; Park et al. 2020). Some of these techniques that match the distributions of source and target signals have already shown success in aligning neural signals from the same monkey performing different behaviors over time (Dyer et al. 2017; Farshchian et al. 2018; Degenhart et al. 2020b; Ma et al. 2022; Karpowicz et al. 2022). However, to date, they have not been successfully applied to match the statistics of neural activity collected from multiple monkeys, each performing varied and unstructured movements characteristic of everyday activities. Perhaps with the combination of

robust decoders, these unsupervised domain adaptation methods could operate in such real-world contexts, making significant headway in translating iBCIs outside the laboratory.

Conclusion

In this thesis, I explored the dimensionality and nonlinearity of M1 manifolds and investigated their implications for the computational principles across tasks and for improving iBCIs. In Chapter 2, I developed a novel, accurate methodology to estimate the dimensionality and nonlinearity of neural manifolds. In Chapter 3, I applied this methodology to M1 manifolds from monkeys performing unconstrained tasks. My findings revealed that the dimensionality of M1 manifolds was only slightly higher during unconstrained tasks and still much lower than the roughly 100 neurons we sampled. Furthermore, I found that the information related to muscle activation was predominantly in nearly linear regions in these slightly nonlinear manifolds. This observation held true for both constrained and unconstrained tasks and could serve as a potential explanation for the surprising effectiveness of linear methods used in understanding M1's function. Chapter 4 bridged these theoretical findings to practical applications in iBCIs. Using linearly approximated M1 manifolds, I explored direct and transfer decoding methods, both of which exhibited similar levels of accuracy and showed promise in aiding the voluntary activation of muscles, a critical limitation of the current observation-based iBCI decoders.

The exploration of neural dimensionality, drawing upon the accurate methods I developed and the definitions and interpretations provided by Gao and Ganguli (P. Gao et al. 2017), and later by Jazayeri and Ostojic (Jazayeri and Ostojic 2021), helped reveal that both intrinsic and embedding dimensionality of M1 are low for both constrained and unconstrained tasks. These findings imply that the computational strategy in M1 favors generalizability over expressivity, enabling the reliable generation of movement commands across a range of task constraints and contexts.

As showed by Gao and Ganguli in an earlier study (P. Gao and Ganguli 2015), the inherent low dimensionality of M1 manifolds enables accurate decoding from a considerably smaller subset of neurons, provided there is sufficient temporal data. In other words, even with millions of active M1 neurons, accurate decoding of movement variables can be achieved with a much smaller sample as long as M1 manifolds are low dimensional.

This insight, coupled with the findings of mild nonlinearity and low dimensionality during unconstrained tasks, has substantial implications for the practical application of iBCIs in real-world settings. It may be possible to develop reliable and efficient iBCIs for unconstrained tasks of daily living, despite the undersampling of the active neurons in the primary motor cortex.

The two approaches I tested for enabling cross-user EMG decoding from the low-dimensional activity on M1 manifolds, namely direct and transfer decoding, demonstrated a proof-of-concept for future applications for muscle-based iBCIs. These two approaches each had their unique strengths and weaknesses. For instance, direct decoding may have issues with adaptability due to its need for retraining. Transfer decoding, despite being slightly less effective, holds potential for broader applications for a wider range of motor behaviors.

An important future direction for my research lies in the development of more sophisticated methods that can extract information from multiple task-specific manifolds and unsupervised domain adaptation methods. These methods could better capture the nuances of neural activity associated with multiple users and task conditions and extend the applicability of iBCIs to a wide array of simple to complex motor tasks. Finally, I propose that we should closely investigate the complexity of tasks and dimensionality of muscle activation patterns measured by EMGs, which could help elucidate the intricate relationship between task complexity, neural manifold geometry, and the role of M1 in driving diverse motor tasks.

References

Abeles, M., H. Bergman, I. Gat, I. Meilijson, E. Seidemann, N. Tishby, and E. Vaadia. 1995. "Cortical Activity Flips among Quasi-Stationary States." *Proceedings of the National Academy of Sciences of the United States of America* 92 (19): 8616–20.

Ahrens, Misha B., Michael B. Orger, Drew N. Robson, Jennifer M. Li, and Philipp J. Keller. 2013. "Whole-Brain Functional Imaging at Cellular Resolution Using Light-Sheet Microscopy." *Nature Methods* 10 (5): 413–20.

Aimon, Sophie, Takeo Katsuki, Tongqiu Jia, Logan Grosenick, Michael Broxton, Karl Deisseroth, Terrence J. Sejnowski, and Ralph J. Greenspan. 2019. "Fast Near-Whole–Brain Imaging in Adult Drosophila during Responses to Stimuli and Behavior." *PLoS Biology* 17 (2): e2006732.

Ajiboye, A. Bolu, Francis R. Willett, Daniel R. Young, William D. Memberg, Brian A. Murphy, Jonathan P. Miller, Benjamin L. Walter, et al. 2017. "Restoration of Reaching and Grasping Movements through Brain-Controlled Muscle Stimulation in a Person with Tetraplegia: A Proof-of-Concept Demonstration." *The Lancet* 389 (10081): 1821–30.

Albergante, Luca, Jonathan Bac, and Andrei Zinovyev. 2019. "Estimating the Effective Dimension of Large Biological Datasets Using Fisher Separability Analysis." *ArXiv [Cs.LG]*. arXiv. http://arxiv.org/abs/1901.06328.

Alexander, G. E., and M. D. Crutcher. 1990. "Neural Representations of the Target (Goal) of Visually Guided Arm Movements in Three Motor Areas of the Monkey." *Journal of Neurophysiology* 64 (1): 164–78.

Allegra, Michele, Elena Facco, Francesco Denti, Alessandro Laio, and Antonietta Mira. 2020. "Data Segmentation Based on the Local Intrinsic Dimension." *Scientific Reports* 10 (1): 16449.

Altan, Ege, Sara A. Solla, Lee E. Miller, and Eric J. Perreault. 2021. "Estimating the Dimensionality of the Manifold Underlying Multi-Electrode Neural Recordings." *PLoS Computational Biology* 17 (11): e1008591.

Andersen, Richard A., Tyson Aflalo, and Spencer Kellis. 2019. "From Thought to Action: The Brain–Machine Interface in Posterior Parietal Cortex." *Proceedings of the National Academy of Sciences* 116 (52): 26274–79.

Anderson, Kim D. 2004. "Targeting Recovery: Priorities of the Spinal Cord-Injured Population." *Journal of Neurotrauma* 21 (10): 1371–83.

Ansuini, Laio, and Macke. 2019. "Intrinsic Dimension of Data Representations in Deep Neural Networks." *Advances in Neural Information Processing Systems*. https://proceedings.neurips.cc/paper/2019/hash/cfcce0621b49c983991ead4c3d4d3b6b-Abstract.html.

Atick, Joseph J., and A. Norman Redlich. 1990. "Towards a Theory of Early Visual Processing." *Neural Computation* 2 (3): 308–20.

Avella, Andrea d', Philippe Saltiel, and Emilio Bizzi. 2003. "Combinations of Muscle Synergies in the Construction of a Natural Motor Behavior." *Nature Neuroscience* 6 (3): 300–308.

Bac, Jonathan, Evgeny M. Mirkes, Alexander N. Gorban, Ivan Tyukin, and Andrei Zinovyev. 2021. "Scikit-Dimension: A Python Package for Intrinsic Dimension Estimation." *Entropy* 23 (10). https://doi.org/10.3390/e23101368.

Badre, David, Apoorva Bhandari, Haley Keglovits, and Atsushi Kikumoto. 2021. "The Dimensionality of Neural Representations for Control." *Current Opinion in Behavioral Sciences* 38 (April): 20–28.

Bajd, Tadej, and Marko Munih. 2010. "Basic Functional Electrical Stimulation (FES) of Extremities: An Engineer's View." *Technology and Health Care: Official Journal of the European Society for Engineering and Medicine* 18 (4–5): 361–69.

Barelli, Renato G., Valter F. Avelino, and Maria Claudia F. Castro. 2022. "STIMGRASP: A Home-Based Functional Electrical Stimulator for Grasp Restoration in Daily Activities." *Sensors* 23 (1). https://doi.org/10.3390/s23010010.

Batty, Eleanor, Josh Merel, Nora Brackbill, Alexander Heitman, Alexander Sher, Alan Litke, E. J. Chichilnisky, and Liam Paninski. 2022. "Multilayer Recurrent Network Models of Primate Retinal Ganglion Cell Responses." *International*. https://openreview.net/pdf?id=HkEI22jeg.

Beiran, Manuel, Alexis Dubreuil, Adrian Valente, Francesca Mastrogiuseppe, and Srdjan Ostojic. 2021. "Shaping Dynamics With Multiple Populations in Low-Rank Recurrent Networks." *Neural Computation* 33 (6): 1572–1615.

Bengio, Yoshua, Olivier Delalleau, and Nicolas Le Roux. 2005. "The Curse of Dimensionality for Local Kernel Machines." *Techn. Rep* 1258: 12.

Bernstein, Nikolai. 1966. "The Co-Ordination and Regulation of Movements." *The Co-Ordination and Regulation of Movements*. https://ci.nii.ac.jp/naid/10008376164/.

Betz, W. 1874. "Anatomischer Nachweis Zweier Gehirncentra." *Zentralbl Med Wiss*.

Bhandari, Apoorva, Christopher Gagne, and David Badre. 2018. "Just above Chance: Is It Harder to Decode Information from Prefrontal Cortex Hemodynamic Activity Patterns?" *Journal of Cognitive Neuroscience* 30 (10): 1473–98.

Bialek, W. 2020. "What Do We Mean by the Dimensionality of Behavior?" *ArXiv Preprint ArXiv:2008.09574*. http://arxiv.org/abs/2008.09574.

Bickel, C. Scott, Chris M. Gregory, and Jesse C. Dean. 2011. "Motor Unit Recruitment during Neuromuscular Electrical Stimulation: A Critical Appraisal." *European Journal of Applied Physiology* 111 (10): 2399–2407.

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): 043002.

Blana, Dimitra, Edward K. Chadwick, Antonie J. van den Bogert, and Wendy M. Murray. 2017. “Real-Time Simulation of Hand Motion for Prosthesis Control.” *Computer Methods in Biomechanics and Biomedical Engineering* 20 (5): 540–49.

Blana, Dimitra, Antonie J. Van Den Bogert, Wendy M. Murray, Amartya Ganguly, Agamemnon Krasoulis, Kianoush Nazarpour, and Edward K. Chadwick. 2020. “Model-Based Control of Individual Finger Movements for Prosthetic Hand Function.” *IEEE Transactions on Neural Systems and Rehabilitation Engineering: A Publication of the IEEE Engineering in Medicine and Biology Society* 28 (3): 612–20.

Bollimunta, Anil, Douglas Totten, and Jochen Ditterich. 2012. “Neural Dynamics of Choice: Single-Trial Analysis of Decision-Related Activity in Parietal Cortex.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 32 (37): 12684–701.

Boon, Mei Ying, Bruce I. Henry, Catherine M. Suttle, and Stephen J. Dain. 2008. “The Correlation Dimension: A Useful Objective Measure of the Transient Visual Evoked Potential?” *Journal of Vision* 8 (1): 6.1-21.

Boots, Byron, and Geoff Gordon. 2012. “Two-Manifold Problems with Applications to Nonlinear System Identification.” *ArXiv [Cs.LG]*. arXiv. http://arxiv.org/abs/1206.4648.

Boser, Bernhard E., Isabelle M. Guyon, and Vladimir N. Vapnik. 1992. “A Training Algorithm for Optimal Margin Classifiers.” In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, 144–52. COLT ’92. New York, NY, USA: Association for Computing Machinery.

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 (7602): 247–50.

Brandman, David M., Tommy Hosman, Jad Saab, Michael C. Burkhart, Benjamin E. Shanahan, John G. Ciancibello, Anish A. Sarma, et al. 2018. “Rapid Calibration of an Intracortical Brain-Computer Interface for People with Tetraplegia.” *Journal of Neural Engineering* 15 (2): 026007.

Brodmann, Korbinian. 1909. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth.

Broome, Bede M., Vivek Jayaraman, and Gilles Laurent. 2006. “Encoding and Decoding of Overlapping Odor Sequences.” *Neuron* 51 (4): 467–82.

Bryden, Anne, Kevin Kilgore, Robert Kirsch, William Memberg, Hunter Peckham, and Michael Keith. 2005. “An Implanted Neuroprosthesis for High Tetraplegia.” *Topics in Spinal Cord Injury Rehabilitation* 10 (3): 38–52.

Buesing, Lars, Jakob H. Macke, and Maneesh Sahani. 2012. “Spectral Learning of Linear Dynamics from Generalised-Linear Observations with Application to Neural Population Data.” *Advances in Neural Information Processing Systems* 25. https://proceedings.neurips.cc/paper/2012/hash/d58072be2820e8682c0a27c0518e805e-Abstract.html.

Buja, A., and N. Eyuboglu. 1992. “Remarks on Parallel Analysis.” *Multivariate Behavioral Research* 27 (4): 509–40.

Camastra, Francesco. 2003. “Data Dimensionality Estimation Methods: A Survey.” *Pattern Recognition* 36 (12): 2945–54.

Camastra, Francesco, and Antonino Staiano. 2016. “Intrinsic Dimension Estimation: Advances and Open Problems.” *Information Sciences* 328 (January): 26–41.

Camastra, Francesco, and A. Vinciarelli. 2002. “Estimating the Intrinsic Dimension of Data with a Fractal-Based Method.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (10): 1404–7.

Caminiti, R., P. B. Johnson, and A. Urbano. 1990. “Making Arm Movements within Different Parts of Space: Dynamic Aspects in the Primate Motor Cortex.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 10 (7): 2039–58.

Campadelli, P., E. Casiraghi, C. Ceruti, and A. Rozza. 2015. “Intrinsic Dimension Estimation: Relevant Techniques and a Benchmark Framework.” *Mathematical Problems in Engineering* 2015 (October). https://doi.org/10.1155/2015/759567.

Carmena, Jose M., Mikhail A. Lebedev, Roy E. Crist, Joseph E. O’Doherty, David M. Santucci, Dragan F. Dimitrov, Parag G. Patil, Craig S. Henriquez, and Miguel A. L. Nicolelis. 2003a. “Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates.” *PLoS Biology* 1 (2): e42.

———. 2003b. “Learning to Control a Brain-Machine Interface for Reaching and Grasping by Primates.” *PLoS Biology* 1 (2): E42.

Ceruti, Claudio, Simone Bassis, Alessandro Rozza, Gabriele Lombardi, Elena Casiraghi, and Paola Campadelli. 2014. “DANCo: An Intrinsic Dimensionality Estimator Exploiting Angle and Norm Concentration.” *Pattern Recognition* 47 (8): 2569–81.

Chandak, Rishabh, and Barani Raman. 2021. “Neural Manifolds for Odor-Driven Innate and Acquired Appetitive Preferences.” *BioRxiv*. https://doi.org/10.1101/2021.08.05.455310.

Chaudhuri, Rishidev, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. 2019. “The Intrinsic Attractor Manifold and Population Dynamics of a Canonical Cognitive Circuit across Waking and Sleep.” *Nature Neuroscience* 22 (9): 1512–20.

Chen, Cathy S., R. Becket Ebitz, Sylvia R. Bindas, A. David Redish, Benjamin Y. Hayden, and Nicola M. Grissom. 2021. “Divergent Strategies for Learning in Males and Females.” *Current Biology: CB* 31 (1): 39-50.e4.

Chen, Hung-Tu, Jeremy R. Manning, and Matthijs A. A. van der Meer. 2021. “Between-Subject Prediction Reveals a Shared Representational Geometry in the Rodent Hippocampus.” *Current Biology: CB* 31 (19): 4293-4304.e5.

Chen, Xiuye, Yu Mu, Yu Hu, Aaron T. Kuan, Maxim Nikitchenko, Owen Randlett, Alex B. Chen, et al. 2018. “Brain-Wide Organization of Neuronal Activity and Convergent Sensorimotor Transformations in Larval Zebrafish.” *Neuron* 100 (4): 876-890.e5.

Cheney, P. D., and E. E. Fetz. 1980. “Functional Classes of Primate Corticomotoneuronal Cells and Their Relation to Active Force.” *Journal of Neurophysiology* 44 (4): 773–91.

Cherian, A., M. O. Krucoff, and L. E. Miller. 2011. “Motor Cortical Prediction of EMG: Evidence That a Kinetic Brain-Machine Interface May Be Robust across Altered Movement Dynamics.” *Journal of Neurophysiology* 106 (2): 564–75.

Chhatbar, Pratik Y., and Joseph T. Francis. 2013. “Towards a Naturalistic Brain-Machine Interface: Hybrid Torque and Position Control Allows Generalization to Novel Dynamics.” *PloS One* 8 (1): e52286.

Chung, S. Y., D. D. Lee, and H. Sompolinsky. 2018. “Classification and Geometry of General Perceptual Manifolds.” *Physical Review X*. https://journals.aps.org/prx/abstract/10.1103/PhysRevX.8.031003.

Churchland, Mark M., John P. Cunningham, Matthew T. Kaufman, Justin D. Foster, Paul Nuyujukian, Stephen I. Ryu, and Krishna V. Shenoy. 2012. “Neural Population Dynamics during Reaching.” *Nature* 487 (7405): 51–56.

Churchland, Mark M., John P. Cunningham, Matthew T. Kaufman, Stephen I. Ryu, and Krishna V. Shenoy. 2010. “Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine?” *Neuron* 68 (3): 387–400.

Churchland, Mark M., and Krishna V. Shenoy. 2007. “Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex.” *Journal of Neurophysiology* 97 (6): 4235–57.

Churchland, Mark M., Byron M. Yu, John P. Cunningham, Leo P. Sugrue, Marlene R. Cohen, Greg S. Corrado, William T. Newsome, et al. 2010. “Stimulus Onset Quenches Neural Variability: A Widespread Cortical Phenomenon.” *Nature Neuroscience* 13 (3): 369–78.

Cisek, Paul. 2006. “Preparing for Speed. Focus on ‘Preparatory Activity in Premotor and Motor Cortex Reflects the Speed of the Upcoming Reach.’” *Journal of Neurophysiology* 96 (6): 2842–43.

Cohen, Uri, Sueyeon Chung, Daniel D. Lee, and Haim Sompolinsky. 2020. “Separability and Geometry of Object Manifolds in Deep Neural Networks.” *Nature Communications* 11 (1): 746.

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.

Cunningham, John P., Vikash Gilja, Stephen I. Ryu, and Krishna V. Shenoy. 2009. “Methods for Estimating Neural Firing Rates, and Their Application to Brain–Machine Interfaces.” *Neural Networks: The Official Journal of the International Neural Network Society* 22 (9): 1235–46.

Cunningham, John P., and Byron M. Yu. 2014. “Dimensionality Reduction for Large-Scale Neural Recordings.” *Nature Neuroscience* 17 (11): 1500–1509.

Dabagia, Max, Konrad P. Kording, and Eva L. Dyer. 2022. “Comparing High-Dimensional Neural Recordings by Aligning Their Low-Dimensional Latent Representations.” *ArXiv [q-Bio.NC]*. arXiv. http://arxiv.org/abs/2205.08413.

Daly, Janis J., Janice Zimbelman, Kristen L. Roenigk, Jessica P. McCabe, Jean M. Rogers, Kristi Butler, Richard Burdsall, John P. Holcomb, E. Byron Marsolais, and Robert L. Ruff. 2011. “Recovery of Coordinated Gait: Randomized Controlled Stroke Trial of Functional Electrical Stimulation (FES) versus No FES, with Weight-Supported Treadmill and over-Ground Training.” *Neurorehabilitation and Neural Repair* 25 (7): 588–96.

De, Anandita, and Rishidev Chaudhuri. 2022. “Common Population Codes Produce Extremely Nonlinear Neural Manifolds.” *BioRxiv*. https://doi.org/10.1101/2022.09.27.509823.

Degenhart, Alan D., William E. Bishop, Emily R. Oby, Elizabeth C. Tyler-Kabara, Steven M. Chase, Aaron P. Batista, and Byron M. Yu. 2020a. “Stabilization of a Brain–Computer Interface via the Alignment of Low-Dimensional Spaces of Neural Activity.” *Nature Biomedical Engineering* 4 (7): 672–85.

———. 2020b. “Stabilization of a Brain-Computer Interface via the Alignment of Low-Dimensional Spaces of Neural Activity.” *Nature Biomedical Engineering* 4 (7): 672–85.

Demas, Jeffrey, Jason Manley, Frank Tejera, Kevin Barber, Hyewon Kim, Francisca Martínez Traub, Brandon Chen, and Alipasha Vaziri. 2021. “High-Speed, Cortex-Wide Volumetric Recording of Neuroactivity at Cellular Resolution Using Light Beads Microscopy.” *Nature Methods* 18 (9): 1103–11.

Denti, Francesco, Diego Doimo, Alessandro Laio, and Antonietta Mira. 2022. “The Generalized Ratios Intrinsic Dimension Estimator.” *Scientific Reports* 12 (1): 20005.

Deo, Darrel R., Francis R. Willett, Donald T. Avansino, Leigh R. Hochberg, Jaimie M. Henderson, and Krishna V. Shenoy. 2023. “Translating Deep Learning to Neuroprosthetic Control.” *BioRxiv : The Preprint Server for Biology*, April. https://doi.org/10.1101/2023.04.21.537581.

DiCarlo, James J., and David D. Cox. 2007. “Untangling Invariant Object Recognition.” *Trends in Cognitive Sciences* 11 (8): 333–41.

Dinno, Alexis. 2009. “Exploring the Sensitivity of Horn’s Parallel Analysis to the Distributional Form of Random Data.” *Multivariate Behavioral Research* 44 (3): 362–88.

Downey, John E., Nathaniel Schwed, Steven M. Chase, Andrew B. Schwartz, and Jennifer L. Collinger. 2018. “Intracortical Recording Stability in Human Brain-Computer Interface Users.” *Journal of Neural Engineering* 15 (4): 046016.

Dyer, Eva L., Mohammad Gheshlaghi Azar, Matthew G. Perich, Hugo L. Fernandes, Stephanie Naufel, Lee E. Miller, and Konrad P. Körding. 2017. “A Cryptography-Based Approach for Movement Decoding.” *Nature Biomedical Engineering* 1 (12): 967–76.

Ebitz, R. Becket, and Benjamin Y. Hayden. 2021. “The Population Doctrine in Cognitive Neuroscience.” *Neuron* 109 (19): 3055–68.

Eckmann, J-P, and D. Ruelle. 1992. “Fundamental Limitations for Estimating Dimensions and Lyapunov Exponents in Dynamical Systems.” *Physica D. Nonlinear Phenomena* 56 (2): 185–87.

Ehrlich, Daniel B., and John D. Murray. 2022. “Geometry of Neural Computation Unifies Working Memory and Planning.” *Proceedings of the National Academy of Sciences of the United States of America* 119 (37): e2115610119.

Einbeck, Jochen, and Zakiah Kalantana. 2013. “Intrinsic Dimensionality Estimation for High-Dimensional Data Sets: New Approaches for the Computation of Correlation Dimension.” *Journal of Emerging Technologies in Web Intelligence* 5 (2). https://doi.org/10.4304/jetwi.5.2.91-97.

Elsayed, Gamaleldin F., and John P. Cunningham. 2017. “Structure in Neural Population Recordings: An Expected Byproduct of Simpler Phenomena?” *Nature Neuroscience* 20 (9): 1310–18.

Elsayed, Gamaleldin F., Antonio H. Lara, Matthew T. Kaufman, Mark M. Churchland, and John P. Cunningham. 2016. “Reorganization between Preparatory and Movement Population Responses in Motor Cortex.” *Nature Communications* 7 (October): 13239.

Eraifej, John, William Clark, Benjamin France, Sebastian Desando, and David Moore. 2017. “Effectiveness of Upper Limb Functional Electrical Stimulation after Stroke for the Improvement of Activities of Daily Living and Motor Function: A Systematic Review and Meta-Analysis.” *Systematic Reviews* 6 (1): 40.

Ethier, C., E. R. Oby, M. J. Bauman, and L. E. Miller. 2012. “Restoration of Grasp Following Paralysis through Brain-Controlled Stimulation of Muscles.” *Nature* 485 (7398): 368–71.

Evarts, E. V. 1966. “Pyramidal Tract Activity Associated with a Conditioned Hand Movement in the Monkey.” *Journal of Neurophysiology* 29 (6): 1011–27.

———. 1968. “Relation of Pyramidal Tract Activity to Force Exerted during Voluntary Movement.” *Journal of Neurophysiology* 31 (1): 14–27.

Facco, Elena, Maria d’Errico, Alex Rodriguez, and Alessandro Laio. 2017. “Estimating the Intrinsic Dimension of Datasets by a Minimal Neighborhood Information.” *Scientific Reports* 7 (1): 12140.

Fagg, Andrew H., Gregory W. Ojakangas, Lee E. Miller, and Nicholas G. Hatsopoulos. 2009. “Kinetic Trajectory Decoding Using Motor Cortical Ensembles.” *IEEE Transactions on Neural Systems and Rehabilitation Engineering: A Publication of the IEEE Engineering in Medicine and Biology Society* 17 (5): 487–96.

Faisal, A. Aldo, Luc P. J. Selen, and Daniel M. Wolpert. 2008. “Noise in the Nervous System.” *Nature Reviews. Neuroscience* 9 (4): 292–303.

Fan, Jianqing, Fang Han, and Han Liu. 2014. “Challenges of Big Data Analysis.” *National Science Review* 1 (2): 293–314.

Farshchian, Ali, Juan A. Gallego, Joseph P. Cohen, Yoshua Bengio, Lee E. Miller, and Sara A. Solla. 2018. “Adversarial Domain Adaptation for Stable Brain-Machine Interfaces.” *ArXiv [Cs.LG]*. arXiv. http://arxiv.org/abs/1810.00045.

Ferrier, D. 1873. “Experimental Researches in Cerebral Physiology and Pathology.” *British Medical Journal* 1 (643): 457.

Fetz, E. E., and M. A. Baker. 1969. “Response Properties of Precentral Neurons in Awake Monkeys.” *The Physiologist*.

Fisher, R. A. 1936. “The Use of Multiple Measurements in Taxonomic Problems.” *Annals of Eugenics* 7 (2): 179–88.

Flint, Robert D., Michael R. Scheid, Zachary A. Wright, Sara A. Solla, and Marc W. Slutzky. 2016. “Long-Term Stability of Motor Cortical Activity: Implications for Brain Machine Interfaces and Optimal Feedback Control.” *Journal of Neuroscience* 36 (12): 3623–32.

Franklin, Scott B., David J. Gibson, Philip A. Robertson, John T. Pohlmann, and James S. Fralish. 1995. “Parallel Analysis: A Method for Determining Significant Principal Components.” *Journal of Vegetation Science: Official Organ of the International Association for Vegetation Science* 6 (1): 99–106.

Fritsch, G. 1870. “The Excitable Cerebral Cortex. Fritsch, G., Hitzig, E. Uber Die Elektrische Erregbarkeit Des Grosshirns.” *Arch. Anat. Physiol. Wissen*.

Fusi, Stefano, Earl K. Miller, and Mattia Rigotti. 2016. “Why Neurons Mix: High Dimensionality for Higher Cognition.” *Current Opinion in Neurobiology* 37 (April): 66–74.

Gallego, Juan A., Matthew G. Perich, Raeed H. Chowdhury, Sara A. Solla, and Lee E. Miller. 2020. “Long-Term Stability of Cortical Population Dynamics Underlying Consistent Behavior.” *Nature Neuroscience* 23 (2): 260–70.

Gallego, Juan A., Matthew G. Perich, Lee E. Miller, and Sara A. Solla. 2017. “Neural Manifolds for the Control of Movement.” *Neuron* 94 (5): 978–84.

Gallego, Juan A., Matthew G. Perich, Stephanie N. Naufel, Christian Ethier, Sara A. Solla, and Lee E. Miller. 2018. “Cortical Population Activity within a Preserved Neural Manifold Underlies Multiple Motor Behaviors.” *Nature Communications* 9 (1): 4233.

Gao, Peiran, and Surya Ganguli. 2015. “On Simplicity and Complexity in the Brave New World of Large-Scale Neuroscience.” *Current Opinion in Neurobiology* 32 (June): 148–55.

Gao, Peiran, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya Ganguli. 2017. “A Theory of Multineuronal Dimensionality, Dynamics and Measurement.” *BioRxiv*. https://doi.org/10.1101/214262.

Gao, Yuanjun, Evan W. Archer, Liam Paninski, and John P. Cunningham. 2016. “Linear Dynamical Neural Population Models through Nonlinear Embeddings.” *Adv. Neural Inf. Process. Syst.* 29.

Gardner, Richard J., Erik Hermansen, Marius Pachitariu, Yoram Burak, Nils A. Baas, Benjamin A. Dunn, May-Britt Moser, and Edvard I. Moser. 2022. “Toroidal Topology of Population Activity in Grid Cells.” *Nature* 602 (7895): 123–28.

Georgopoulos, A. P., R. Caminiti, J. F. Kalaska, and J. T. Massey. 1983. “Spatial Coding of Movement: A Hypothesis Concerning the Coding of Movement Direction by Motor Cortical Populations.” In *Neural Coding of Motor Performance*, 49:327–36. Berlin, Heidelberg: Springer Berlin Heidelberg.

Georgopoulos, A. P., J. F. Kalaska, R. Caminiti, and J. T. Massey. 1982. “On the Relations between the Direction of Two-Dimensional Arm Movements and Cell Discharge in Primate Motor Cortex.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 2 (11): 1527–37.

Georgopoulos, A. P., A. B. Schwartz, and R. E. Kettner. 1986. “Neuronal Population Coding of Movement Direction.” *Science* 233 (4771): 1416–19.

Glaser, Joshua I., Ari S. Benjamin, Raeed H. Chowdhury, Matthew G. Perich, Lee E. Miller, and Konrad P. Kording. 2020. “Machine Learning for Neural Decoding.” *ENeuro* 7 (4). https://doi.org/10.1523/ENEURO.0506-19.2020.

Glaser, Joshua I., Ari S. Benjamin, Roozbeh Farhoodi, and Konrad P. Kording. 2019. “The Roles of Supervised Machine Learning in Systems Neuroscience.” *Progress in Neurobiology* 175 (April): 126–37.

Glorot, Xavier, Antoine Bordes, and Yoshua Bengio. 2011. “Deep Sparse Rectifier Neural Networks.” In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, edited by Geoffrey Gordon, David Dunson, and Miroslav Dudík, 15:315–23. Proceedings of Machine Learning Research. Fort Lauderdale, FL, USA: PMLR.

Gobbo, Massimiliano, Nicola A. Maffiuletti, Claudio Orizio, and Marco A. Minetto. 2014. “Muscle Motor Point Identification Is Essential for Optimizing Neuromuscular Electrical Stimulation Use.” *Journal of Neuroengineering and Rehabilitation* 11 (February): 17.

Gorban, A. N., and I. Y. Tyukin. 2018. “Blessing of Dimensionality: Mathematical Foundations of the Statistical Physics of Data.” *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences* 376 (2118). https://doi.org/10.1098/rsta.2017.0237.

Gorban, Alexander N., Valery A. Makarov, and Ivan Y. Tyukin. 2020. “High-Dimensional Brain in a High-Dimensional World: Blessing of Dimensionality.” *Entropy* 22 (1). https://doi.org/10.3390/e22010082.

Granat, M. H., A. C. Ferguson, B. J. Andrews, and M. Delargy. 1993. “The Role of Functional Electrical Stimulation in the Rehabilitation of Patients with Incomplete Spinal Cord Injury--Observed Benefits during Gait Studies.” *Paraplegia* 31 (4): 207–15.

Grassberger, Peter. 1983. “Generalized Dimensions of Strange Attractors.” *Physics Letters. A* 97 (6): 227–30.

Grassberger, Peter, and Itamar Procaccia. 1983. “Measuring the Strangeness of Strange Attractors.” *Physica D. Nonlinear Phenomena* 9 (1): 189–208.

Harvey, Christopher D., Philip Coen, and David W. Tank. 2012. “Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task.” *Nature* 484 (7392): 62–68.

Hennig, Jay A., Matthew D. Golub, Peter J. Lund, Patrick T. Sadtler, Emily R. Oby, Kristin M. Quick, Stephen I. Ryu, et al. 2018. “Constraints on Neural Redundancy.” *ELife* 7 (August). https://doi.org/10.7554/eLife.36774.

Hepp-Reymond, M. C., E. J. Hüsler, M. A. Maier, and H. X. Ql. 1994. “Force-Related Neuronal Activity in Two Regions of the Primate Ventral Premotor Cortex.” *Canadian Journal of Physiology and Pharmacology* 72 (5): 571–79.

Hinton, G. E., and R. R. Salakhutdinov. 2006. “Reducing the Dimensionality of Data with Neural Networks.” *Science* 313 (5786): 504–7.

Hobby, J., P. N. Taylor, and J. Esnouf. 2001. “Restoration of Tetraplegic Hand Function by Use of the Neurocontrol Freehand System.” *Journal of Hand Surgery* 26 (5): 459–64.

Hochberg, Leigh R., Daniel Bacher, Beata Jarosiewicz, Nicolas Y. Masse, John D. Simeral, Joern Vogel, Sami Haddadin, et al. 2012. “Reach and Grasp by People with Tetraplegia Using a Neurally Controlled Robotic Arm.” *Nature* 485 (7398): 372–75.

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.

Holdefer, R. N., and L. E. Miller. 2002. “Primary Motor Cortical Neurons Encode Functional Muscle Synergies.” *Experimental Brain Research. Experimentelle Hirnforschung. Experimentation Cerebrale* 146 (2): 233–43.

Hong, Guosong, and Charles M. Lieber. 2019. “Novel Electrode Technologies for Neural Recordings.” *Nature Reviews. Neuroscience* 20 (6): 330–45.

Horn, J. L. 1965. “A Rationale and Test For The Number Of Factors In Factor Analysis.” *Psychometrika* 30 (June): 179–85.

Huang, Lawrence, Peter Ledochowitsch, Ulf Knoblich, Jérôme Lecoq, Gabe J. Murphy, R. Clay Reid, Saskia Ej de Vries, et al. 2021. “Relationship between Simultaneously Recorded Spiking Activity and Fluorescence Signal in GCaMP6 Transgenic Mice.” *ELife* 10 (March). https://doi.org/10.7554/eLife.51675.

Humphrey, D. R., E. M. Schmidt, and W. D. Thompson. 1970. “Predicting Measures of Motor Performance from Multiple Cortical Spike Trains.” *Science* 170 (3959): 758–62.

Ibitoye, Morufu Olusola, Nur Azah Hamzaid, Nazirah Hasnan, Ahmad Khairi Abdul Wahab, and Glen M. Davis. 2016. “Strategies for Rapid Muscle Fatigue Reduction during FES Exercise in Individuals with Spinal Cord Injury: A Systematic Review.” *PloS One* 11 (2): e0149024.

Im, Daniel Im, Sungjin Ahn, Roland Memisevic, and Yoshua Bengio. 2017. “Denoising Criterion for Variational Auto-Encoding Framework.” *Proceedings of the ... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence* 31 (1). https://doi.org/10.1609/aaai.v31i1.10777.

Jazayeri, Mehrdad, and Srdjan Ostojic. 2021. “Interpreting Neural Computations by Examining Intrinsic and Embedding Dimensionality of Neural Activity.” *Current Opinion in Neurobiology* 70 (October): 113–20.

Kalantan, Zakiah, and Jochen Einbeck. 2012. “On the Computation of the Correlation Integral for Fractal Dimension Estimation.” In *2012 International Conference on Statistics in Science, Business and Engineering (ICSSBE)*, 1–6.

Kalaska, J. F., and M. L. Hyde. 1985. “Area 4 and Area 5: Differences between the Load Direction-Dependent Discharge Variability of Cells during Active Postural Fixation.” *Experimental Brain Research. Experimentelle Hirnforschung. Experimentation Cerebrale* 59 (1): 197–202.

Karpowicz, Brianna M., Yahia H. Ali, Lahiru N. Wimalasena, Andrew R. Sedler, Mohammad Reza Keshtkaran, Kevin Bodkin, Xuan Ma, Lee E. Miller, and Chethan Pandarinath. 2022. “Stabilizing Brain-Computer Interfaces through Alignment of Latent Dynamics.” *BioRxiv*. https://doi.org/10.1101/2022.04.06.487388.

Kato, Saul, Harris S. Kaplan, Tina Schrödel, Susanne Skora, Theodore H. Lindsay, Eviatar Yemini, Shawn Lockery, and Manuel Zimmer. 2015. “Global Brain Dynamics Embed the Motor Command Sequence of Caenorhabditis Elegans.” *Cell* 163 (3): 656–69.

Kaufman, Matthew T., Marcus K. Benna, Mattia Rigotti, Fabio Stefanini, Stefano Fusi, and Anne K. Churchland. 2022. “The Implications of Categorical and Category-Free Mixed Selectivity on Representational Geometries.” *Current Opinion in Neurobiology* 77 (December): 102644.

Kaufman, Matthew T., Mark M. Churchland, Stephen I. Ryu, and Krishna V. Shenoy. 2014. “Cortical Activity in the Null Space: Permitting Preparation without Movement.” *Nature Neuroscience* 17 (3): 440–48.

Kégl, Balázs. 2002. “Intrinsic Dimension Estimation Using Packing Numbers.” *Adv. Neural Inf. Process. Syst.* 15.

Kilgore, Kevin L., Harry A. Hoyen, Anne M. Bryden, Ronald L. Hart, Michael W. Keith, and P. Hunter Peckham. 2008. “An Implanted Upper-Extremity Neuroprosthesis Using Myoelectric Control.” *The Journal of Hand Surgery* 33 (4): 539–50.

Kingma, Diederik P., and Max Welling. 2019. “An Introduction to Variational Autoencoders.” *Foundations and Trends® in Machine Learning* 12 (4): 307–92.

Kleinfeld, David, Lan Luan, Partha P. Mitra, Jacob T. Robinson, Rahul Sarpeshkar, Kenneth Shepard, Chong Xie, and Timothy D. Harris. 2019. “Can One Concurrently Record Electrical Spikes from Every Neuron in a Mammalian Brain?” *Neuron* 103 (6): 1005–15.

Kobayashi, T., K. Misaki, H. Nakagawa, S. Madokoro, T. Ota, H. Ihara, K. Tsuda, Y. Umezawa, J. Murayama, and K. Isaki. 2000. “Correlation Dimension of the Human Sleep Electroencephalogram.” *Psychiatry and Clinical Neurosciences* 54 (1): 11–16.

Kohn, Adam, Anna I. Jasper, João D. Semedo, Evren Gokcen, Christian K. Machens, and Byron M. Yu. 2020. “Principles of Corticocortical Communication: Proposed Schemes and Design Considerations.” *Trends in Neurosciences* 43 (9): 725–37.

Kralj, A., T. Bajd, and R. Turk. 1988. “Enhancement of Gait Restoration in Spinal Injured Patients by Functional Electrical Stimulation.” *Clinical Orthopaedics and Related Research*, no. 233 (August): 34–43.

Kulkarni, Jayant E., and Liam Paninski. 2007. “Common-Input Models for Multiple Neural Spike-Train Data.” *Network* 18 (4): 375–407.

Langdon, Christopher, Mikhail Genkin, and Tatiana A. Engel. 2023. “A Unifying Perspective on Neural Manifolds and Circuits for Cognition.” *Nature Reviews. Neuroscience*, April. https://doi.org/10.1038/s41583-023-00693-x.

Lanore, Frederic, N. Alex Cayco-Gajic, Harsha Gurnani, Diccon Coyle, and R. Angus Silver. 2021. “Cerebellar Granule Cell Axons Support High-Dimensional Representations.” *Nature Neuroscience* 24 (8): 1142–50.

Le Morvan, Marine, Andrei Zinovyev, and Jean-Philippe Vert. 2017. “NetNorM: Capturing Cancer-Relevant Information in Somatic Exome Mutation Data with Gene Networks for Cancer Stratification and Prognosis.” *PLoS Computational Biology* 13 (6): e1005573.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. “Deep Learning.” *Nature* 521 (7553): 436–44.

Lee, J. A., and M. Verleysen. 2007. *Nonlinear Dimensionality Reduction*. Springer New York.

Lee, John. 2010. *Introduction to Topological Manifolds*. Springer Science & Business Media.

Lemon, R. N., R. S. Johansson, and G. Westling. 1995. “Corticospinal Control during Reach, Grasp, and Precision Lift in Man.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 15 (9): 6145–56.

Levina, Elizaveta, and Peter Bickel. 2004. “Maximum Likelihood Estimation of Intrinsic Dimension.” *Adv. Neural Inf. Process. Syst.* 17.

Li, Chunyuan, Heerad Farkhoor, Rosanne Liu, and Jason Yosinski. 2018. “Measuring the Intrinsic Dimension of Objective Landscapes.” *ArXiv [Cs.LG]*. arXiv. http://arxiv.org/abs/1804.08838.

Li, Peter H., Jeffrey L. Gauthier, Max Schiff, Alexander Sher, Daniel Ahn, Greg D. Field, Martin Greschner, Edward M. Callaway, Alan M. Litke, and E. J. Chichilnisky. 2015. “Anatomical Identification of Extracellularly Recorded Cells in Large-Scale Multielectrode Recordings.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 35 (11): 4663–75.

Li, Yangang, Xinyun Zhu, Yu Qi, and Yueming Wang. 2023. “Revealing Unexpected Complex Encoding but Simple Decoding Mechanisms in Motor Cortex via Separating Task-Relevant Neural Signals.” *BioRxiv*. https://doi.org/10.1101/2022.11.13.515644.

Libby, Alexandra, and Timothy J. Buschman. 2021. “Rotational Dynamics Reduce Interference between Sensory and Memory Representations.” *Nature Neuroscience* 24 (5): 715–26.

Litwin-Kumar, Ashok, Kameron Decker Harris, Richard Axel, Haim Sompolinsky, and L. F. Abbott. 2017. “Optimal Degrees of Synaptic Connectivity.” *Neuron* 93 (5): 1153-1164.e7.

Lombardi, Gabriele, Alessandro Rozza, Claudio Ceruti, Elena Casiraghi, and Paola Campadelli. 2011. “Minimum Neighbor Distance Estimators of Intrinsic Dimension.” *Machine Learning and Knowledge Discovery in Databases*. https://doi.org/10.1007/978-3-642-23783-6_24.

Low, Ryan J., Sam Lewallen, Dmitriy Aronov, Rhino Nevers, and David W. Tank. 2018. “Probing Variability in a Cognitive Map Using Manifold Inference from Neural Dynamics.” *BioRxiv*. https://doi.org/10.1101/418939.

Luo, Shiyu, Haonan Xu, Yi Zuo, Xiaogang Liu, and Angelo H. All. 2020. “A Review of Functional Electrical Stimulation Treatment in Spinal Cord Injury.” *Neuromolecular Medicine* 22 (4): 447–63.

Lynch, Cheryl L., and Milos R. Popovic. 2008. “Functional Electrical Stimulation.” *IEEE Control Systems Magazine* 28 (2): 40–50.

Ma, Xuan, Fabio Rizzoglio, Eric J. Perreault, Lee E. Miller, and Ann Kennedy. 2022. “Using Adversarial Networks to Extend Brain Computer Interface Decoding Accuracy over Time.” *BioRxiv*. https://doi.org/10.1101/2022.08.26.504777.

Maass, Wolfgang. 2016. “Searching for Principles of Brain Computation.” *Current Opinion in Behavioral Sciences* 11 (October): 81–92.

Machens, Christian K., Ranulfo Romo, and Carlos D. Brody. 2010. “Functional, But Not Anatomical, Separation of ‘What’ and ‘When’ in Prefrontal Cortex.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 30 (1): 350–60.

Macke, Jakob H., Lars Buesing, John P. Cunningham, Byron M. Yu, Krishna V. Shenoy, and Maneesh Sahani. 2011. “Empirical Models of Spiking in Neural Populations.” *Advances in Neural Information Processing Systems* 24. https://proceedings.neurips.cc/paper/2011/hash/7143d7fbadfa4693b9eec507d9d37443-Abstract.html.

Maier, M. A., K. M. Bennett, M. C. Hepp-Reymond, and R. N. Lemon. 1993. “Contribution of the Monkey Corticomotoneuronal System to the Control of Force in Precision Grip.” *Journal of Neurophysiology* 69 (3): 772–85.

Mann, Kevin, Courtney L. Gallen, and Thomas R. Clandinin. 2017. “Whole-Brain Calcium Imaging Reveals an Intrinsic Functional Network in Drosophila.” *Current Biology: CB* 27 (15): 2389-2396.e4.

Mante, Valerio, David Sussillo, Krishna V. Shenoy, and William T. Newsome. 2013. “Context-Dependent Computation by Recurrent Dynamics in Prefrontal Cortex.” *Nature* 503 (7474): 78–84.

Marblestone, Adam H., Bradley M. Zamft, Yael G. Maguire, Mikhail G. Shapiro, Thaddeus R. Cybulski, Joshua I. Glaser, Dario Amodei, et al. 2013. “Physical Principles for Scalable Neural Recording.” *Frontiers in Computational Neuroscience* 7 (October): 137.

Marquez-Chin, Cesar, and Milos R. Popovic. 2020. “Functional Electrical Stimulation Therapy for Restoration of Motor Function after Spinal Cord Injury and Stroke: A Review.” *Biomedical Engineering Online* 19 (1): 34.

Mastrogiuseppe, Francesca, and Srdjan Ostojic. 2018. “Linking Connectivity, Dynamics, and Computations in Low-Rank Recurrent Neural Networks.” *Neuron* 99 (3): 609-623.e29.

Mazor, Ofer, and Gilles Laurent. 2005. “Transient Dynamics versus Fixed Points in Odor Representations by Locust Antennal Lobe Projection Neurons.” *Neuron* 48 (4): 661–73.

Mazzucato, Luca, Alfredo Fontanini, and Giancarlo La Camera. 2016. “Stimuli Reduce the Dimensionality of Cortical Activity.” *Frontiers in Systems Neuroscience* 10 (February): 11.

McFarland, Daniel C., Benjamin I. Binder-Markey, Jennifer A. Nichols, Sarah J. Wohlman, Marije de Bruin, and Wendy M. Murray. 2023. “A Musculoskeletal Model of the Hand and Wrist Capable of Simulating Functional Tasks.” *IEEE Transactions on Bio-Medical Engineering* 70 (5): 1424–35.

McInnes, Leland, John Healy, and James Melville. 2018. “UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction.” *ArXiv [Stat.ML]*. arXiv. http://arxiv.org/abs/1802.03426.

Melbaum, Svenja, Eleonora Russo, David Eriksson, Artur Schneider, Daniel Durstewitz, Thomas Brox, and Ilka Diester. 2022. “Conserved Structures of Neural Activity in Sensorimotor Cortex of Freely Moving Rats Allow Cross-Subject Decoding.” *Nature Communications* 13 (1): 7420.

Mitchell-Heggs, Rufus, Seigfred Prado, Giuseppe P. Gava, Mary Ann Go, and Simon R. Schultz. 2023. “Neural Manifold Analysis of Brain Circuit Dynamics in Health and Disease.” *Journal of Computational Neuroscience* 51 (1): 1–21.

Morrow, M. M., L. R. Jordan, and L. E. Miller. 2007. “Direct Comparison of the Task-Dependent Discharge of M1 in Hand Space and Muscle Space.” *Journal of Neurophysiology* 97 (2): 1786–98.

Musall, Simon, Matthew T. Kaufman, Ashley L. Juavinett, Steven Gluf, and Anne K. Churchland. 2019. “Single-Trial Neural Dynamics Are Dominated by Richly Varied Movements.” *Nature Neuroscience* 22 (10): 1677–86.

Musallam, S., B. D. Corneil, B. Greger, H. Scherberger, and R. A. Andersen. 2004. “Cognitive Control Signals for Neural Prosthetics.” *Science* 305 (5681): 258–62.

Musk, Elon, and Neuralink. 2019. “An Integrated Brain-Machine Interface Platform With Thousands of Channels.” *Journal of Medical Internet Research* 21 (10): e16194.

Mussa-Ivaldi, F. A. 1988. “Do Neurons in the Motor Cortex Encode Movement Direction? An Alternative Hypothesis.” *Neuroscience Letters* 91 (1): 106–11.

Musslick, Sebastian, and Jonathan D. Cohen. 2019. “A Mechanistic Account of Constraints on Control-Dependent Processing: Shared Representation, Conflict and Persistence.” cogsci.mindmodeling.org. 2019. https://cogsci.mindmodeling.org/2019/papers/0161/0161.pdf.

Musslick, Sebastian, Andrew Saxe, Kayhan Özcimder, Biswadip Dey, Greg Henselman, and Jonathan D. Cohen. 2017. “Multitasking Capability versus Learning Efficiency in Neural Network Architectures.” In . Cognitive Science Society. https://ora.ox.ac.uk/objects/uuid:5d234221-e892-4765-8741-775ae0b74043.

Nelder, J. A., and R. W. M. Wedderburn. 1972. “Generalized Linear Models.” *Journal of the Royal Statistical Society. Series A* 135 (3): 370.

Nieh, Edward H., Manuel Schottdorf, Nicolas W. Freeman, Ryan J. Low, Sam Lewallen, Sue Ann Koay, Lucas Pinto, Jeffrey L. Gauthier, Carlos D. Brody, and David W. Tank. 2021. “Geometry of Abstract Learned Knowledge in the Hippocampus.” *Nature* 595 (7865): 80–84.

Nordhausen, C. T., E. M. Maynard, and R. A. Normann. 1996. “Single Unit Recording Capabilities of a 100 Microelectrode Array.” *Brain Research* 726 (1–2): 129–40.

Obaid, Abdulmalik, Mina-Elraheb Hanna, Yu-Wei Wu, Mihaly Kollo, Romeo Racz, Matthew R. Angle, Jan Müller, et al. 2020. “Massively Parallel Microwire Arrays Integrated with CMOS Chips for Neural Recording.” *Science Advances* 6 (12): eaay2789.

Oby, Emily R., Christian Ethier, and Lee E. Miller. 2013. “Movement Representation in the Primary Motor Cortex and Its Contribution to Generalizable EMG Predictions.” *Journal of Neurophysiology* 109 (3): 666–78.

Opie, Nicholas. 2021. “The StentrodeTM Neural Interface System.” In *Brain-Computer Interface Research: A State-of-the-Art Summary 9*, edited by Christoph Guger, Brendan Z. Allison, and Michael Tangermann, 127–32. Cham: Springer International Publishing.

Pandarinath, Chethan, K. Cora Ames, Abigail A. Russo, Ali Farshchian, Lee E. Miller, Eva L. Dyer, and Jonathan C. Kao. 2018. “Latent Factors and Dynamics in Motor Cortex and Their Application to Brain-Machine Interfaces.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 38 (44): 9390–9401.

Pandarinath, Chethan, K. Cora Ames, Abigail A. Russo, Ali Farshchian, Lee E. Miller, Eva L. Dyer, and Jonathan C. Kao. 2018. “Latent Factors and Dynamics in Motor Cortex and Their Application to Brain–Machine Interfaces.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 38 (44): 9390–9401.

Pandarinath, Chethan, Vikash Gilja, Christine H. Blabe, Paul Nuyujukian, Anish A. Sarma, Brittany L. Sorice, Emad N. Eskandar, Leigh R. Hochberg, Jaimie M. Henderson, and

Krishna V. Shenoy. 2015. “Neural Population Dynamics in Human Motor Cortex during Movements in People with ALS.” *ELife* 4 (June): e07436.

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 (February). https://doi.org/10.7554/eLife.18554.

Pandarinath, Chethan, Daniel J. O’Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D. Stavisky, Jonathan C. Kao, Eric M. Trautmann, et al. 2018. “Inferring Single-Trial Neural Population Dynamics Using Sequential Auto-Encoders.” *Nature Methods* 15 (10): 805–15.

Pang, Rich, Benjamin J. Lansdell, and Adrienne L. Fairhall. 2016. “Dimensionality Reduction in Neuroscience.” *Current Biology: CB* 26 (14): R656-60.

Paninski, Liam, Shy Shoham, Matthew R. Fellows, Nicholas G. Hatsopoulos, and John P. Donoghue. 2004. “Superlinear Population Encoding of Dynamic Hand Trajectory in Primary Motor Cortex.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 24 (39): 8551–61.

Park, Taesung, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. 2020. “Contrastive Learning for Unpaired Image-to-Image Translation.” *ArXiv [Cs.CV]*. arXiv. http://arxiv.org/abs/2007.15651.

Pashkovski, Stan L., Giuliano Iurilli, David Brann, Daniel Chicharro, Kristen Drummey, Kevin M. Franks, Stefano Panzeri, and Sandeep Robert Datta. 2020. “Structure and Flexibility in Cortical Representations of Odour Space.” *Nature* 583 (7815): 253–58.

Peckham, P. H., M. W. Keith, K. L. Kilgore, J. H. Grill, K. S. Wuolle, G. B. Thrope, P. Gorman, et al. 2001. “Efficacy of an Implanted Neuroprosthesis for Restoring Hand Grasp in Tetraplegia: A Multicenter Study.” *Archives of Physical Medicine and Rehabilitation* 82 (10): 1380–88.

Peckham, P. H., E. B. Marsolais, and J. T. Mortimer. 1980. “Restoration of Key Grip and Release in the C6 Tetraplegic Patient through Functional Electrical Stimulation.” *The Journal of Hand Surgery* 5 (5): 462–69.

Penfield, Wilder, and Edwin Boldrey. 1937. “Somatic Motor and Sensory Representation in the Cerebral Cortex of Man as Studied by Electrical Stimulation.” *Brain: A Journal of Neurology* 60 (4): 389–443.

Perge, János A., Mark L. Homer, Wasim Q. Malik, Sydney Cash, Emad Eskandar, Gerhard Friehs, John P. Donoghue, and Leigh R. Hochberg. 2013. “Intra-Day Signal Instabilities Affect Decoding Performance in an Intracortical Neural Interface System.” *Journal of Neural Engineering* 10 (3): 036004.

Perich, Matthew G., Juan A. Gallego, and Lee E. Miller. 2018. “A Neural Population Mechanism for Rapid Learning.” *Neuron* 100 (4): 964-976.e7.

Perkel, D. H., G. L. Gerstein, and G. P. Moore. 1967a. “Neuronal Spike Trains and Stochastic Point Processes. I. The Single Spike Train.” *Biophysical Journal* 7 (4): 391–418.

———. 1967b. “Neuronal Spike Trains and Stochastic Point Processes. II. Simultaneous Spike Trains.” *Biophysical Journal* 7 (4): 419–40.

Perkins, Sean M., John P. Cunningham, Qi Wang, and Mark M. Churchland. 2023. “Simple Decoding of Behavior from a Complicated Neural Manifold.” *BioRxiv*. https://doi.org/10.1101/2023.04.05.535396.

Pinamonti, Giovanni, Jianbo Zhao, David E. Condon, Fabian Paul, Frank Noè, Douglas H. Turner, and Giovanni Bussi. 2017. “Predicting the Kinetics of RNA Oligonucleotides Using Markov State Models.” *Journal of Chemical Theory and Computation* 13 (2): 926–34.

Poliakov, A. V., and M. H. Schieber. 1999. “Limited Functional Grouping of Neurons in the Motor Cortex Hand Area during Individuated Finger Movements: A Cluster Analysis.” *Journal of Neurophysiology* 82 (6): 3488–3505.

Pollock, Eli, and Mehrdad Jazayeri. 2020. “Engineering Recurrent Neural Networks from Task-Relevant Manifolds and Dynamics.” *PLoS Computational Biology* 16 (8): e1008128.

Ponce-Alvarez, Adrián, Verónica Nácher, Rogelio Luna, Alexa Riehle, and Ranulfo Romo. 2012. “Dynamics of Cortical Neuronal Ensembles Transit from Decision Making to Storage for Later Report.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 32 (35): 11956–69.

Pope, Phillip, Chen Zhu, Ahmed Abdelkader, Micah Goldblum, and Tom Goldstein. 2021. “The Intrinsic Dimension of Images and Its Impact on Learning.” *ArXiv [Cs.CV]*. arXiv. http://arxiv.org/abs/2104.08894.

Popovic, M. R., T. Keller, I. P. I. Papas, V. Dietz, and M. Morari. 2001. “Surface-Stimulation Technology for Grasping and Walking Neuroprostheses.” *IEEE Engineering in Medicine and Biology Magazine: The Quarterly Magazine of the Engineering in Medicine & Biology Society* 20 (1): 82–93.

Popović, Mirjana B. 2003. “Control of Neural Prostheses for Grasping and Reaching.” *Medical Engineering & Physics* 25 (1): 41–50.

Popovic, Mirjana B., Dejan B. Popovic, Thomas Sinkjaer, Aleksandra Stefanovic, and Laszlo Schwirtlich. 2002. “Restitution of Reaching and Grasping Promoted by Functional Electrical Therapy.” *Artificial Organs* 26 (3): 271–75.

Prochazka, Arthur. 2017. “Neurophysiology and Neural Engineering: A Review.” *Journal of Neurophysiology* 118 (2): 1292–1309.

Pryluk, Raviv, Yoav Kfir, Hagar Gelbard-Sagiv, Itzhak Fried, and Rony Paz. 2019. “A Tradeoff in the Neural Code across Regions and Species.” *Cell* 176 (3): 597-609.e18.

Purves, D., G. J. Augustine, D. Fitzpatrick, and W. C. Hall. 2008. “Neuroscience. 4th.” *Sunderland, Mass.: Sinauer*.

Rastogi, Anisha, Carlos E. Vargas-Irwin, Francis R. Willett, Jessica Abreu, Douglas C. Crowder, Brian A. Murphy, William D. Memberg, et al. 2020. “Neural Representation of Observed, Imagined, and Attempted Grasping Force in Motor Cortex of Individuals with Chronic Tetraplegia.” *Scientific Reports* 10 (1): 1429.

Rathelot, Jean-Alban, and Peter L. Strick. 2009. “Subdivisions of Primary Motor Cortex Based on Cortico-Motoneuronal Cells.” *Proceedings of the National Academy of Sciences of the United States of America* 106 (3): 918–23.

Recanatesi, Stefano, Matthew Farrell, Guillaume Lajoie, Sophie Deneve, Mattia Rigotti, and Eric Shea-Brown. 2021. “Predictive Learning as a Network Mechanism for Extracting Low-Dimensional Latent Space Representations.” *Nature Communications* 12 (1): 1417.

Remington, Evan D., Devika Narain, Eghbal A. Hosseini, and Mehrdad Jazayeri. 2018. “Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics.” *Neuron* 98 (5): 1005-1019.e5.

Rigotti, Mattia, Omri Barak, Melissa R. Warden, Xiao-Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi. 2013. “The Importance of Mixed Selectivity in Complex Cognitive Tasks.” *Nature* 497 (7451): 585–90.

Ringach, Dario L. 2019. “The Geometry of Masking in Neural Populations.” *Nature Communications* 10 (1): 4879.

Roweis, S. T., and L. K. Saul. 2000. “Nonlinear Dimensionality Reduction by Locally Linear Embedding.” *Science* 290 (5500): 2323–26.

Rozza, A., G. Lombardi, C. Ceruti, E. Casiraghi, and P. Campadelli. 2012. “Novel High Intrinsic Dimensionality Estimators.” *Machine Learning* 89 (1): 37–65.

Rubin, Alon, Liron Sheintuch, Noa Brande-Eilat, Or Pinchasof, Yoav Rechavi, Nitzan Geva, and Yaniv Ziv. 2019. “Revealing Neural Correlates of Behavior without Behavioral Measurements.” *Nature Communications* 10 (1): 4745.

Russo, Abigail A., Sean R. Bittner, Sean M. Perkins, Jeffrey S. Seely, Brian M. London, Antonio H. Lara, Andrew Miri, et al. 2018. “Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response.” *Neuron* 97 (4): 953-966.e8.

Russo, Abigail A., Ramin Khajeh, Sean R. Bittner, Sean M. Perkins, John P. Cunningham, L. F. Abbott, and Mark M. Churchland. 2020. “Neural Trajectories in the Supplementary Motor Area and Motor Cortex Exhibit Distinct Geometries, Compatible with Different Classes of Computation.” *Neuron* 107 (4): 745-758.e6.

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.

Safaie, Mostafa, Joanna C. Chang, Junchol Park, Lee E. Miller, Joshua T. Dudman, Matthew G. Perich, and Juan A. Gallego. 2022. “Preserved Neural Population Dynamics across Animals Performing Similar Behaviour.” *BioRxiv*. https://doi.org/10.1101/2022.09.26.509498.

Sagiv, Yotam, Sebastian Musslick, Yael Niv, and Jonathan D. Cohen. 2020. “Efficiency of Learning vs. Processing: Towards a Normative Theory of Multitasking.” *ArXiv [q-Bio.NC]*. arXiv. http://arxiv.org/abs/2007.03124.

Saha, Debajit, Kevin Leong, Chao Li, Steven Peterson, Gregory Siegel, and Baranidharan Raman. 2013. “A Spatiotemporal Coding Mechanism for Background-Invariant Odor Recognition.” *Nature Neuroscience* 16 (12): 1830–39.

Salkoff, David B., Edward Zagha, Erin McCarthy, and David A. McCormick. 2020. “Movement and Performance Explain Widespread Cortical Activity in a Visual Detection Task.” *Cerebral Cortex* 30 (1): 421–37.

Santello, M., and J. F. Soechting. 2000. “Force Synergies for Multifingered Grasping.” *Experimental Brain Research. Experimentelle Hirnforschung. Experimentation Cerebrale* 133 (4): 457–67.

Santhanam, Gopal, Byron M. Yu, Vikash Gilja, Stephen I. Ryu, Afsheen Afshar, Maneesh Sahani, and Krishna V. Shenoy. 2009. “Factor-Analysis Methods for Higher-Performance Neural Prostheses.” *Journal of Neurophysiology* 102 (2): 1315–30.

Saxena, Shreya, and John P. Cunningham. 2019. “Towards the Neural Population Doctrine.” *Current Opinion in Neurobiology* 55 (April): 103–11.

Saxena, Shreya, Abigail A. Russo, John Cunningham, and Mark M. Churchland. 2022. “Motor Cortex Activity across Movement Speeds Is Predicted by Network-Level Strategies for Generating Muscle Activity.” *ELife* 11 (May). https://doi.org/10.7554/eLife.67620.

Schieber, M. H. 1999. “Somatotopic Gradients in the Distributed Organization of the Human Primary Motor Cortex Hand Area: Evidence from Small Infarcts.” *Experimental Brain Research. Experimentelle Hirnforschung. Experimentation Cerebrale* 128 (1–2): 139–48.

———. 2001. “Constraints on Somatotopic Organization in the Primary Motor Cortex.” *Journal of Neurophysiology* 86 (5): 2125–43.

Schieber, M. H., and L. S. Hibbard. 1993. “How Somatotopic Is the Motor Cortex Hand Area?” *Science* 261 (5120): 489–92.

Schieber, Marc H. 1996. “5 - Individuated Finger Movements: Rejecting the Labeled-Line Hypothesis.” In *Hand and Brain*, edited by Alan M. Wing, Patrick Haggard, and J. Randall Flanagan, 81–98. San Diego: Academic Press.

Schrödel, Tina, Robert Prevedel, Karin Aumayr, Manuel Zimmer, and Alipasha Vaziri. 2013. “Brain-Wide 3D Imaging of Neuronal Activity in Caenorhabditis Elegans with Sculpted Light.” *Nature Methods* 10 (10): 1013–20.

Schuessler, F., A. Dubreuil, and F. Mastrogiuseppe. 2020. “Dynamics of Random Recurrent Networks with Correlated Low-Rank Structure.” *Physical Review*. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.2.013111.

Schwartz, A. B., and D. W. Moran. 2000. “Arm Trajectory and Representation of Movement Processing in Motor Cortical Activity.” *The European Journal of Neuroscience* 12 (6): 1851–56.

Scott, S. H., and J. F. Kalaska. 1997. “Reaching Movements with Similar Hand Paths but Different Arm Orientations. I. Activity of Individual Cells in Motor Cortex.” *Journal of Neurophysiology* 77 (2): 826–52.

Scott, Stephen H. 2004. “Optimal Feedback Control and the Neural Basis of Volitional Motor Control.” *Nature Reviews. Neuroscience* 5 (7): 532–46.

Sedler, Andrew R., Christopher Versteeg, and Chethan Pandarinath. 2022. “Expressive Architectures Enhance Interpretability of Dynamics-Based Neural Population Models.” *ArXiv [q-Bio.NC]*. arXiv. http://arxiv.org/abs/2212.03771.

Seidemann, E., I. Meilijson, M. Abeles, H. Bergman, and E. Vaadia. 1996. “Simultaneously Recorded Single Units in the Frontal Cortex Go through Sequences of Discrete and Stable States in Monkeys Performing a Delayed Localization Task.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 16 (2): 752–68.

Semedo, João D., Amin Zandvakili, Christian K. Machens, Byron M. Yu, and Adam Kohn. 2019. “Cortical Areas Interact through a Communication Subspace.” *Neuron* 102 (1): 249-259.e4.

Sergio, Lauren E., Catherine Hamel-Pâquet, and John F. Kalaska. 2005. “Motor Cortex Neural Correlates of Output Kinematics and Kinetics during Isometric-Force and Arm-Reaching Tasks.” *Journal of Neurophysiology* 94 (4): 2353–78.

Sergio, Lauren E., and John F. Kalaska. 2003. “Systematic Changes in Motor Cortex Cell Activity with Arm Posture during Directional Isometric Force Generation.” *Journal of Neurophysiology* 89 (1): 212–28.

Serruya, M. D., N. G. Hatsopoulos, L. Paninski, and M. R. Fellows. 2002. “Instant Neural Control of a Movement Signal. Hands-Free Operation of a Cursor Can Be Achieved by a Few Neurons in the Motor Cortex.” *Nature*.

Shenoy, Krishna V., Maneesh Sahani, and Mark M. Churchland. 2013. “Cortical Control of Arm Movements: A Dynamical Systems Perspective.” *Annual Review of Neuroscience* 36 (July): 337–59.

Shih, Jerry J., Dean J. Krusienski, and Jonathan R. Wolpaw. 2012. “Brain-Computer Interfaces in Medicine.” *Mayo Clinic Proceedings. Mayo Clinic* 87 (3): 268–79.

Siegle, Joshua H., Peter Ledochowitsch, Xiaoxuan Jia, Daniel J. Millman, Gabriel K. Ocker, Shiella Caldejon, Linzy Casal, et al. 2021. “Reconciling Functional Differences in Populations of Neurons Recorded with Two-Photon Imaging and Electrophysiology.” *ELife* 10 (July). https://doi.org/10.7554/eLife.69068.

Simoncelli, E. P., and B. A. Olshausen. 2001. “Natural Image Statistics and Neural Representation.” *Annual Review of Neuroscience* 24: 1193–1216.

Smith, Anne C., and Emery N. Brown. 2003. “Estimating a State-Space Model from Point Process Observations.” *Neural Computation* 15 (5): 965–91.

Smith, Leonard A. 1988. “Intrinsic Limits on Dimension Calculations.” *Physics Letters. A* 133 (6): 283–88.

Snášel, Václav, Jana Nowaková, Fatos Xhafa, and Leonard Barolli. 2017. “Geometrical and Topological Approaches to Big Data.” *Future Generations Computer Systems: FGCS* 67 (February): 286–96.

Snoek, G. J., M. J. IJzerman, F. A. in ’t Groen, T. S. Stoffers, and G. Zilvold. 2000. “Use of the NESS Handmaster to Restore Handfunction in Tetraplegia: Clinical Experiences in Ten Patients.” *Spinal Cord* 38 (4): 244–49.

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). https://doi.org/10.1126/science.abf4588.

Steinmetz, Nicholas A., Christof Koch, Kenneth D. Harris, and Matteo Carandini. 2018. “Challenges and Opportunities for Large-Scale Electrophysiology with Neuropixels Probes.” *Current Opinion in Neurobiology* 50 (June): 92–100.

Steinmetz, Nicholas A., Peter Zatka-Haas, Matteo Carandini, and Kenneth D. Harris. 2019. “Distributed Coding of Choice, Action and Engagement across the Mouse Brain.” *Nature* 576 (7786): 266–73.

Stevenson, Ian H., and Konrad P. Kording. 2011. “How Advances in Neural Recording Affect Data Analysis.” *Nature Neuroscience* 14 (2): 139–42.

Stopfer, Mark, Vivek Jayaraman, and Gilles Laurent. 2003. “Intensity versus Identity Coding in an Olfactory System.” *Neuron* 39 (6): 991–1004.

Stringer, Carsen, Michalis Michaelos, Dmitri Tsyboulski, Sarah E. Lindo, and Marius Pachitariu. 2021. “High-Precision Coding in Visual Cortex.” *Cell* 184 (10): 2767-2778.e15.

Stringer, Carsen, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D. Harris. 2019. “High-Dimensional Geometry of Population Responses in Visual Cortex.” *Nature* 571 (7765): 361–65.

Stringer, Carsen, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. 2019. “Spontaneous Behaviors Drive Multidimensional, Brainwide Activity.” *Science* 364 (6437): 255.

Sun, Baochen, Jiashi Feng, and Kate Saenko. 2015. “Return of Frustratingly Easy Domain Adaptation.” *ArXiv [Cs.CV]*. arXiv. http://arxiv.org/abs/1511.05547.

Sun, Baochen, and Kate Saenko. 2016. “Deep CORAL: Correlation Alignment for Deep Domain Adaptation.” *ArXiv [Cs.CV]*. arXiv. http://arxiv.org/abs/1607.01719.

Sun, Zheng, Weiqing Xing, Wenjun Guo, Seungwook Kim, Hongze Li, Wenye Li, Jianru Wu, Yiwen Zhang, Bin Cheng, and Shenghui Cheng. 2020. “A Survey on Dimension Reduction Algorithms in Big Data Visualization.” In *Cloud Computing, Smart Grid and Innovative Frontiers in Telecommunications*, 375–95. Springer International Publishing.

Sussillo, David, Mark M. Churchland, Matthew T. Kaufman, and Krishna V. Shenoy. 2015. “A Neural Network That Finds a Naturalistic Solution for the Production of Muscle Activity.” *Nature Neuroscience* 18 (7): 1025–33.

Sussillo, David, Sergey D. Stavisky, Jonathan C. Kao, Stephen I. Ryu, and Krishna V. Shenoy. 2016. “Making Brain-Machine Interfaces Robust to Future Neural Variability.” *Nature Communications* 7 (December): 13749.

Taylor, P., J. Esnouf, and J. Hobby. 2002. “The Functional Impact of the Freehand System on Tetraplegic Hand Function. Clinical Results.” *Spinal Cord* 40 (11): 560–66.

Tenenbaum, J. B., V. de Silva, and J. C. Langford. 2000. “A Global Geometric Framework for Nonlinear Dimensionality Reduction.” *Science* 290 (5500): 2319–23.

Terrell, George R., and David W. Scott. 1992. “Variable Kernel Density Estimation.” *Annals of Statistics* 20 (3): 1236–65.

Thach, W. T. 1978. “Correlation of Neural Discharge with Pattern and Force of Muscular Activity, Joint Position, and Direction of Intended next Movement in Motor Cortex and Cerebellum.” *Journal of Neurophysiology* 41 (3): 654–76.

Thrasher, T. A., H. M. Flett, and M. R. Popovic. 2006. “Gait Training Regimen for Incomplete Spinal Cord Injury Using Functional Electrical Stimulation.” *Spinal Cord* 44 (6): 357–61.

Todorov, E., and Z. Ghahramani. 2004. “Analysis of the Synergies Underlying Complex Hand Manipulation.” *Conference Proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference* 2004: 4637–40.

Trautmann, Eric M., Sergey D. Stavisky, Subhaneil Lahiri, Katherine C. Ames, Matthew T. Kaufman, Daniel J. O’Shea, Saurabh Vyas, et al. 2019. “Accurate Estimation of Neural Population Dynamics without Spike Sorting.” *Neuron* 103 (2): 292-308.e4.

Triolo, R. J., C. Bieri, J. Uhlir, R. Kobetic, A. Scheiner, and E. B. Marsolais. 1996. “Implanted Functional Neuromuscular Stimulation Systems for Individuals with Cervical Spinal Cord Injuries: Clinical Case Reports.” *Archives of Physical Medicine and Rehabilitation* 77 (11): 1119–28.

Urai, Anne E., Brent Doiron, Andrew M. Leifer, and Anne K. Churchland. 2022. “Large-Scale Neural Recordings Call for New Insights to Link Brain and Behavior.” *Nature Neuroscience* 25 (1): 11–19.

Venugopalan, L., P. N. Taylor, J. E. Cobb, and I. D. Swain. 2015. “Upper Limb Functional Electrical Stimulation Devices and Their Man–Machine Interfaces.” *Journal of Medical Engineering & Technology* 39 (8): 471–79.

Vyas, Saurabh, Matthew D. Golub, David Sussillo, and Krishna V. Shenoy. 2020. “Computation Through Neural Population Dynamics.” *Annual Review of Neuroscience* 43 (July): 249–75.

Wärnberg, Emil, and Arvind Kumar. 2019. “Perturbing Low Dimensional Activity Manifolds in Spiking Neuronal Networks.” *PLoS Computational Biology* 15 (5): e1007074.

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.

Willett, Francis R., Daniel R. Young, Brian A. Murphy, William D. Memberg, Christine H. Blabe, Chethan Pandarinath, Sergey D. Stavisky, et al. 2019. “Principled BCI Decoder Design and Parameter Selection Using a Feedback Control Model.” *Scientific Reports* 9 (1): 8881.

Williams, Alex H., Tony Hyun Kim, Forea Wang, Saurabh Vyas, Stephen I. Ryu, Krishna V. Shenoy, Mark Schnitzer, Tamara G. Kolda, and Surya Ganguli. 2018. “Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis.” *Neuron* 98 (6): 1099-1115.e8.

Williamson, Ryan C., Benjamin R. Cowley, Ashok Litwin-Kumar, Brent Doiron, Adam Kohn, Matthew A. Smith, and Byron M. Yu. 2016. “Scaling Properties of Dimensionality Reduction for Neural Populations and Network Models.” *PLoS Computational Biology* 12 (12): e1005141.

Williamson, Ryan C., Brent Doiron, Matthew A. Smith, and Byron M. Yu. 2019. “Bridging Large-Scale Neuronal Recordings and Large-Scale Network Models Using Dimensionality Reduction.” *Current Opinion in Neurobiology* 55 (April): 40–47.

Willsey, Matthew S., Samuel R. Nason-Tomaszewski, Scott R. Ensel, Hisham Temmar, Matthew J. Mender, Joseph T. Costello, Parag G. Patil, and Cynthia A. Chestek. 2022. “Real-Time Brain-Machine Interface in Non-Human Primates Achieves High-Velocity Prosthetic Finger Movements Using a Shallow Feedforward Neural Network Decoder.” *Nature Communications* 13 (1): 6899.

Wodlinger, B., J. E. Downey, E. C. Tyler-Kabara, A. B. Schwartz, M. L. Boninger, and J. L. Collinger. 2014. “Ten-Dimensional Anthropomorphic Arm Control in a Human Brain-Machine Interface: Difficulties, Solutions, and Limitations.” *Journal of Neural Engineering* 12 (1): 016011.

———. 2015. “Ten-Dimensional Anthropomorphic Arm Control in a Human Brain-Machine Interface: Difficulties, Solutions, and Limitations.” *Journal of Neural Engineering* 12 (1): 016011.

Wolfe, Patrick J. 2013. “Making Sense of Big Data.” *Proceedings of the National Academy of Sciences of the United States of America*. National Acad Sciences.

World Health Organization. 2006. *Neurological Disorders: Public Health Challenges*. World Health Organization.

———. 2013. “Spinal Cord Injury.” World Health Organization - Fact Sheets. November 19, 2013. https://www.who.int/news-room/fact-sheets/detail/spinal-cord-injury.

Wu, Anqi, Nicholas A. Roy, Stephen Keeley, and Jonathan W. Pillow. 2017. “Gaussian Process Based Nonlinear Latent Structure Discovery in Multivariate Spike Train Data.” *Advances in Neural Information Processing Systems* 30 (December): 3496–3505.

Yamamoto, Wataru, and Rafael Yuste. 2020. “Whole-Body Imaging of Neural and Muscle Activity during Behavior in Hydra Vulgaris: Effect of Osmolarity on Contraction Bursts.” *ENeuro* 7 (4). https://doi.org/10.1523/ENEURO.0539-19.2020.

Yan, Yuke, James M. Goodman, Dalton D. Moore, Sara A. Solla, and Sliman J. Bensmaia. 2020. “Unexpected Complexity of Everyday Manual Behaviors.” *Nature Communications* 11 (1): 3564.

Yianilos, Peter N. 2000. “Locally Lifting the Curse of Dimensionality for Nearest Neighbor Search.” In *SODA*, 361–70. pnylab.com.

Yu, Byron M., John P. Cunningham, Gopal Santhanam, Stephen I. Ryu, Krishna V. Shenoy, and Maneesh Sahani. 2009. “Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity.” *Journal of Neurophysiology* 102 (1): 614–35.

Zhao, Yuan, and Il Memming Park. 2017. “Variational Latent Gaussian Process for Recovering Single-Trial Dynamics from Population Spike Trains.” *Neural Computation* 29 (5): 1293–1316.

Zhu, Jun-Yan, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. “Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks.” In *Proceedings of the IEEE International Conference on Computer Vision*, 2223–32. openaccess.thecvf.com.

Zimnik, Andrew J., and Mark M. Churchland. 2021. “Independent Generation of Sequence Elements by Motor Cortex.” *Nature Neuroscience* 24 (3): 412–24.