THE UNIVERSITY OF CHICAGO

# DECOMPOSING COMPLEX MOVEMENT USING DISCRETE DYNAMICAL POPULATION STATES

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES

AND THE PRITZKER SCHOOL OF MEDICINE

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL NEUROSCIENCE

BY

CALEB SCHEFFER SPONHEIM

CHICAGO, ILLINOIS

AUGUST 2023

Table of Contents

# List of Figures

# List of Tables

# Acknowledgements

To members of the Hatsopoulos Lab, past and present: Vassilis, Freddie, Taka, Kathikeyan, Alex, Carrie, Hide, Wei, Marina, Jeff, Paul, Courtney, Rebecca, Taylor, Ariana, Dalton, Rashi, Carmen, and Milan. Your support and input have been invaluable.

To the workers of the Animal Resource Center: Husbandry staff for keeping our animals happy and comfortable, vets and vet techs for keeping them healthy. Thank you for your dedication, facilitation, and input to our research.

To my dissertation committee, your advice and guidance have been instrumental. I want to specifically acknowledge Matt Kaufman for donating hours of his time to my math conundrums and scientific strategy.

To my nonhuman primate colleagues, those I have met or otherwise: Rockstar, Raju, Breaux, Kris, Theseus, Hermes.

To Nicho, for the steadfast advocacy and willingness to defend my time throughout my tenure in the Hatsopoulos lab. Thank you for your unyielding mentorship.

To my mother and father, two sources of academic and personal inspiration. To my brothers, Emmett and Aidan, whose perspectives, choices, and lived experiences will always make me grateful to know them. I love you all very much.

To my friends, past and present, whose positive influences, support, and warmth has buoyed me throughout this process. From long table-top role-playing sessions to twenty-mile runs, my life is leagues better having spent these years with y'all.

To my love, Anna; your light and purpose enriches my life. It excites me to no end to realize that we get to start the next chapter of our lives together, and that we get to bring our cats along for the ride.

# Abstract

The long history of research on primary motor cortex has led to a consensus that low-dimensional dynamics can characterize motor population activity, with categorically different activity regimens during distinct phases of movement. Here I build from previous work in the previous work to strengthen the claim that primary motor cortex uses distinct patterns of population neural activity which switch during movement execution. I replicate previous findings using parametric Hidden Markov Models and evolve our understanding of these population states by implementing recurrent switching linear dynamical systems models on data from nonhuman primates executing two-dimensional planar reaching movements. Neural population states mapped onto accelerative and decelerative directional components of motion, and linear dynamics of discrete states exhibit relationships with the kinematics they produce. These results further support the view that movement representations in M1 populations decompose movements into accelerative/decelerative directional elements instead of bell-shaped sub movements. Here I also present potential applications and evolutions of these findings in future experiments and contexts.

# 0: Introduction

The study of brain function has historically focused on the points at which the nervous system interacts with our environment. Scientists have conducted research to understand our bodies' physical interactions with the world for hundreds of years, and the ability for scientists to measure the relationship between brain and behavior in a sophisticated manner has increased in scale and scope substantially over the past fifty years. Much of this work is possible by collaborating with human participants; however, questions about cellular processes and neural circuits usually require invasive experiments involving animal models.

A large body of work has been conducted and continues to be pursued with nonhuman primates (NHPs), specifically Rhesus Macaques. There are a few reasons why Rhesus Macaques are still used heavily in motor neuroscience research: Macaques can be trained on complex tasks combining visuospatial and motor components that are difficult if not impossible for other species. They also possess opposable thumbs, and a neocortex evolutionarily proximate to our own. As with any model species, researchers have also designated Rhesus Macaques as a preferred model simply due to the large body of existing research with them. This long history of behavioral training has made macaques a preferred species to relate activity in primary motor cortex to movement and motor behaviors. Recent developments in markerless tracking and calcium imaging have increased the usage and popularity of rodent animal models for addressing questions in motor neuroscience (Mathis et al. 2018); these species are extremely nascent in their potential to provide deep insights, but the focus of the present dissertation is with experiments conducted in conjunction with rhesus macaques.

One of the main focuses of researchers has been the function of neocortex, specifically primary motor cortex (M1). Neuroscientists designate M1 as an extremely important cortical area for motor signals; it is an area which coordinates and structures motor signals, containing activity which correlates with kinematics, kinetics, and muscle activity, among other movement-related variables (Dum and Strick 1991). In primate species with anatomically distinct M1 areas (namely rhesus macaques), scientists have found M1 to be functionally related to motor learning and dexterous motor behavior (Apostolos P. Georgopoulos, Schwartz, and Kettner 1986; A P Georgopoulos et al. 1982; Schwartz 1994; Tanji and Evarts 1975; Jackson, Mavoori, and Fetz 2006; Cheney and Fetz 1980; Fetz and Finocchio 1971). M1, in the case of primates, can even send commands directly to the spinal cord via special brain cells called corticomotoneuronal cells (Rathelot and Strick 2009). Other commands travel out of M1 to other important subcortical areas of the brain, including thalamic regions, striatum, and cerebellum (Gibson, Houk, and Kohlerman 1985). Researchers study M1 via a myriad of technical approaches: lesion studies, acute inactivations, single neuron recordings, types of imaging, as well as chronic neural population recordings. By combining multiple single neuron recordings, researchers have been able to characterize the response properties of M1 neural populations (Apostolos P. Georgopoulos, Schwartz, and Kettner 1986; A P Georgopoulos et al. 1982).

As soon as they became available, researchers took advantage of advances in electronics fabrication and manufacturing to create dense, high-performance electrode arrays, simultaneously recording large populations of separable neurons (Sponheim et al. 2021; Nordhausen, Maynard, and Normann 1996). Chronically recording M1 during animal behavior allowed researchers to test additional hypotheses for the link between M1 activity and resulting behavior. Alongside innovation in array manufacturing, increased computational power allowed

for the practical use of machine learning to model neural activity, leading to powerful tools which infer the activity of an entire population based of the small subset of neurons to which scientists have access (Pandarinath et al. 2018).

## Dynamical Systems and Optimal Control Theory

There are two major theoretical approaches which inform the research presented in this dissertation: Optimal Control Theory and Dynamical Systems Theory. It should be noted that neither of these theories are mutually exclusive. Optimal Control Theory interprets the human motor control system by describing its impressive capabilities and explicating a control system that could accomplish such functionality (Scott and An 2004; Todorov and Jordan 2002). Arm movements are smooth, controlled, and can exhibit startling consistency in task-relevant settings. This highly performant system is also simultaneously noisy, approximate, subject to delay, and variable in many aspects. For primary motor cortex to be useful in coordinating motor commands, it must be capable of simultaneously taking in higher-level target and goal information and sending commands, all while incorporating constant sensory feedback.

Todorov and Jordan posited that for a system to overcome the stochastic nature of a squishy biological body, it must allow variability in a fashion, and utilize feedback to correct for errors (Scott and An 2004; Todorov and Jordan 2002). As a result, their hypothetical control system only corrects errors that negatively impact performance on a goal-directed task, whereas variability in other dimensions or domains of movement are not. For example, if a finger reaches the “send” button of a text message, it does not matter the trajectory it took to get there, variable, or otherwise. Todorov and Jordan also propose that a set of control rules could exist for different behavioral contexts, prioritizing certain open loop responses. An evolution of this optimal

control theory is one in which command structures based on feedback are not adjusted or changed constantly – rather, motor commands and mapping would change intermittently, to allow for feedback and error correction to be more precise and predictable.

Intermittent Control Theory is a subset of Optimal Control Theory and supposes that certain behaviors humans exert would benefit from a system that observes constantly, yet acts intermittently, such as reaching or eye movements (Gawthrop et al. 2011; Karniel 2013; Leib, d'Avella, and Nisky 2017). By switching intermittently between certain patterns of activity based on feedback, a system like primary motor cortex would be more energy efficient and resilient to noise. While this framework for an intermittent control system has not been adopted by systems neuroscientists, in particular experimentalists, it serves as a potential explanation for the results found by work published in 2019 by Naama Kadmon Harpaz (Kadmon Harpaz et al. 2019). Additionally, this idea of intermittent changes in control signals may align with how certain movements are composed already.

Scientists have found utility in modeling neural activity using dynamical systems approaches. In the context of M1 population activity, significant evidence has been established that the current state of a neural system, plus some noise, can determine the future state of that system, while correlating strongly with produced movement (Shenoy, Sahani, and Churchland 2013). A key observation which has benefitted this interpretation is that neural activity often inhabits a lower-dimensional subspace, plane, or manifold within a higher-dimensional space of all possible neural activity (Vyas et al. 2020); using dimensionality reduction techniques, the state of a neural dynamical system can be more readily described (Yu et al. 2009).

## Movement Compositionality

The concept of movement compositionality posits that movements, particularly complex movements such as speaking and writing, are composed of a series of movement elements. Behavioral evidence for movement compositionality has been shown at the kinematic and dynamics levels (Doeringer and Hogan 1998; D'Avella, Saltiel, and Bizzi 2003; Hatsopoulos, Xu, and Amit 2007). The idea of a reusable set of movement elements combined in diverse ways provides a computationally efficient method for generating a rich variety of different movements, much like the productive capacity of language in expressing an infinite number of ideas with a limited set of words. If this concept of compositionality forms the foundation for the neural control of movement, one should expect to find evidence for discrete neural states associated with movement elements. The neural activity responsible for the preparation of a movement has consistently been identified as separable and different from primary motor activity during movement itself (Churchland and Shenoy 2007; Kaufman et al. 2016; Elsayed et al. 2016). M1 also transitions between a passive and active state, or a preparation and movement state, during a given behavioral task, with the transition between these two sets of dynamics initiated by a large condition-invariant signal (Kaufman et al. 2014). The Shenoy group even used a Hidden Markov Model (HMM) to define a discrete transition between a preparatory and execution state (Kemere et al. 2008).

## Population States in Primary Motor Cortex

Researchers have identified that M1 displays distinctive low dimensional linear dynamics during the execution of reaching movements; it also displays activity in linearly separable neural

subspaces when an animal prepares to move and initiates a movement (Elsayed et al. 2016). Taken together, primary motor cortical activity is often designated as a nonlinear dynamical system to support these different dynamics (Vyas et al. 2020). Unfortunately, modeling nonlinear dynamical systems is an extremely difficult and often intractable challenge; as such, researchers have developed approaches which approximate the hypothetical nonlinear dynamical system primary motor cortex may use to generate movement, without having to solve a complicated and hairy optimization problem (Sabatini and Kaufman 2021; Linderman et al. 2017; Glaser et al. 2020; E. Fox et al. 2008; Pandarinath et al. 2018). I present an application of one such solution in Chapter 1 and discuss other approaches in the conclusion of this dissertation.

While differential equations and linear algebra were brought to bear on population activity, an important body of work has focused on identifying the fundamental building blocks of movement, often called movement primitives. Classical perspectives on movement primitives posit that full sub-movements, those being a reach whose directionless speed profile is bell-shaped, remains the coherent unit of movement coding for which primary motor cortex codes (A P Georgopoulos et al. 1982; Morasso 1981).

Recent research has built upon this previous work by identifying neural patterns which imply a finer-grain decomposition, a small unit of basic building block of movement (Kadmon Harpaz et al. 2019). This work has shown consistent changes in population activity corresponding to speed extrema as opposed to speed minima. Without using kinematic data as an input, researchers used an HMM and found distinct population states under two different planar reaching paradigms, each of which corresponded to accelerative or decelerative segments moving in a particular direction. Moreover, using simple linear models, another study showed that a minority of neurons in motor cortex exhibit a unique encoding of movement during a

double step reaching paradigm (Dickey, Amit, and Hatsopoulos 2013); the neurons aborted their activity when the target moved positions, and seemingly restarted its activity as if beginning a new reaching movement; the remainder of the neural population continued coding for the end of the first movement, however. In contrast, there exists evidence that primary motor cortex simultaneously codes for rapid sequences of reaches, not building movement coding using fused elements (Zimnik and Churchland 2021). Evidence has been provided for mid-movement state transitions (Suway et al. 2018a), showing consistent stability in preferred directions during a three-dimensional center-out reaching task. Specifically, researchers found three epochs of stable PDs, between which populations discretely transitioned. Researchers have also showed that M1 activity distinctly shifts consistently in the middle of reaching, potentially due to a change in population activity state (Takei et al. 2018).

A key component of deciphering the relationship between brain and behavior is identifying the structure and nature of the brain activity itself. In M1, conversations about the main components of activity, the building blocks which are utilized to generate movement, have been in contention for decades. This dissertation's goal is to further investigate the possibility that M1 population activity could be approximated with different blocks of low-dimensional patterns, particularly during movement.

# 1: Discrete Neural Population States Decompose Movement with Different Switching Linear Dynamics

## Introduction

Prominent theories of motor control agree that populations of neurons in the primary motor cortex (M1) work in tandem to transmit commands via the midbrain and spinal cord to muscle groups, generating movement (Omrani et al. 2017). Research from a variety of scientists describe these coordinated signals as a dynamical system - that is, population activity that dynamically evolves over time and whose current state can predict its future state, with noise (Churchland et al. 2012; Vyas et al. 2020; Shenoy, Sahani, and Churchland 2013). Additionally, studies have identified different and separable dynamics in neural populations depending on the behavioral state of an animal (Elsayed et al. 2016). When an animal prepares to move, M1 population activity displays distinct regimens of activity, which significantly shift when the animal begins to execute a movement (Kaufman et al. 2016).These changes in coordinated activity patterns have been identified using various analyses and have historically decomposed neural population "states" into preparatory and execution-related epochs. These movements often correspond with a set of rotational dynamics identified in low-dimensional latent space (Churchland et al. 2012). The ballistic reaches often observed in center-out reaching tasks can also be described as a "submovement" with an associated bell-shaped speed profile, i.e., an acceleration away from the center, paired with a deceleration as the hand approaches the target.

Recent work suggests that neural population activity and coding significantly changes during movement itself and the bell-shaped speed profiles associated with a simple point-to-point

reaching movement could be separated into multiple neural states. Suway and colleagues (Suway et al. 2018b) demonstrated that the preferred directions of individual neurons quickly change between two periods of stability roughly associated with accelerative and decelerative components of point-to-point reaching movements. Moreover, Hidden Markov Models (HMMs) trained solely on neural activity in motor and parietal cortices have demonstrated that population activity reliably decomposes reaching kinematics into accelerative and decelerative components of movement (Kadmon Harpaz et al. 2019).

Hidden Markov Models are Markovian processes which predict the probability that a system's state is one of a set of possible states that cannot be directly observed. The hidden states of the system are inferred from observed emissions. The fundamental assumption of applying HMMs to modeling neural activity is that we can infer the overall hidden population state from the noisy signal of neural activity.

However, while an HMM provides behaviorally meaningful decompositions of M1 population activity, it has distinct limitations. In particular, the emission probabilities of a given hidden state are static - they cannot account for neural dynamics that occur within the state. Without extending the HMM framework, these models fail to describe whether and how the neural population follows predictable laws of motion within a given hidden state. As many recent studies have demonstrated, neural dynamics are important to incorporate into our understanding of discrete population states and their decomposition of resultant kinematics (Sabatini and Kaufman 2021; Musall et al. 2019; Omrani et al. 2017). Extending past patterns of kinematic decomposition using models that more faithfully capture the temporal evolution of population activity would bolster the case for mid-movement neural population states.

In order to incorporate neural dynamics while accounting for the possible discrete nature of population states in M1, we modeled population activity as a recurrent switching dynamical system (rSLDS); (Linderman et al. 2017). This class of generative models decomposes population activity into a set of discrete population states, each of which is modeled as a unique linear dynamical system, thereby approximating the underlying nonlinear dynamics as a piecewise set of linear dynamical systems. These models not only allowed us to decompose kinematics according to the most likely inferred population state at any given point in time, but they also enabled us to model distinct linear dynamics associated with each discrete state. Using rSLDS to model neural activity during continuous, sequential reaching movements performed by nonhuman primates, we show here that neural state transitions occur at speed extrema such that states map to accelerative and decelerative movement segments in particular motion directions, as has been shown with simpler models. More importantly, neural states associated with accelerative and decelerative movement segments possess qualitatively different characteristics in their neural dynamics.

## Methods

### Subjects and array implantation

Experiments were conducted with three male rhesus macaques ("RJ", "RS", and "BX"). Multi-electrode Utah arrays were implanted in the arm/hand area of primary motor cortex of each subject. For RJ and RS, one 96-channel array (1.0 mm long electrodes) was implanted whereas for BX, two 64-channel arrays (1.5 mm long electrodes) were implanted parallel to the central sulcus on the precentral gyrus. Using intracortical microstimulation on the arrays in BX,

the medial array evoked responses in the proximal arm and wrist, whereas the lateral array evoked responses in the wrist and hand. All surgical and behavioral procedures were approved by the University of Chicago Institutional Animal Care and Use Committee and conform to the principles outlined in the *Guide for the Care and Use of Laboratory Animals*.

**Random Target Pursuit Task**

Subjects performed a Random Target Pursuit (RTP) task, which presented the animal with a sequence of seven pseudo randomly placed targets, one at a time (Figure 1A) (Kadmon Harpaz et al. 2019; Hatsopoulos, Joshi, and O'Leary 2004). Subjects moved a cursor to the targets on a screen by controlling the position of a handle in the horizontal plane which was placed at the end of a two-link exoskeletal robot (Kinarm Kingston, Ontario) on which their arm rested. Once the cursor controlled by the animal intersected with a target, it would disappear and the next target in the sequence would appear. After seven target intersections, the trial was marked successful and the animal was rewarded. We only included and analyzed successful trials.

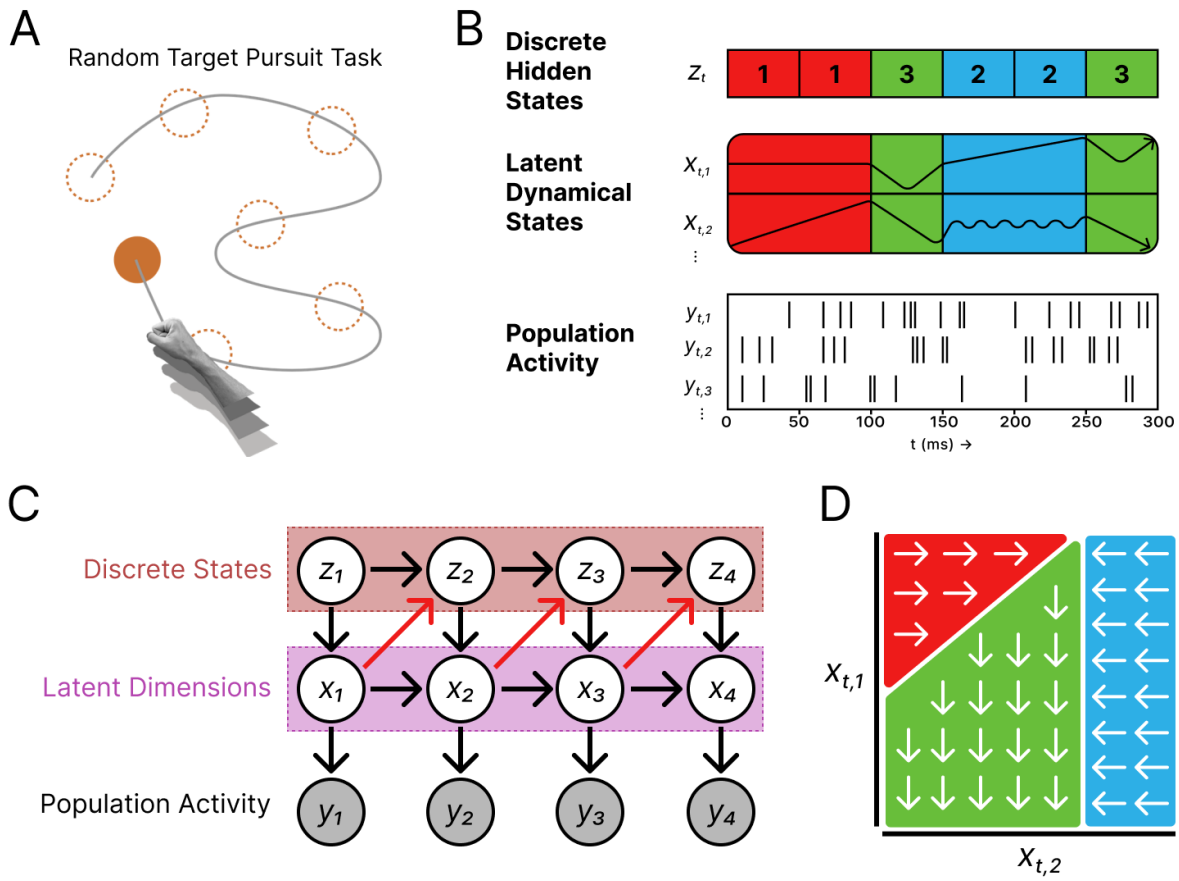


**Figure 1. Random Target Pursuit Task, rSLDS overview, and state-space visualization.** A. Subjects were presented with seven targets sequentially at random locations, reward was delivered upon intersection with the final target. B. Neural population activity was summed across 50ms bins, and modeled using a recurrent switching linear dynamical system to identify and describe common activity patterns using multiple linear dynamical systems. C. Graphical model of the recurrent switching linear dynamical system. Recurrent connections are labeled in red. D. Hypothetical example of a two-dimensional, three-state model flow field; visualizing the piecewise linear approximation of nonlinear dynamics.

**Data Collection**

Detailed explanations of data collection and handling can be found in previous work from our research group (Hatsopoulos, Xu, and Amit 2007). The exoskeleton used by the subjects tracked their shoulder, elbow, and hand position, translating their movements to a cursor on a screen in their line of sight (BKIN Technologies). X and Y hand positions were gathered using

forward kinematics calculations from motors attached to the shoulder and elbow joints of the exoskeleton. RS and RJ kinematics were sampled at 500Hz; BX kinematics were collected at 2kHz.

Neural data was gathered at 30kHz with Blackrock Utah microelectrode arrays implanted in Primary Motor Cortex. Signals from individual separable neurons were sorted offline, with the number of units recorded ranging from 70 units for subject RJ, 100 units for subject RS and 141 units for subject BX.

**Data Processing**

Data from RJ and RS were spike sorted using Offline Sorter ("Offline Sorter" 2017); data from BX was spike sorted using Kilosort2 and manually curated using Phy (Pachitariu, Sridhar, and Stringer 2023). Spiketimes detected from spike-sorted waveforms were collected into bins with the width of 50 ms. Neural and kinematic data were included in trials defined as a window of time from the initial target intersection to the last target intersection.

Kinematics were resampled to 1kHz for all subjects and bidirectionally low pass filtered at 6hz using a 6th order Butterworth filter. Directionless speed and acceleration were calculated from x and y velocity data. Based on the discrete state transition times derived from the appropriate rSLDS model, we broke up each trial's kinematics into snippets of movement according to each discrete state visit over 100ms. We collated these snippets according to each discrete state identity, across trials. We calculated kinematic metrics to characterize each discrete state's kinematic decomposition, including mean snippet duration (as seen in Figure 4B), normalized speed profiles (Figure 2A), and state snippet direction (Figure 2A).

**rSLDS Implementation**

The primary goals of the present work is to verify the reliability of previous results found by Kadmon Harpaz (2019) using more sophisticated modeling approaches. An essential part of that is to verify those results using the same analytical approach. We began by replicating the results from previous work using the same models and different datasets. We utilized the tabular nonparametric hidden Markov model described in Kadmon Harpaz (2019) and verified similar results. This type of model, instead of designating a particular emission model for some distribution of emission values (such as a Poisson emission model, which designates a probability of spiking), fits a parameter to each unique spike count integer present in the dataset across the neural population for each discrete state. As with the models we present in our main results, the number of discrete states for these tabular models were provided *a priori*; unlike our main results, the optimal number of states were determined using minimum Akaike Information Criterion (AIC), calculated separately for each subject's dataset. Specifically, we were able to verify that the discrete state timing modeled using a nonparametric tabular HMM demonstrated consistent state transitions at speed extrema, and each discrete state showed accelerative or decelerative trends in speed.

The main results presented here are based on state space models developed by Scott Linderman's research group under the umbrella term "ssm" (State Space Models). We present results mainly from the recurrent Switching Linear Dynamical Systems (rSLDS) class of models included in ssm.

Linear Dynamical Systems (LDS) are commonly used to characterize neural population activity (Shenoy, Sahani, and Churchland 2013). In the context of summarizing and describing

brain activity, LDS models often identify latent neural subspaces/manifolds within high-dimensional neural activity and use a linear set of differential equations to describe the low-dimensional latent activity. Switching Linear Dynamical Systems incorporate Markovian discrete states and low-dimensional latent continuous states, whose probabilities and transitions are modeled with a Hidden Markov Model; these switching models allow each discrete state to be associated with a separate set of low-dimensional linear differential equations (Linderman et al. 2017). By identifying areas of latent space where specific population states have the highest probability, SLDS models break up low-dimensional latent space into "pie slices" of separate, individual linear dynamical systems with distinct linear dynamics. Recurrent SLDS models (rSLDS) extend this approach by allowing the linear dynamics of each discrete state to influence the probabilities of discrete state transitions (Equation 2). Put another way, rSLDS models are allowed to learn the dependence between a continuous state in a certain location in state space, and discrete state switch timing; if a latent neural trajectory maintains its presence in a particular area of state space and exhibits similar dynamics, it will resist switching discrete states due to that behavior, even if discrete state Markovian transition probability to switch is high.

rSLDS models characterize population dynamics using latent variables with lower dimensionality than the entire neural population. A Generalized Linear Model (GLM) is established and fit to describe a linear transformation between high-dimensional neural activity and low-dimensional latent space and is trained alongside continuous and discrete states using expectation maximization and a structured mean field variational posterior. The details of optimization procedures can be found in Linderman et al., (2017) or Glaser et al,. (2020) (Linderman et al. 2017; Glaser et al. 2020).

The rSLDS models we used here assumed diagonal gaussian dynamics (the diagonal and gaussian characteristics only extend to the noise term included in fitting the dynamics) with a Poisson emission distribution of spike times. All models for each subject were trained on available neural data from 80% of behavioral trials and tested using the remaining 20% of trials. In addition, during cross-validation, models were trained on 75% of the recorded neurons, with 25% held out to solely test on.

The low-dimensional latent trajectory of a given trial is determined by the dynamics A, a bias term b, and a gaussian noise covariance term Q at each timestep *t*.

$$(1)\ x_t \approx A^{(z_t)}x_{t-1} \quad + b^{(z_t)} + Q^{(z_t)}$$

Discrete state switching is determined using a Markovian process; the most probable discrete state at any point in time and at any given point in low-dimensional latent space is determined by a combination of markov state probabilities R and recurrent weights x from the previous timestep:

$$(2)\ z_t \sim Cat(\pi_t), \pi_t \ = \ softmax(R_{z_{t-1}}x_{t-1} + r_{z_{t-1}}),$$

Discrete state timing was estimated for each test trial, and mapped to kinematic data with a 100ms lag to account for the delay between the motor cortical activity and limb movement (Schwartz 1994; Paninski et al. 2004). Kinematic data was segmented into separate snippets according to neural state timing. State snippets were collected and analyzed for various characteristics and compared with the corresponding discrete states' linear dynamical systems.

**Cross Validation Procedures**

rSLDS models utilize advanced optimization procedures to identify local minima solutions, but they require a small number of hyperparameters to be specified *a priori*. For our

purposes, the most important hyperparameters to identify were the optimal number of discrete states and the dimensionality of the resulting latent state model where the continuous states and linear dynamics for each discrete state would be defined. We performed a grid search across a range of potential discrete states and latent dimensions (2-40 discrete states, 2-80 latent dimensions), calculating the mean of the log-likelihood of held-out trials across five separate folds of test and train data (using the same 80%, 20% test split). We implemented neural population co-smoothing by additionally holding out 25% of the neural population during training and calculated the log-likelihood of the neural activity of only those neurons in held-out test trials given the trained model. Stated another way, when calculating the probability of test data, models were trained on 80% of available trials; for those 80% of trials, 75% of the available neural population were used for training. When testing these trained models, we calculated the likelihood that the held out 20% of trials, and the 25% of neurons within those trials, could be produced given the trained models. This method of training latent variables models on a portion of a neural population is referred to as Co-Smoothing. Co-smoothing allows for a more accurate assessment of cross-validated model performance to infer neural population activity (Pei et al. 2021).

**Dynamics Analysis**

In addition to providing a more sophisticated analysis of population activity, we used rSLDS models to investigate any connection between the linear dynamics of each discrete state and the kinematics they generate. We focused our analysis on the speed of latent trajectories of all trials in our dataset, as well as the eigenvalues and eigenvectors of the linear dynamics matrices $A$ for each discrete state for each subject. Additionally, we assessed the dimensions

accounting for the highest portion of variance for each discrete state by projecting latent values from each state onto their respective eigenvectors. We selected a subset of dimensions for each discrete state based on their highest variability up to 90% cumulative variance. The results presented in Figure 4 were calculated using that subset of dimensions for each subject and discrete state. A key component of our analyses included examining the low-dimensional latent trajectories of individual trials. We calculated the speed of these trajectories as they evolved over time through different discrete states (and therefore different sets of linear dynamics) and compared those collections of trajectories between different classes of states.

## Results

### rSLDS model fitting

We analyzed population activity from M1 in three non-human primates who were trained to perform a random-target pursuit (RTP) task in the horizontal plane. The RTP task required the animal to reach continuously to a sequence of randomly positioned targets, as opposed to single ballistic movements often seen in center-out reaching tasks (Figure 1A). We used rSLDS models to analyze the latent dynamical structure of neural activity during each animal's reaching behavior, decomposing kinematics into discrete states solely based on neural population activity (Figure 1B-D).

Utilizing a cross-validated and co-smoothed grid search from 2-40 discrete states and 2-80 latent dimensions, we were able to identify the following optimal states and dimensions for each subject's datasets. We used 10 states and 25 dimensions for subject RS, 14 states and 22 dimensions for subject RJ, and 10 states and 30 dimensions for subject BX. Individual cross-

validated performance plots of log likelihood grid searches can be found in supplementary materials (Figure S2).

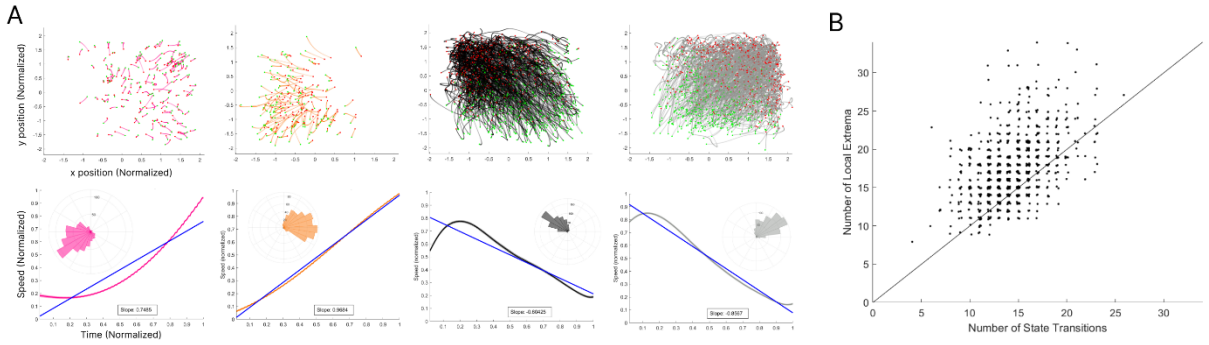


**Figure 2. Decomposition of kinematics using rSLDS**. A. a subset of discrete states, showing the kinematic segments and the normalized mean speed profiles of the segments corresponding to each state for subject RS. The inset polar histogram shows the directionality of each state's segments. B. The relationship between the number of state transitions and speed extrema across trials. Each point corresponds to a single RTP trial.

## Kinematic Decomposition

We assessed whether the transitions in discrete states (and therefore linear dynamics) occurred at relevant and consistent points in behavior. The rSLDS model identified states that decomposed continuous, sequential reaching movements into segments moving directions that were either accelerative or decelerative (Figure 2A). To determine whether a state was classified as accelerative or decelerative, we computed a linear regression to the normalized speed profile of each states' movement snippets; if the slope of the linear fit was higher than 0.2, it was classified as accelerative. If it was below -0.2, it was classified as decelerative.

To evaluate whether neural transitions occurred close to speed extrema, we first computed the correlation between the number of speed extrema in each trial and the number of discrete states transitions in each trial. We calculated the threshold for speed extrema detection by equating the number of speed extrema to the number of discrete state transitions across the

entire dataset, not per trial. the number of speed extrema per trial were significantly correlated with the number of state transitions ($r=0.46$, $p<0.0001$ for RS, $r=0.16$, $p<.01$ for Bx. and $r=0.73$, $p<0.0001$ for RJ) (Figure 2C; Figure S2D-F).

For each instance of a speed extrema, we also calculated metrics of Precision and Recall utilized in Kadmon Harpaz (2019) and compared Precision and Recall scores to scored calculated using a null model. In our analysis, a "true positive" was registered if a speed extremum occurred within 50ms of a neural transition, a "false positive" if a state transition occurred without a speed extremum nearby, and a "false negative" if a speed extremum was detected without a state transition nearby. Precision was defined as:

$$\frac{\#\ True\ Positives}{\#\ True\ Positives + \#\ False\ Positives}$$

Recall was defined as:

$$\frac{\#\ True\ Positives}{\#\ True\ Positives + \#\ False\ Negatives}$$

These metrics were calculated using 2000ms windows, across all trials concatenated together, within each subject separately. Our null model was calculated by randomly shuffling the order of the windows of neural transition timepoints, such that the state transition windows were no longer matched with their original location in the dataset. The best outcome of prevision and recall would be a value of one; we found that for all three subjects' data, precision and recall were significantly greater for the real state transitions and kinematics than the null shuffled model ($p < .001$) (Figure 2B; Figure S2A-C).

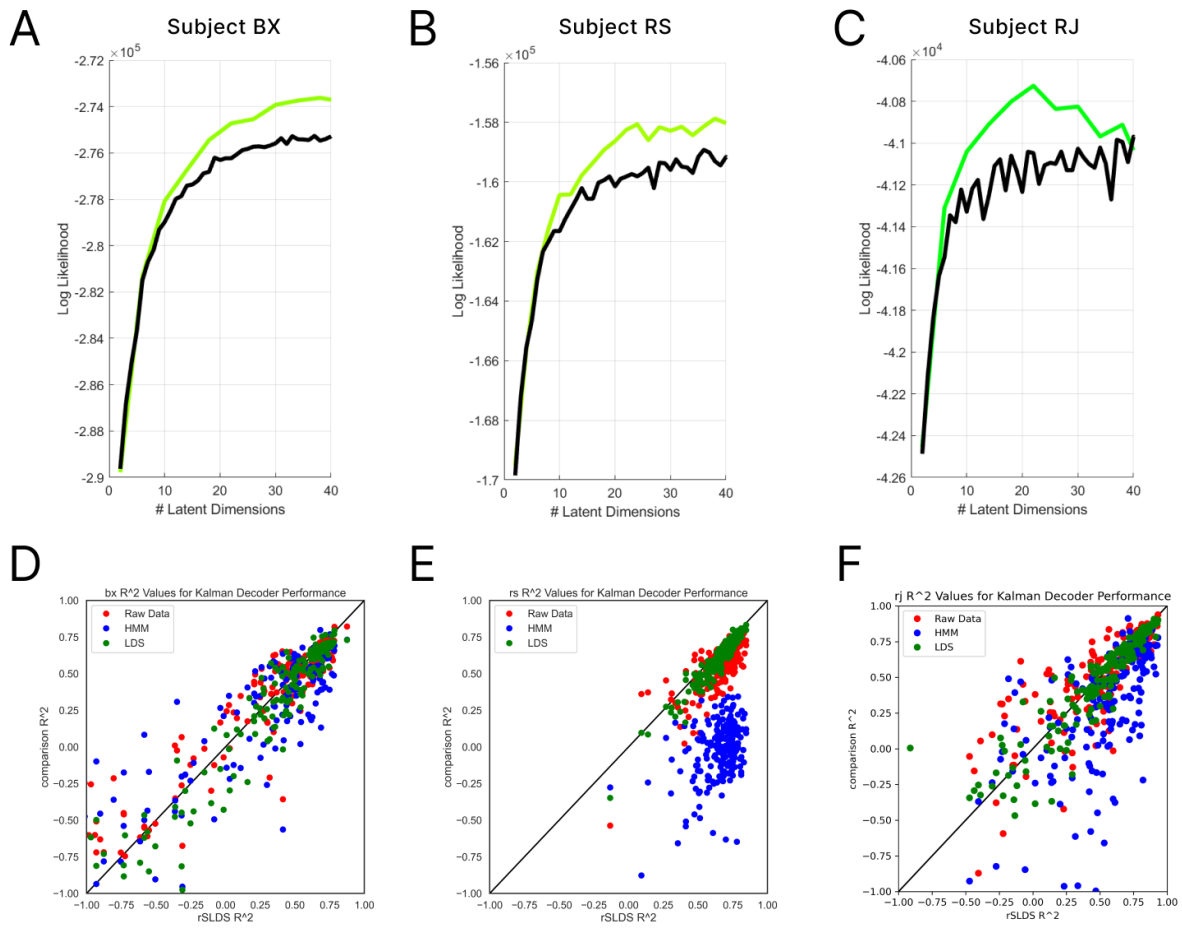


**Figure 3. Decoding and cross-validated performance results**. A,B,C. Cross-validated and co-smoothed log-likelihood, for a linear dynamical system (black), versus the optimal discrete state rSLDS model (green). Across all three subjects, rSLDS models had a higher probability of held-out test data then LDS models. D,E,F. Treating latent variables from each model (save raw neural data) as the inputs, linear Kalman filters were used to predict hand velocity on portions of held-out test data. rSLDS model $R^2$ is plotted on the x-axis, compared with one of three alternative models on the y-axis. Dots plotted below the diagonal indicate superior rSLDS model performance.

**Comparison with other models**

Next, we compared rSLDS with other simpler models in terms of capturing the spatio-temporal structure of M1 population activity, and of decoding movement of the hand. These other models included a single, non-switching linear dynamical system (LDS), as well as a Poisson parametric Hidden Markov Model (HMM) that does not consider dynamics. Using a

parameter sweep (through a range of latent dimensions for latent variable models, and a range of discrete states for HMM models), we found the optimal number of discrete states for the rSLDS and HMM models using the log likelihood of held-out trials / neurons from cross-validated and co-smoothed data. We then systematically varied the number of latent dimensions for the rSLDS and LDS models and compared log likelihood of held-out test data (Figure 3A-C). The log likelihood for the optimal rSLDS models were higher than that of the two other two models for all three subjects, indicating the recurrent switching model more accurately captured the structure of the data. Hidden Markov Models do not employ low-dimensional neural trajectories, so we were unable to include cross-validated performance metrics for these classes of models in Figures 3A-C.

We then examined whether the latent trajectory values from the rSLDS model could more accurately decode hand velocity than latent trajectory values from other models. We used a Kalman filter (KordingLab python decoding package), treating latent dimension values from the LDS and rSLDS model as the inputs, with x- and y- components of hand velocity as the predicted outputs. Decoders were trained on ninety-nine percent of available data, with one percent held out to evaluate predictive fit using the fraction of variance accounted for metric ($R^2$). Test performance was evaluated across 100 folds of data. Three different models were chosen to benchmark rSLDS latent value decoding performance: the LDS latent dimension value, the parametric HMM's mean spike count parameters from the emission probabilities, and the raw high-dimensional neural data (evaluated in 50ms bins). In the HMMs' case, in lieu of low-dimensional latent variables as decoder inputs, we used the Poisson mean spike count parameters for each neuron associated with the discrete state, at each time point for which the given state was most probable; for example, if an HMM were to predict state 4 as most probable,

the input values for the decoder would be the mean spike count parameter for each neuron associated with state 4, for as long as that state is most probable and active, switching to a different set of mean spike count parameters if another state became more probable.

We compared $R^2$ across equivalent cross-validated test folds of data for these four types of models, including rSLDS (Figure 3). Across all three subjects, rSLDS models decoded kinematics as well or better than other models; using a Wilcoxon rank-sum test, across all five cross-validation folds of data, rSLDS performed significantly better than the other tested inputs ($p < 0.01$, Wilcoxon rank-sum).

**State Dynamics**

Each discrete state modeled by rSLDS is associated with its own unique linear dynamical system, which can be characterized using the eigenvalues and eigenvectors of the dynamics matrix of each state. We found that the real eigenvalues of each state, for each subject, ranged from 0 to 1, indicating convergence and decay to their respective fixed points. We also verified that the corresponding imaginary eigenvalues were all larger than zero, indicating rotation of the states' linear dynamics around their given fixed points (Figure 4A). We found that discrete states with short-duration kinematic snippets, meaning that certain discrete states had smaller timelengths of continuous state visits, displayed lower real eigenvalue magnitudes (Figure 4B). This trend demonstrates that fast-decaying dynamics tend to be associated with short state visits, and therefore shorter movement snippet in kinematics. This trend was shown across all three subjects, and implied that short movements may be driven by fast-decaying state dynamics.

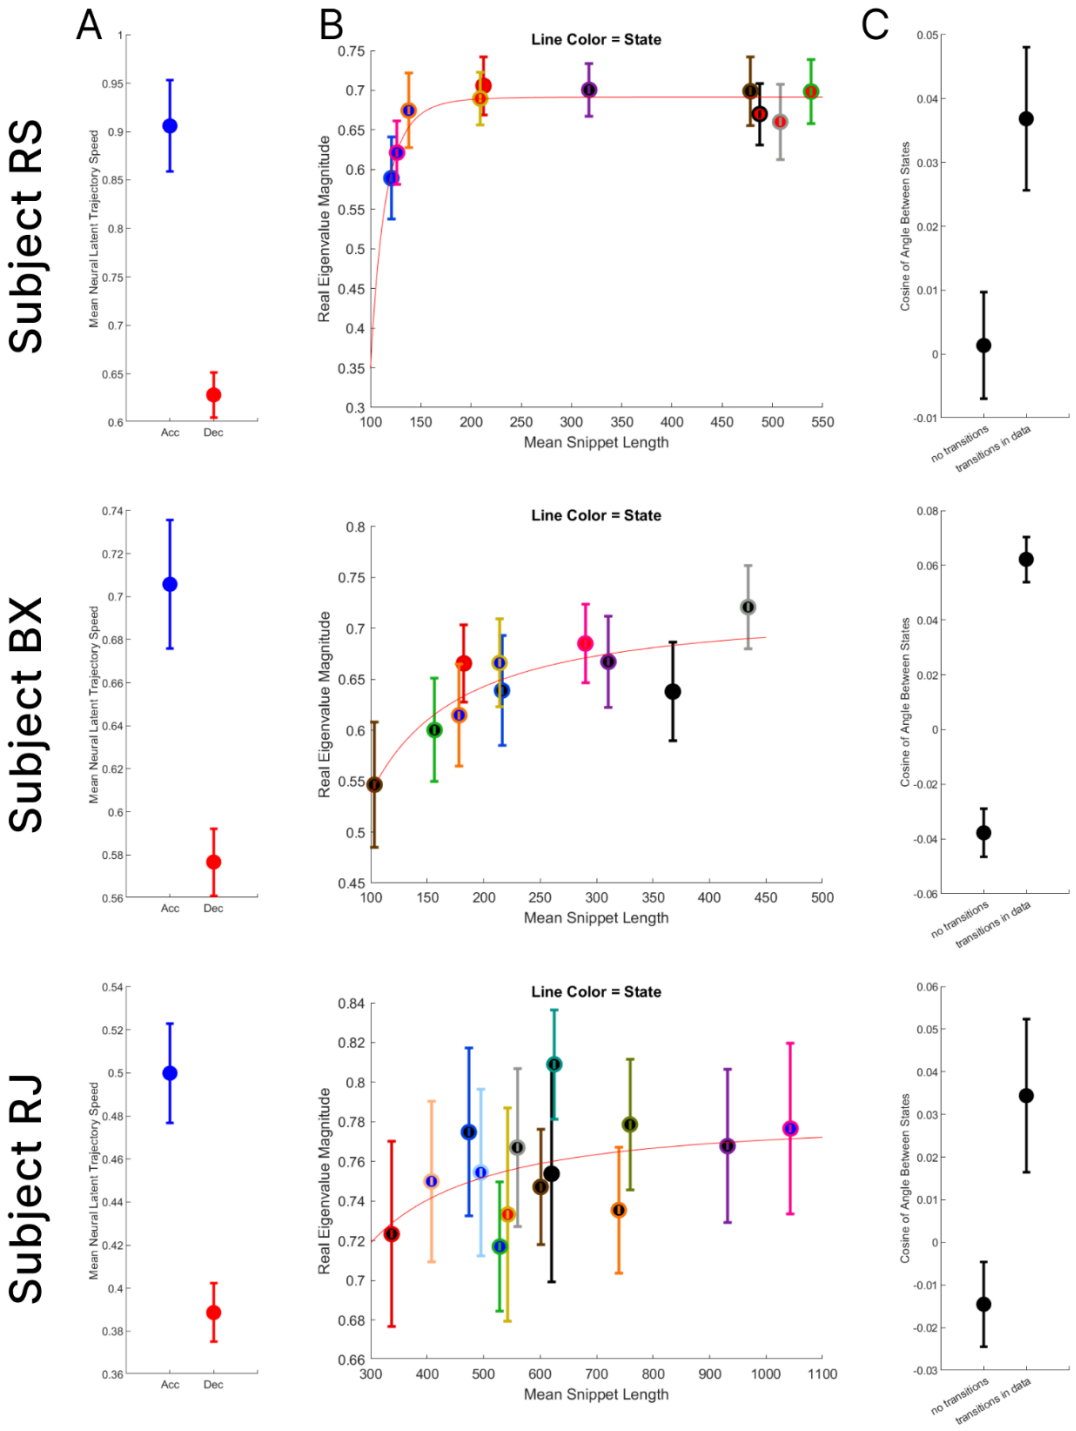


**Figure 4. Differences in dynamics of discrete states**; top row shows results from subject RS, middle row subject BX, and bottom row subject RJ. A. speed of neural trajectories belonging to accelerative and decelerative eigenvalues. B. Comparison of individual state real eigenvalue magnitude and average snippet length C. Cosine of the angle of state eigenvectors which transition between one another, compared to angles between all other states.

States were classified as “accelerative” based on the slope of mean normalized speed profiles of respective kinematic segments. Accelerative states had significantly lower real eigenvalues ($p < 0.05$, test) and demonstrated higher (though not significantly so) imaginative components of complex eigenvalues than decelerative states. We also compared the directionless “speed” of neural trajectories (calculated using the derivative of the latent value across dimensions) in accelerative versus decelerative states and found significant differences (Figure 4A). Neural trajectories in accelerative states were over 25% faster than those in decelerative states. All three subjects displayed faster trajectories in accelerative states than decelerative states.

**Neighboring States Similarities**

While rSLDS models approximate a lower-dimensional nonlinear dynamical system using multiple discrete piecewise linear dynamical systems, the neural trajectories which travel through latent space and across discrete state boundaries are smooth and continuous. We wanted to find out whether the direction of different discrete state dynamics was more similar for transitions between adjacent states for which we observed neural trajectories from our data. We found that the directions of dynamics as measured by the eigenvectors of adjacent states, whose boundaries supported neural trajectories observed in data, were more similar to one another than states where transitions never occurred (Wilcoxon Rank Sum Test, $p < .05$)(Figure 4C). This result implies that transitions represent a change in dynamics strength and direction but a continuity of neural trajectory; these discrete states do not require neural trajectories to support discontinuities in low dimensional latent space.

## Discussion

By incorporating neural dynamics in a state-based population model, we have provided further evidence that population activity in M1 transitions across discrete states, each of which maps to an accelerative or decelerative directional segment of upper-limb movement. Each state is characterized by a unique linear dynamical system whose properties differ qualitatively between associated accelerative and decelerative segments. These results suggest alterations in neural population coding over the course of a continuous movement.

Why distinct neural population states decompose continuous movements into accelerative and decelerative movement segments is unclear. One possibility is that these different segments are associated with distinct agonist/antagonist muscle groups. However, our analysis identified distinct population states for accelerations in one direction and decelerations in the opposite direction which presumably recruit similar muscle groups (See orange and gray states in Figure 2A). Another possibility is that tuning properties of individual neurons become unstable across the population at certain points in a reaching movement thus resulting in population state change. In particular, one particularly relevant study indicated that directional tuning across a neural population remains stable for brief periods of time punctuated by instabilities at one or two points during a point-to-point reaching task (Suway et al. 2018b). The first of these instabilities occurs after movement onset at ~100 ms before peak speed which, given a ~100 ms lag between M1 and limb movement (Schwartz 1994; Paninski et al. 2004), may correspond to a population state transition we observed from acceleration to deceleration. The second of these instabilities is more prominent in reaching movement that require stopping at the target which does not apply in the random-target pursuit task where movements move through the targets without stopping.

Mid-movement neural decomposition implies a potentially different view of what constitutes a movement primitive at least at the level of primary motor cortex. Classical perspectives on movement primitives propose that submovements with bell-shaped speed profiles (i.e. a combination of acceleration and deceleration) constitute the kinematic units of reaching movements (A P Georgopoulos et al. 1982; Morasso 1981). Our results suggest that population cortical activity decomposes movements at a finer grain and forms neural building blocks that parse movements at speed extrema as opposed to whole submovements.

We observed linear rotational dynamics within each discrete state, in line with past studies identifying rotational dynamics associated with movement (Churchland et al. 2012; Sabatini and Kaufman 2021; Musall et al. 2019). However, given that population dynamics in motor cortex is almost certainly non-linear (Sabatini and Kaufman 2021), we argue here that an accurate description of population dynamics can be obtained by approximating the non-linearity as a piecewise linear system Of the available sophisticated generative models that combine discrete state-based structure with neural dynamics (E. B. Fox et al. 2010; E. Fox et al. 2008; Glaser et al. 2020; Zoltowski, Pillow, and Linderman 2020), we moved ahead with models which incorporated switching behavior.

**Limitations**

Generative models such as rSLDS provide a unique opportunity to predict the evolution of neural activity over time. Because of time and availability, we were unable to implement additional analyses to evaluate the predictive ability of these rSLDS models or implement comparisons to more sophisticated neural decoders such as LFADS.

## Future Directions

This project represents the application of rSLDS to a random target pursuit behavioral task; these tools and analytic approaches could be applied to any collection of time series behavior and population activity. There is no reason to believe that rSLDS could not be applied to orofacial behavior, grasping behavior, or even naturalistic locomotion with neural network kinematic tracking such as DeepLabCut (Mathis et al. 2018). All in all, It is important to find out whether the primary motor cortex uses similar dynamical modes to produce different behaviors; by expanding the collection of behaviors, similarities between discrete states may be identified across categories of movements.

More sophisticated dynamics analysis would help elucidate meaningful differences or similarities between discrete states' linear dynamical systems. Since our study began, more advanced versions of recurrent dynamical systems models have been released including faster implementations of rSLDS (Chang et al. 2022). Other researchers have also extended rSLDS to model multiple separate neural populations (Glaser et al. 2020). We look forward to extending and validating the results presented in the current study using these recently developed methods. Other approaches such as LDR are converging on similar dynamical phenomena, approximating nonlinear dynamical systems to drive behavior (Sabatini and Kaufman 2021).

# 2: Conclusion

Recurrent Switching Linear Dynamical Systems (rSLDS), applied to neural data associated with more naturalistic and complex behavior such as the Random Target Pursuit Task, provides a tractable and interpretable set of computational tools to approximate the nonlinear dynamical system of primary motor cortical activity. The work presented in this dissertation demonstrates that rSLDS and similar models can be used to model and structure neural data. Also, the dynamics and structure identified by these models can be linked to behavior. In particular, the work presented here contributes to convergent support that M1 utilizes significantly different building blocks of neural activity to drive arm movements during accelerative and decelerative reaches in different directions.

M1 certainly engages with complex nonlinear dynamics for at least three crucial reasons. M1 must incorporate sensory feedback from multiple brain regions at diverse types of delays; it must also receive signals from higher-level motor areas containing sequence (Supplementary Motor Area) and preparatory (dorsal Premotor Cortex) information. It needs to incorporate these inputs while also coordinating the generation of neural patterns to generate movement via output pathways (e.g., red nucleus, thalamus, spinal cord). As such, it would make sense for M1 to generate different patterns of activity at different periods of time, potentially even during the execution of movement.

While M1 may drive movement using a nonlinear set of dynamics, computationally modeling such a system in a nonlinear way proves difficult. Optimizing and fitting a nonlinear dynamical system to neural data is currently intractable, so rSLDS sidesteps the issue by approximating a nonlinear system with a piecewise combination of linear dynamical systems

with discrete borders. Fitting, modeling, and analyzing sets of linear differential equations is much easier, therefore allowing researchers to evaluate hypotheses and present results in a more effective way. rSLDS is just one example of a collection of converging neural population activity analysis methods, but there are other approaches worth considering, especially in the pursuit of understanding the basic building blocks of activity that primary motor cortex may use to generate movement.

## Future Directions

While my dissertation is mostly focused on a set of models developed by the Linderman group in 2018, development has continued apace. As of writing, the Linderman group recently released a new set of models, underpinned by Google's JAX framework, dubbed DYNAMAX (Chang et al. 2022). DYNAMAX promises to make fitting and applying probabilistic latent state space models easier, all while enabling compatibility with modern machine learning capabilities and toolboxes. While not defined as a neuroscientific endeavor, groups like Linderman's continue to push the frontier of probabilistic latent state-space models, for all aspects of applications, not only neural activity.

A computational approach aimed squarely at uncovering the complexity of motor cortical activity comes from Matthew Kaufman's lab here at the University of Chicago, with their development of their Location Dependent Rotations (LDR) model of motor activity (Sabatini and Kaufman 2021). LDR posits that M1 activity inhabits a low-dimensional curved latent manifold, which supports many locations for rotational dynamics, depending on the direction and type of movement. By describing neural dynamics as a "record player in a large bowl," Sabatini and Kaufman put forward another way to approximate nonlinear dynamics in a tractable way. To

make considerable progress in understanding the structure of M1 activity and its connection to movement, strong and equitable collaborations must occur between experimentalists and computational groups. Collaborations must occur to ground progress in experimental work, and research labs which combine experiments with theory should continue to be supported.

**Other Behavioral Contexts**

One of the major potential benefits from probabilistic generative models such as rSLDS is that they do not rely upon trial-averaged data (Linderman et al. 2017). This characteristic allowed us to analyze RTP data, but it also means that these models are extensible enough to identify and model high-dimensional neural data from any brain area. Our implementation of rSLDS decomposes behavioral data using the structure identified purely from neural data and could be applied to other tasks and behavioral contexts. This implementation would allow us to address one of the largest potential follow-up questions from this dissertation: to what extent do these neural population states generalize to other behavioral contexts?

In our lab, we were able to record animals performing two different tasks (center-out reaching and RTP) within the same recording session. This means that we should be able to analyze the same cortical population generating two distinct tasks and compare the results. These data are a valuable tool for comparison, since it is still difficult to isolate neurons across recording sessions, especially with chronically implanted arrays. If distinct population states are a common characteristic of motor command structure, one would expect finding consistent state structure in different behavioral tasks. If states found in one task fail to be identified in another, it may hint at population states as a more limited and constrained phenomena. Additionally, we must further explore the nature of the transitions between said population states; while our

previous method displays a state transition between two 50ms time bins, the timescale and nature of transitions is certainly more complex, as evidenced by Schwartz' work (Suway et al. 2018b). One benefit of dual-task recordings is the ability to train neural models on data from one task and test the model on data from the other task in the same session. While this approach only works in situations where the same neural population is recorded in multiple tasks, it allows for a direct test of latent state applicability for these models.

Even without the ability to apply models from one task directly to different behavioral context, there remains enormous benefit for analyzing existing neural and behavioral data with new tools. For example, the Hatsopoulos Lab has gathered neural and kinematic data associated with animals reaching and grasping various objects (Saleh, Takahashi, and Hatsopoulos 2012). Additionally, the lab has a long-standing interest in orofacial control, including a large amount of data concerning tongue and jaw movement (Laurence-Chasen et al. 2023). Bringing these new probabilistic models to bear on additional rich datasets could prove extremely fruitful for identifying links between neural activity and behavior.

**Additional Brain Areas**

Just as the computational approaches mentioned here can be applied to other behavioral tasks, they can just as effectively be applied to other brain areas for the same tasks. In fact, population states have already been identified in visual cortex (Engel et al. 2016) and parietal cortex using Hidden Markov Models (Diomedi et al. 2021). While we could conduct new experiments, many chronic recording implants and resulting datasets from our lab contain neural populations from premotor cortex. Dorsal premotor cortex contains higher-order motor information about target location and movement preparation; it may demonstrate significantly

different neural population state structure and dynamics if modeled using rSLDS (Dekleva, Kording, and Miller 2018). Premotor cortex may utilize population states for motor planning during the preparation stages of movement but not during execution; however, little research has been conducted to confirm this. Conversely, premotor cortex may not exhibit state-based structure at all, implying M1 is unique in its temporal structure.

**Analyzing Muscle Activity**

Another important ambiguity to be interrogated is the relationship between neural activity in neocortex, and muscle activity. We still do not understand the signal transformations that occur along the motor pathway between these two parts of the nervous system, but by analyzing the temporal structure of activity at both points, we may gain some insight without needing to record directly from the spinal cord. There is evidence to suggest that intermittent changes in a motor control signal (in this case, a firing rate transient) are interpreted by the spinal cord and transformed into continuous activity in muscles (Shalit et al. 2012), but this concept has not been examined in the context of state-based population activity and transitions.

A possible approach to analyzing muscle activity would be to simply decompose signals from recorded muscles in the same manner as kinematics. However, a more interesting analysis would involve using EMG recordings from muscles as inputs to a switching linear dynamical system model instead of neural activity; these probabilistic models, although discrete, may be able to identify the latent structure of the transformed motor signal, if it exists. We would expect not to observe similar state structure at this point in the motor system, but it may uncover other latent structure to these signals.

## Overall Conclusion

From the very beginning, my intention with my PhD research and dissertation work were two-fold: to become a responsible, high-performing scientist, and to contribute in some way to the ongoing conversation about motor control. With these experiments, I have gained the skills necessary to conduct science well. I have also been able to contribute to the scientific conversation in two ways: there are consistent and significant differences in the neural dynamics which drive accelerative versus decelerative movements in various reach directions; in addition, I have demonstrated that using discrete state and linear dynamical system categorization to decompose kinematics is a useful, extensible, and tractable way to analyze complex, naturalistic behavioral data.

# References


Barrese, James C., Naveen Rao, Kaivon Paroo, Corey Triebwasser, Carlos Vargas-Irwin, Lachlan Franquemont, and John P. Donoghue. 2013. “Failure Mode Analysis of Silicon-Based Intracortical Microelectrode Arrays in Non-Human Primates.” *Journal of Neural Engineering* 10 (6). https://doi.org/10.1088/1741-2560/10/6/066014.

Black, Bryan J., Aswini Kanneganti, Alexandra Joshi-Imre, Rashed Rihani, Bitan Chakraborty, Justin Abbott, Joseph J. Pancrazio, and Stuart F. Cogan. 2018. “Chronic Recording and Electrochemical Performance of Utah Microelectrode Arrays Implanted in Rat Motor Cortex.” *Journal of Neurophysiology* 120 (4): 2083–90. https://doi.org/10.1152/jn.00181.2018.

Blackrock Microsystems. n.d. “Personal Communication.”

Bullard, Autumn J., Brianna C. Hutchison, Jiseon Lee, Cynthia A. Chestek, and Parag G. Patil. 2020. “Estimating Risk for Future Intracranial, Fully Implanted, Modular Neuroprosthetic Systems: A Systematic Review of Hardware Complications in Clinical Deep Brain Stimulation and Experimental Human Intracortical Arrays.” *Neuromodulation* 23 (4): 411–26. https://doi.org/10.1111/ner.13069.

Buzsáki, György. 2004. “Large-Scale Recording of Neuronal Ensembles.” *Nature Neuroscience* 7 (5): 446–51. https://doi.org/10.1038/nn1233.

Center for Devices and Radiological Health. n.d. “Implanted Brain-Computer Interface (BCI) Devices for Patients with Paralysis or Amputation - Non-Clinical Testing and Clinical Considerations.” https://www.fda.gov/regulatory-information/search-fda-guidance-documents/implanted-brain-computer-interface-bci-devices-patients-paralysis-or-amputation-non-clinical-testing.

Chang, Peter, Giles Harper-Donnelly, Aleyna Kara, Xinglong Li, Scott Linderman, and Kevin Murphy. 2022. “Welcome to Dynamax!” Github. 2022.

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. https://doi.org/10.1152/jn.1980.44.4.773.

Chestek, Cynthia A., Vikash Gilja, Paul Nuyujukian, Justin D. Foster, Joline M. Fan, Matthew T. Kaufman, Mark M. Churchland, et al. 2011. “Long-Term Stability of Neural Prosthetic Control Signals from Silicon Cortical Arrays in Rhesus Macaque Motor Cortex.” *Journal of Neural Engineering* 8 (4). https://doi.org/10.1088/1741-2560/8/4/045005.

Churchland, Mark M., John P. Cunningham, Matthew T. Kaufman, Justin D. Foster, Paul Nuyujukian, Stephen I. Ryu, Krishna V. Shenoy, and Krishna V. Shenoy. 2012. “Neural Population Dynamics during Reaching.” *Nature* 487 (7405): 51–56. https://doi.org/10.1038/nature11129.

Churchland, Mark M., and Krishna V. Shenoy. 2007. “Delay of Movement Caused by Disruption of Cortical Preparatory Activity.” *Journal of Neurophysiology* 97 (1): 348–59. https://doi.org/10.1152/jn.00808.2006.

Collinger, Jennifer L. 2021. “Personal Communication.”

Collinger, Jennifer L., Brian Wodlinger, John E. Downey, Wei Wang, Elizabeth C. Tyler-Kabara, Douglas J. Weber, Angus J.C. McMorland, Meel Velliste, Michael L. Boninger, and Andrew B. Schwartz. 2013. “High-Performance Neuroprosthetic Control by an Individual with Tetraplegia.” *The Lancet* 381 (9866): 557–64. https://doi.org/10.1016/S0140-6736(12)61816-9.

D’Avella, Andrea, Philippe Saltiel, and Emilio Bizzi. 2003. “Combinations of Muscle Synergies in the Construction of a Natural Motor Behavior.” *Nature Neuroscience* 6 (3): 300–308. https://doi.org/10.1038/nn1010.

Dekleva, Brian M., Konrad P. Kording, and Lee E. Miller. 2018. “Single Reach Plans in Dorsal Premotor Cortex during a Two-Target Task.” *Nature Communications* 9 (1). https://doi.org/10.1038/s41467-018-05959-y.

Dickey, Adam S, Yali Amit, and Nicholas G Hatsopoulos. 2013. “Heterogeneous Neural Coding of Corrective Movements in Motor.” *Front. Neural Circuits* 7 (April): 51.

Diomedi, S., F. E. Vaccari, C. Galletti, K. Hadjidimitrakis, and P. Fattori. 2021. “Motor-like Neural Dynamics in Two Parietal Areas during Arm Reaching.” *Progress in Neurobiology* 205 (October). https://doi.org/10.1016/j.pneurobio.2021.102116.

Doeringer, Joseph A., and Neville Hogan. 1998. “Serial Processing in Human Movement Production.” *Neural Networks* 11 (7–8): 1345–56. https://doi.org/10.1016/S0893-6080(98)00083-5.

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). https://doi.org/10.1088/1741-2552/aab7a0.

Dum, RP, and PL Strick. 1991. “The Origin of Corticospinal Projections from the Premotor Areas in the Frontal Lobe.” *The Journal of Neuroscience* 11 (3): 667–89. https://doi.org/10.1523/JNEUROSCI.11-03-00667.1991.

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 (1): 1–15. https://doi.org/10.1038/ncomms13239.

Engel, Tatiana A., Nicholas A. Steinmetz, Marc A. Gieselmann, Alexander Thiele, Tirin Moore, and Kwabena Boahen. 2016. “Selective Modulation of Cortical State during Spatial Attention.” *Science* 354 (6316): 1140–44. https://doi.org/10.1126/science.aag1420.

Fetz, Eberhard E., and Dom V. Finocchio. 1971. “Operant Conditioning of Specific Patterns of Neural and Muscular Activity.” *Science* 174 (4007): 431–35. https://doi.org/10.1126/science.174.4007.431.

Fox, Emily B, Erik B Sudderth, Michael I Jordan, and Alan S Willsky. 2010. “Bayesian Nonparametric Methods for Learning Markov Switching.” *IEEE Signal Process. Mag.* 27 (6): 43–54.

Fox, Emily, Erik Sudderth, Michael Jordan, and Alan Willsky. 2008. “Nonparametric Bayesian Learning of Switching Linear Dynamical.” *Adv. Neural Inf. Process. Syst.* 21.

Fraser, George W., and Andrew B. Schwartz. 2012. “Recording from the Same Neurons Chronically in Motor Cortex.” *Journal of Neurophysiology* 107 (7): 1970–78. https://doi.org/10.1152/jn.01012.2010.

Gawthrop, Peter, Ian Loram, Martin Lakie, and Henrik Gollee. 2011. “Intermittent Control: A Computational Theory of Human Control.” *Biological Cybernetics* 104 (1–2): 31–51. https://doi.org/10.1007/s00422-010-0416-4.

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

Georgopoulos, Apostolos P., Andrew B. Schwartz, and Ronald E. Kettner. 1986. “Neuronal Population Coding of Movement Direction.” *Science*. https://doi.org/10.1126/science.3749885.

Gibson, A. R., J. C. Houk, and N. J. Kohlerman. 1985. “Magnocellular Red Nucleus Activity during Different Types of Limb Movement in the Macaque Monkey.” *The Journal of Physiology* 358 (1): 527–49. https://doi.org/10.1113/jphysiol.1985.sp015565.

Glaser, Joshua I, Matthew Whiteway, John P Cunningham, Liam Paninski, and Scott W Linderman. 2020. “Recurrent Switching Dynamical Systems Models for Multiple Interacting Neural Populations.” *BioRxiv*, no. NeurIPS: 2020.10.21.349282. https://doi.org/10.1101/2020.10.21.349282.

Hatsopoulos, Nicholas G., Jignesh Joshi, and John G. O’Leary. 2004. “Decoding Continuous and Discrete Motor Behaviors Using Motor and Premotor Cortical Ensembles.” *Journal of Neurophysiology* 92 (2): 1165–74. https://doi.org/10.1152/jn.01245.2003.

Hatsopoulos, Nicholas G, Qingqing Xu, and Yali Amit. 2007. “Encoding of Movement Fragments in the Motor Cortex.” *J. Neurosci.* 27 (19): 5105–14.

Hughes, Christopher L., Sharlene N. Flesher, Jeffrey M. Weiss, John E. Downey, Jennifer L. Collinger, and Robert A. Gaunt. 2020. “Neural Stimulation and Recording Performance in Human Somatosensory Cortex over 1500 Days.” *MedRxiv*. https://doi.org/10.1101/2020.01.21.20018341.

Jackson, Andrew, Jaideep Mavoori, and Eberhard E. Fetz. 2006. “Long-Term Motor Cortex Plasticity Induced by an Electronic Neural Implant.” *Nature* 444 (7115): 56–60. https://doi.org/10.1038/nature05226.

Jones, Kelly E., Patrick K. Campbell, and Richard A. Normann. 1992. “A Glass/Silicon Composite Intracortical Electrode Array.” *Annals of Biomedical Engineering* 20 (4): 423–37. https://doi.org/10.1007/BF02368134.

Jun, James J., Nicholas A. Steinmetz, Joshua H. Siegle, Daniel J. Denman, Marius Bauza, Brian Barbarits, Albert K. Lee, et al. 2017. “Fully Integrated Silicon Probes for High-Density Recording of Neural Activity.” *Nature* 551 (7679): 232–36. https://doi.org/10.1038/nature24636.

Kadmon Harpaz, Naama, David Ungarish, Nicholas G Hatsopoulos, and Tamar Flash. 2019. “Movement Decomposition in the Primary Motor Cortex.” *Cereb. Cortex* 29 (4): 1619–33.

Karniel, Amir. 2013. “The Minimum Transition Hypothesis for Intermittent Hierarchical Motor Control.” *Frontiers in Computational Neuroscience* 7 (FEB): 1–8. https://doi.org/10.3389/fncom.2013.00012.

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

Kaufman, Matthew T., Jeffrey S. Seely, David Sussillo, Stephen I. Ryu, Krishna V. Shenoy, and Mark M. Churchland. 2016. “The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type.” *Eneuro* 3 (4): ENEURO.0085-16.2016. https://doi.org/10.1523/ENEURO.0085-16.2016.

Kemere, Caleb, Gopal Santhanam, Byron M. Yu, Afsheen Afshar, Stephen I. Ryu, Teresa H. Meng, and Krishna V. Shenoy. 2008. “Detecting Neural-State Transitions Using Hidden Markov Models for Motor Cortical Prostheses.” *Journal of Neurophysiology* 100 (4): 2441–52. https://doi.org/10.1152/jn.00924.2007.

Kozai, Takashi D.Y., Kasey Catt, Xia Li, Zhannetta V. Gugel, Valur T. Olafsson, Alberto L. Vazquez, and X. Tracy Cui. 2015. “Mechanical Failure Modes of Chronically Implanted Planar Silicon-Based Neural Probes for Laminar Recording.” *Biomaterials* 37: 25–39. https://doi.org/10.1016/j.biomaterials.2014.10.040.

Laurence-Chasen, Jeffrey D., Callum F. Ross, Fritzie I. Arce-McShane, and Nicholas G. Hatsopoulos. 2023. “Robust Cortical Encoding of 3D Tongue Shape during Feeding in Macaques.” *Nature Communications* 14 (1): 2991. https://doi.org/10.1038/s41467-023-38586-3.

Leib, Raz, Andrea d’Avella, and Llana Nisky. 2017. “An Intermittent Control Model Predicts the Triphasic Muscles Activity during Hand Reaching Raz,” October. https://doi.org/http://dx.doi.org/10.1101/206565.

Linderman, Scott, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski. 2017. “Bayesian Learning and Inference in Recurrent Switching Linear Systems.” In *Proceedings of the 20th International Conference on Artificial and Statistics*, edited by Aarti Singh and Jerry Zhu, 54:914–22. Proceedings of Machine Learning Research. PMLR.

Mathis, Alexander, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, and Matthias Mathis Mackenzie Weygandt and Bethge. 2018. “DeepLabCut: Markerless Pose Estimation of User-Defined Body with Deep Learning.” *Nat. Neurosci.* 21 (9): 1281–89.

Morasso, P. 1981. “Spatial Control of Arm Movements.” *Exp. Brain Res.* 42 (2): 223–27.

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

Musk, Elon. 2019. “An Integrated Brain-Machine Interface Platform with Thousands of Channels.” *BioRxiv*, January, 703801. https://doi.org/10.1101/703801.

Negi, S., R. Bhandari, L. Rieth, and F. Solzbacher. 2010. “In Vitro Comparison of Sputtered Iridium Oxide and Platinum-Coated Neural Implantable Microelectrode Arrays.” *Biomedical Materials (Bristol, England)* 5 (1): 15007. https://doi.org/10.1088/1748-6041/5/1/015007.

Nordhausen, Craig T., Edwin M. Maynard, and Richard A. Normann. 1996. “Single Unit Recording Capabilities of a 100 Microelectrode Array.” *Brain Research* 726 (1–2): 129–40. https://doi.org/10.1016/0006-8993(96)00321-6.

“Offline Sorter.” 2017. *Plexon*.

Omrani, Mohsen, Matthew T Kaufman, Nicholas G Hatsopoulos, and Paul D Cheney. 2017. “Perspectives on Classical Controversies about the Motor Cortex.” *J. Neurophysiol.* 118 (3): 1828–48.

Pachitariu, Marius, Shashwat Sridhar, and Carsen Stringer. 2023. “Solving the Spike Sorting Problem with Kilosort.” *BioRxiv*.

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. https://doi.org/10.1038/s41592-018-0109-9.

Paninski, Liam, Matthew R Fellows, Hatsopoulos Nicholas G, and John P Donoghue. 2004. "Spatiotemporal Tuning of Motor Cortical Neurons for Hand position and Velocity." *J. Neurophysiol.* 91 (1): 515–32.

Pei, Felix, Joel Ye, David Zoltowski, Raeed H Wu Anqi and Chowdhury, Hansem Sohn, Joseph E O'Doherty, Krishna V Shenoy, et al. 2021. "Neural Latents Benchmark '21: Evaluating Latent Variable of Neural Population Activity," September.

Piech, David K., Benjamin C. Johnson, Konlin Shen, M. Meraj Ghanbari, Ka Yiu Li, Ryan M. Neely, Joshua E. Kay, Jose M. Carmena, Michel M. Maharbiz, and Rikky Muller. 2020. "A Wireless Millimetre-Scale Implantable Neural Stimulator with Ultrasonically Powered Bidirectional Communication." *Nature Biomedical Engineering* 4 (2): 207–22. https://doi.org/10.1038/s41551-020-0518-9.

Polikov, Vadim S., Patrick A. Tresco, and William M. Reichert. 2005a. "Response of Brain Tissue to Chronically Implanted Neural Electrodes." *Journal of Neuroscience Methods* 148 (1): 1–18. https://doi.org/10.1016/j.jneumeth.2005.08.015.

———. 2005b. "Response of Brain Tissue to Chronically Implanted Neural Electrodes." *Journal of Neuroscience Methods*. https://doi.org/10.1016/j.jneumeth.2005.08.015.

Rathelot, J.-A., and P. L. Strick. 2009. "Subdivisions of Primary Motor Cortex Based on Cortico-Motoneuronal Cells." *Proceedings of the National Academy of Sciences* 106 (3): 918–23. https://doi.org/10.1073/pnas.0808362106.

Rousche, Patrick J., and Richard A. Normann. 1998. "Chronic Recording Capability of the Utah Intracortical Electrode Array in Cat Sensory Cortex." *Journal of Neuroscience Methods* 82 (1): 1–15. https://doi.org/10.1016/S0165-0270(98)00031-4.

Sabatini, David A, and Matthew T Kaufman. 2021. "A Curved Manifold Orients Rotational Dynamics in Motor Cortex." *BioRxiv*.

Salatino, Joseph W., Kip A. Ludwig, Takashi D.Y. Kozai, and Erin K. Purcell. 2017. "Glial Responses to Implanted Electrodes in the Brain." *Nature Biomedical Engineering* 1 (11): 862–77. https://doi.org/10.1038/s41551-017-0154-1.

Saleh, Maryam, Kazutaka Takahashi, and Nicholas G. Hatsopoulos. 2012. "Encoding of Coordinated Reach and Grasp Trajectories in Primary Motor Cortex." *The Journal of Neuroscience* 32 (4): 1220–32. https://doi.org/10.1523/JNEUROSCI.2438-11.2012.

Santhanam, Gopal, Michael D. Linderman, Vikash Gilja, Afsheen Afshar, Stephen I. Ryu, Teresa H. Meng, and Krishna V. Shenoy. 2007. “HermesB: A Continuous Neural Recording System for Freely Behaving Primates.” *IEEE Transactions on Biomedical Engineering*. https://doi.org/10.1109/TBME.2007.895753.

Schwartz, A B. 1994. “Direct Cortical Representation of Drawing.” *Science* 265 (5171): 540–42.

Scott, Stephen H., and S An. 2004. “Optimal Feedback Control and the Neural Basis of Volitional Motor Control.” *Nature Reviews Neuroscience* 5 (7): 532–46. https://doi.org/10.1038/nrn1427.

Shalit, Uri, Nofya Zinger, Mati Joshua, and Yifat Prut. 2012. “Descending Systems Translate Transient Cortical Commands into a Sustained Muscle Activation Signal.” *Cerebral Cortex* 22 (8): 1904–14. https://doi.org/10.1093/cercor/bhr267.

Shenoy, Krishna V., Maneesh Sahani, and Mark M. Churchland. 2013. “Cortical Control of Arm Movements: A Dynamical Systems Perspective.” *Annual Review of Neuroscience* 36 (1): 337–59. https://doi.org/10.1146/annurev-neuro-062111-150509.

Sponheim, Caleb, Vasileios Papadourakis, Jennifer L. Collinger, John Downey, Jeffrey Weiss, Lida Pentousi, Kaisa Elliott, and Nicholas G. Hatsopoulos. 2021. “Longevity and Reliability of Chronic Unit Recordings Using the Utah, Intracortical Multi-Electrode Arrays.” *Journal of Neural Engineering* 18 (6). https://doi.org/10.1088/1741-2552/ac3eaf.

Suner, Selim, Matthew R. Fellows, Carlos Vargas-Irwin, Gordon Kenji Nakata, and John P. Donoghue. 2005. “Reliability of Signals from a Chronically Implanted, Silicon-Based Electrode Array in Non-Human Primate Primary Motor Cortex.” *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 13 (4): 524–41. https://doi.org/10.1109/TNSRE.2005.857687.

Suway, S. B., J. Orellana, A. J.C. McMorland, G. W. Fraser, Z. Liu, M. Velliste, S. M. Chase, R. E. Kass, and A. B. Schwartz. 2018a. “Temporally Segmented Directionality in the Motor Cortex.” *Cerebral Cortex* 28 (7): 2326–39. https://doi.org/10.1093/cercor/bhx133.

———. 2018b. “Temporally Segmented Directionality in the Motor Cortex.” *Cerebral Cortex* 28 (7): 2326–39. https://doi.org/10.1093/cercor/bhx133.

Takei, Tomohiko, Troy M Herter, Kevin P Cross, and Stephen H Scott. 2018. “Correlations between Primary Motor Cortex Activity with Recent Past and Future Limb Motion during Unperturbed Reaching Title ( 16 / 50 Words ) Correlations between Primary Motor Cortex Activity with Recent Past and Future Limb Motion during Unperturbed Re.” https://doi.org/10.1523/JNEUROSCI.2667-17.2018.

Tanji, Jun, and Edward V Evarts. 1975. “Anticipatory Activity of Motor Cortex Neurons in Relation to Direction of an Intended Movement.”

Todorov, Emanuel, and Michael I. Jordan. 2002. “Optimal Feedback Control as a Theory of Motor Coordination.” *Nature Neuroscience* 5 (11): 1226–35. https://doi.org/10.1038/nn963.

Vyas, Saurabh, Matthew D. Golub, David Sussillo, and Krishna V. Shenoy. 2020. “Computation Through Neural Population Dynamics.” *Annual Review of Neuroscience* 43 (1): 249–75. https://doi.org/10.1146/annurev-neuro-092619-094115.

Walker, Jeff, Jason MacLean, and Nicholas G. Hatsopoulos. 2017. “The Marmoset as a Model System for Studying Voluntary Motor Control.” *Developmental Neurobiology* 77 (3): 273–85. https://doi.org/10.1002/dneu.22461.

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. https://doi.org/10.1152/jn.90941.2008.

Zimnik, Andrew J., and Mark M. Churchland. 2021. “Independent Generation of Sequence Elements by Motor Cortex.” *Nature Neuroscience* 24 (3): 412–24. https://doi.org/10.1038/s41593-021-00798-5.

Zoltowski, David M., Jonathan W. Pillow, and Scott W. Linderman. 2020. “Unifying and Generalizing Models of Neural Dynamics during Decision-Making,” January.

# Appendix

## A. Longevity and Reliability of Chronic Unit Recordings


### Abstract

Objective. Microelectrode arrays are standard tools for conducting chronic electrophysiological experiments, allowing researchers to simultaneously record from large numbers of neurons. Specifically, Utah electrode arrays (UEAs) have been utilized by scientists in many species, including rodents, rhesus macaques, marmosets, and human participants. The field of clinical human brain-computer interfaces currently relies on the UEA as a number of research groups have FDA clearance for this device through the investigational device exemption pathway. Despite its widespread usage in systems neuroscience, few studies have comprehensively evaluated the reliability and signal quality of the Utah array over long periods of time in a large dataset. Approach. We collected and analyzed over six thousand recorded datasets from various cortical areas spanning almost 9 years of experiments, totaling seventeen rhesus macaques (Macaca Mulatta) and two human subjects, and fifty-five separate microelectrode Utah arrays. The scale of this dataset allowed us to evaluate the average life of these arrays, based primarily on the signal-to-noise ratio of each electrode over time. Main Results. Using implants in primary motor, premotor, prefrontal, and somatosensory cortices, we found that the average lifespan of available recordings from UEAs was 622 days, although we provide several examples of these UEAs lasting over one thousand days and one up to 9 years; human implants were also shown to last longer than non-human primate implants. We also found

that electrode length did not affect longevity and quality, but iridium oxide metallization on the electrode tip exhibited superior longevity and quality as compared to platinum metallization.

Significance. Understanding longevity and reliability of microelectrode array recordings allows researchers to set expectations and plan experiments accordingly and maximize the amount of high-quality data gathered. Our results suggest that one can expect chronic unit recordings to last at least two years, with the possibility for arrays to last the better part of a decade.

## 1. Introduction

First developed in 1992, a microelectrode array recording platform now called the “Utah Electrode Array” (UEA) has provided the opportunity for researchers to record from populations of neurons simultaneously, in-vivo, in a chronic preparation (Jones, Campbell, and Normann 1992). Today, UEAs are a standard data collection method for systems neuroscience and represent a common choice for chronic multi-electrode array recordings, particularly in animal models such as rodents and non-human primates (NHPs), including macaques and marmosets (Walker, MacLean, and Hatsopoulos 2017; Black et al. 2018; Barrese et al. 2013). With the popularity of such a tool comes increased scrutiny on the reliability and failure modes of the technology. UEAs, like all chronic neural implants that interface with brain tissue, usually degrade in their recording capabilities over time (Barrese et al. 2013); in some cases, these implants can experience more acute issues, either during surgical implantation or shortly thereafter, whereas in many other cases these implants can reliably record unit activity for several years (Barrese et al. 2013). These implants have been in widespread use for many years,

yet there have been few attempts to quantify the reliability of this technology in-vivo (Barrese et al. 2013). Most previous efforts to quantify the signal quality of UEAs over time, while laudable, have focused heavily on the consistency of neuronal characteristics in a small sample size of implants and subjects (Rousche and Normann 1998; Suner et al. 2005; Chestek et al. 2011; Downey et al. 2018; Fraser and Schwartz 2012). Other studies have examined the longevity and reliability of other chronic electrophysiological recording technologies, to help researchers manage expectations and select the right tool for the job (Kozai et al. 2015; Salatino et al. 2017; Polikov, Tresco, and Reichert 2005a; Buzsáki 2004). However, one previous study has examined the failure modes and reliability of UEAs from a large sample size, albeit with UEAs undergoing significant changes in manufacturing processes (4). The current work aims to extend this previous study by considering more recent UEAs whose manufacturing process has become relatively stable and considers UEAs with iridium oxide metallization optimized for stimulation as well as platinum metallization on the electrode tips.

UEAs currently remain the only FDA-cleared intracortical implant for chronic human neuroscientific studies under Investigational Device Exemptions (IDE), an allowance which is held by several groups across the United States. There have been at least eighteen human subjects who have received a chronic UEA implant for brain computer interface (BCI) studies reported in the literature as of 2019 (Bullard et al. 2020), and we estimate that at least another seven people have been implanted since then (Collinger 2021). Human studies require ongoing consideration of the risk-benefit ratio of implanting human subjects, which will be informed by more information on the expected reliability and longevity of these implants over time. A majority of human studies expect participants to remain in the study for at least 1 year, which is supported by a significant body of evidence demonstrating the safety and efficacy of these

devices for long term recording (Walker, MacLean, and Hatsopoulos 2017; Black et al. 2018). Here we extend this earlier work by examining a large number of array implants over extended durations up to 9 years in both humans and non-human primate subjects.

It is likely that the UEA will continue to be the dominant neural recording platform for human intracortical BCI studies in the near future. FDA clearance for investigational approval still requires the collection of significant longitudinal safety and efficacy data along with the institutional knowledge for a successful clinical trial design (Center for Devices and Radiological Health, n.d.). A number of efforts are currently underway to develop more advanced tools for chronic implants across non-human and human primates alike (Jun et al. 2017; Musk 2019; Piech et al. 2020), but given the prevalence of the UEA for chronic recording and stimulation, it is imperative to characterize the reliability of these arrays across and within species, to take full advantage of their capabilities.

Here, using data from over 6000 recording sessions in 55 array implants across human (n=2) and non-human primate (n=17) subjects, we assess the reliability of long-term unit recordings in primary motor, dorsal and ventral premotor, prefrontal, and primary somatosensory cortices. We show that nearly fifty percent of implants exhibit year-long recordings with a yield greater than 40% of the total available electrodes displaying satisfactory quality (SNR>1.5). Moreover, long-term recordings from ~1000 days and up to nine years are shown to be possible in some cases. We also examine the influence of electrode length and electrode-tip metallization on the longevity and quality of recordings.

# 2. Methods

## 2.1 Non-Human Primate Subjects

Recordings collected during the period of 2003-2020 from a total of 55 UEAs implanted in seventeen rhesus macaques (*Macaca Mulatta*) (10 female, seven male) were analyzed. This project analyzed data from the maximum possible number of recordings throughout the history of our research group. All animals were implanted with at least one (often two or more) UEAs (Blackrock Microsystems, Inc. Salt Lake City, UT). The earliest implants were manufactured by Cyberkinetics Neurotechnology Systems, Inc., Foxboro MA whose research business was sold to Blackrock Microsystems, Inc. in 2008. The metallization, number of electrodes, and electrode length of these arrays varied depending on brain area, scientific need, and implantation year, and individual array details are included in Table 1. Briefly, arrays were either implanted as an 8x8 or 10x10 grid, with either 1.5 mm or 1.0 mm length electrodes (although never mixed within a given array), with electrode tips metallized with either platinum or iridium-oxide. Since our research group focuses primarily on upper limb motor tasks, UEAs were most often implanted in primary motor (M1), dorsal premotor (PMd), or ventral premotor cortices (PMv). Other implanted areas included somatosensory (S1), orofacial primary motor (M1o), and prefrontal cortex (PFC). We saw no appreciable difference in signal longevity or quality between the two areas which had large enough numbers of implants to compare (dorsal/ventral premotor versus primary motor areas). All implants were connected to CerePort connectors with analog headstages produced by Blackrock Microsystems. All arrays were wired in a consistent manner; none were hand soldered.

Data collection was conducted using Blackrock's Cerebus data acquisition system. Neural data recorded as analog signals were amplified with a gain of five thousand, bandpass filtered using the built-in hardware filter between 0.3 Hz and 7.5 kHz and digitized at 30 kHz. Spike waveforms were detected and saved into the ".nev" file format at 46 samples per waveform. Spike events were detected using a global root mean squared signal energy (RMS) threshold set by the experimenter; occasionally the experimenter manually adjusted individual electrodes' thresholds.

Due to the broad date range of these recordings, training protocols and behavioral tasks varied widely across our subjects. Most subjects were trained to perform upper limb motor tasks, either involving two-dimensional planar reaching tasks using a two-link exoskeletal robot (Kinarm, BKIN Technologies, Ltd., Kingston, Ontario), unconstrained three-dimensional reaching, grasping, and reach-to-grasp tasks, or brain-computer interface (BCI) tasks. One animal also engaged in orofacial behavior.

### 2.2 Human Subjects

Two 96-channel arrays were implanted in participant P1, and two 88-channel arrays were implanted in participant P2. Arrays were implanted in the hand knob area of motor cortex (precentral gyrus). P2 also had two thirty-two channel arrays implanted in somatosensory cortex for the purposes of electrical stimulation; recording and stimulation performance for these arrays has been previously described (Hughes et al. 2020). Participant P1 contributed to 350 recording sessions over 2.7 years. Participant P2 contributed to 721 recordings sessions over 5.3 years (Downey et al. 2018). Data collection with participant P2 is ongoing. Neural data were recorded

at the beginning of each BCI testing session while the participants were at rest or speaking with experimenters.

### 2.3 Neural Data Processing

Each ".nev" file was processed to evaluate the number of electrodes that displayed spike events over the entire recording. A channel's spikes were included in an SNR calculation if at least fourteen spike events were detected throughout the entire recording. Each electrode's signal-to-noise ratio (SNR) was calculated by measuring the average peak-to-trough amplitude of the electrode's waveforms and dividing it by two times the average standard deviation of the voltage of the waveforms. The average standard deviation was computed by taking the mean of standard deviation values over all forty-six samples of the waveform. If the electrode's SNR exceeded 1.5, the electrode was designated as "good" and counted towards that recording's array yield. This threshold guarantees that a channel's signal is 50% higher than the noise of the channel. This threshold may be seen as too high by some, but we preferred to have confidence in a "good" channel, rather than include channels of nebulous quality. It should be noted that neural signals with an SNR of less than 1.5 can still carry information that may be useful for BCI or scientific applications. Moreover, each recording was manually inspected and erroneous "good" channels, often having inaccurate SNR measurements due to artifacts, were removed from the "good" electrode count. No spike sorting was performed so our SNR values were conservative and were lower than would be expected from individually sorted units. Also, due to the lack of spike sorting, we could not address the degree to which array recordings change in their capability to detect separable units within a given electrode's signal. All statistics and figures displaying SNR were calculated across all electrodes of the array in which spikes were detected

After manual inspection, our search resulted in data from nineteen subjects (17 NHPs and two humans) and 55 Utah microelectrode arrays, totaling 6132 recording sessions. Due to different array sizes (8x8, 10x10), array yield of arrays was not presented as the absolute number of electrodes but rather as the percentage of total recordable electrodes**,** thus allowing for comparisons across different electrode configurations. SNR was also calculated as an average across all electrodes within a given recording session.

It is important to note that the number of arrays we could examine decreased with time due to loss of signals, hardware failure, infections surrounding the percutaneous connector, and experiment termination (see Table 1). Some analyses took loss of arrays into account to represent the number of arrays that survived past a certain point in time (see Figure 7). Other analyses prioritized calculating the reliability and yield of functioning arrays, considering only arrays from which viable recordings were possible (see Figure 9).

| Name | Sex | Brain Area | Implant Date | Electrode Length (mm) | Size | Metallization | Number of Recordings | Reason for Termination |
|---|---|---|---|---|---|---|---|---|
| AtPmv | F | PMv | 80123 | 1 | 96 | Platinum | 47 | Signal Quality |
| AtM1b | F | M1 | 130305 | 1 | 128 | Iridium Oxide | 40 | Signal Quality |
| AtM1o | F | M1o | 80123 | 1 | 96 | Platinum | 13 | Medical |
| BiM1 | M | M1o | 100824 | 1 | 96 | Iridium Oxide | 10 | Signal Quality |
| BoPMda | M | PMd | 30128 | 1 | 96 | Platinum | 317 | Signal Quality |
| BoM1b | M | M1 | 80408 | 1.5 | 96 | Iridium Oxide | 68 | Signal Quality |
| BoPMdb | M | PMd | 80408 | 1 | 96 | Platinum | 68 | Signal Quality |
| BoM1a | M | M1 | 30128 | 1 | 96 | Platinum | 6 | Signal Quality |
| CoPMd | F | PMd | 90413 | 1 | 96 | Platinum | 82 | Medical |
| CoM1 | F | M1 | 90413 | 1.5 | 96 | Platinum | 68 | Medical |
| CoPMv | F | PMv | 90413 | 1.5 | 96 | Platinum | 54 | Medical |
| JaPMv | M | PMv | 121002 | 1.5 | 96 | Iridium Oxide | 27 | Signal Quality |
| JaPMd | M | PMd | 121002 | 1 | 96 | Iridium Oxide | 17 | Signal Quality |
| KiM1ips | F | M1ipsi | 120604 | 1 | 96 | Iridium Oxide | 327 | Study End |
| KiM1con | F | M1con | 120604 | 1 | 96 | Iridium Oxide | 51 | Study End |
| LeM1c | M | M1l | 150720 | 1 | 128 | Iridium Oxide | 202 | Signal Quality |
| LeM1b | M | M1m | 150720 | 1 | 128 | Iridium Oxide | 63 | Signal Quality |
| LeM1a | M | M1 | 110711 | 1.5 | 96 | Iridium Oxide | 38 | Signal Quality |
| LePMd | M | PMd | 110711 | 1.5 | 96 | Platinum | 32 | Signal Quality |
| LePMv | M | PMv | 110711 | 1.5 | 96 | Platinum | 25 | Signal Quality |
| MkM1c | M | M1m | 71128 | 1 | 96 | Iridium Oxide | 501 | Hardware Failure |
| MkM1b | M | M1l | 60906 | 1 | 96 | Platinum | 32 | Hardware Failure |
| MkM1d | M | M1l | 71128 | 1 | 96 | Platinum | 32 | Hardware Failure |
| MkM1a | M | M1m | 60906 | 1 | 96 | Platinum | 28 | Hardware Failure |
| MkPMd | M | PMd | 60906 | 1 | 96 | Platinum | 14 | Hardware Failure |
| NkM1b | F | M1contra | 140303 | 1 | 96 | Iridium Oxide | 42 | Study End |
| NkPFC | F | PFCcontra | 140303 | 1 | 96 | Iridium Oxide | 13 | Study End |
| NkM1a | F | M1ipsi | 140303 | 1 | 96 | Iridium Oxide | 12 | Study End |
| NiPMv | F | PMv | 41103 | 1 | 96 | Platinum | 91 | Signal Quality |
| NiM1 | F | M1 | 41103 | 1 | 96 | Platinum | 86 | Signal Quality |
| NiPMd | F | PMd | 41103 | 1 | 96 | Platinum | 21 | Signal Quality |
| OrM1a | F | M1 | 70501 | 1.5 | 96 | Platinum | 52 | Signal Quality |
| OrM1b | F | M1o | 70501 | 1 | 96 | Platinum | 52 | Signal Quality |
| OrPmv | F | PMv | 70501 | 1 | 96 | Platinum | 18 | Signal Quality |
| RjPMd | M | PMd | 31027 | 1 | 96 | Platinum | 156 | Signal Quality |
| RjM1a | M | M1 | 31027 | 1 | 64 | Platinum | 81 | Signal Quality |
| RjM1b | M | M1 | 50804 | 1 | 96 | Platinum | 21 | Hardware Failure |
| RjPMv | M | PMv | 50804 | 1 | 96 | Platinum | 10 | Hardware Failure |
| RoM1a | M | M1 | 40602 | 1 | 96 | Platinum | 211 | Signal Quality |
| RoPMd | M | PMd | 40602 | 1 | 96 | Platinum | 34 | Signal Quality |
| RxM1a | F | M1 | 50328 | 1 | 96 | Platinum | 97 | Signal Quality |
| RxPMva | F | PMv | 50328 | 1 | 96 | Platinum | 80 | Signal Quality |
| RxM1b | F | M1 | 60712 | 1.5 | 96 | Platinum | 55 | Signal Quality |
| RxPMdb | F | PMd | 60712 | 1 | 96 | Platinum | 42 | Signal Quality |
| RxM1c | F | M1 | 90729 | 1.5 | 96 | Iridium Oxide | 33 | Signal Quality |
| RxPMvb | F | PMv | 60712 | 1 | 96 | Platinum | 27 | Signal Quality |
| RxPMda | F | PMd | 50328 | 1 | 96 | Platinum | 10 | Signal Quality |
| RxS1a | F | S1 | 90729 | 1.5 | 96 | Iridium Oxide | 6 | Signal Quality |
| VePMv | F | PMv | 50809 | 1 | 32 | Platinum | 132 | Medical |
| VeM1a | F | M1 | 50809 | 1 | 96 | Platinum | 110 | Medical |
| VePMd | F | PMd | 50809 | 1 | 96 | Platinum | 48 | Medical |
| ZiM1a | F | M1contra | 120730 | 1 | 96 | Iridium Oxide | 384 | Study End |
| ZiM1c | F | M1ipsi | 120730 | 1 | 96 | Iridium Oxide | 5 | Study End |

**Table 1. List of microelectrode arrays included in the current study.**

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| P1_A | F | M1 | 120210 | 1.5 | 96 | Platinum | 350 | Medical |
| P1_P | F | M1 | 120210 | 1.5 | 96 | Platinum | 335 | Medical |
| P2_A | M | M1 | 150504 | 1.5 | 88 | Platinum | 722 | N/A |
| P2_P | M | M1 | 150504 | 1.5 | 88 | Platinum | 718 | N/A |

**Table 1 Continued. List of microelectrode arrays included in the current study.** Column A indicates the abbreviation used for the identity of the animal and brain area. Many animals were implanted with multiple arrays. Column B indicates the sex of the animal in which the array was implanted. Column C indicates the brain area in which the array was implanted. Column D indicates the date of initial array implantation in YYMMDD format. The rightmost column is reserved for reasons for explant. "Medical" encompasses infections, skin retraction, or other medical issues which required the implant to be removed, unrelated to signal quality. "Signal Quality" refers to poor signal quality or yield. "Study End" indicates that the arrays were explanted not due to poor signal quality, but due to the end of a given experiment. "Hardware failure" indicates non-array components failing such that recording ability was compromised. The final four rows of Table 1 are reserved for the human participant arrays which were included in the study.

# 3. Results

### 3.4 Lifetime of chronic recordings

A total of over six thousand recording sessions over fifty-five arrays in seventeen animals and two human subjects were analyzed by computing SNR values of unsorted waveforms for each electrode on the array (Figure 5). A subset of electrodes was then selected as "good" if their SNR exceeded 1.5 and passed our manual examination from which overall yield was assessed over time. Over the fifty-five array implants, there was a large variance in lifetime of recordings from 44 days to over three thousand days (Figure 2). Recordings from UEAs were terminated for several reasons including loss of signals, electrode assembly hardware failure (i.e. the titanium connector pedestal detached from the skull), infections surrounding the connector pedestal, or completion of an experiment (Table 1).

To gauge the likelihood that a UEA would provide high quality signals over time, we examined the proportion of arrays that exceeded a fixed percentage yield at month-to-month

intervals (Figure 7). We observed a slight increase in the proportion of arrays exceeding a given yield in the first month of recording post-implantation followed by a steady decrease over 36 months. Nearly seventy percent of arrays displayed at least a 40% yield in the first three months of recordings and fifty percent of arrays exceeded the same yield threshold after one year of recording. A long tail of reliability was observed with nearly ten percent of arrays exceeding the same yield threshold out to at least 36 months post-implantation.

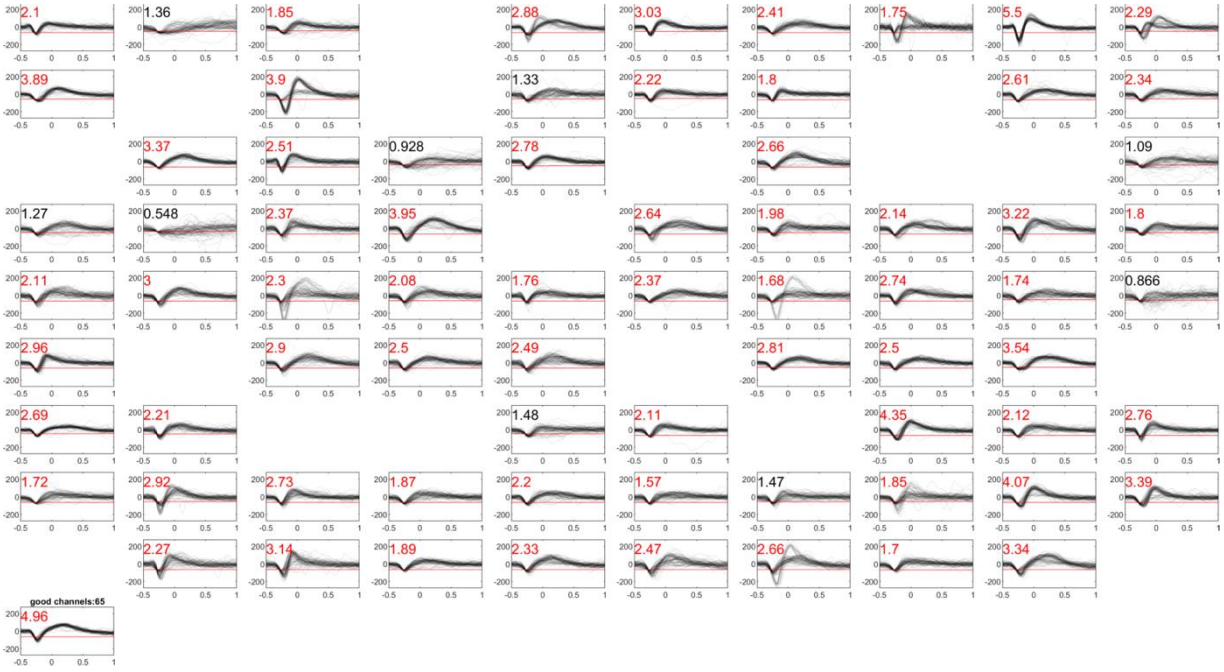


**Figure 5. Example waveforms and signal-to-noise ratio (SNR) from a UEA recording session** (Subject Mk, array MkM1c). SNR values are presented in the top left corner of each panel. Red SNR values indicate electrodes which exceeded the requisite SNR threshold of 1.5 to be considered as "good" electrodes included in subsequent analyses of array yield. The red line in each panel indicates the threshold across which a waveform must cross to be considered an action potential. Empty panels indicate electrodes in which less than fourteen spikes were detected within the given recording session, and therefore were not included in analysis.

### 3.2 Extended long-term recordings

To examine the day-to-day variability in signal quality among very long-term recordings, we focused on a subset of arrays that continued to record good signals beyond 950 days (Figure

8a). Sixteen out of fifty-five arrays displayed longevity well over eight hundred days of which three lasted well into five years of service (two of the three arrays were implanted in human subject P2). One array (Mk) exhibited extended recording capabilities to nearly nine years, only needing to be explanted due to an infection near the connector. This subset of arrays generally displayed a gradual decrease in signal quality over time with variations from recording to recording presumably due to several factors such as micro-movements of the array that may have occurred in the home cage such as sudden head accelerations

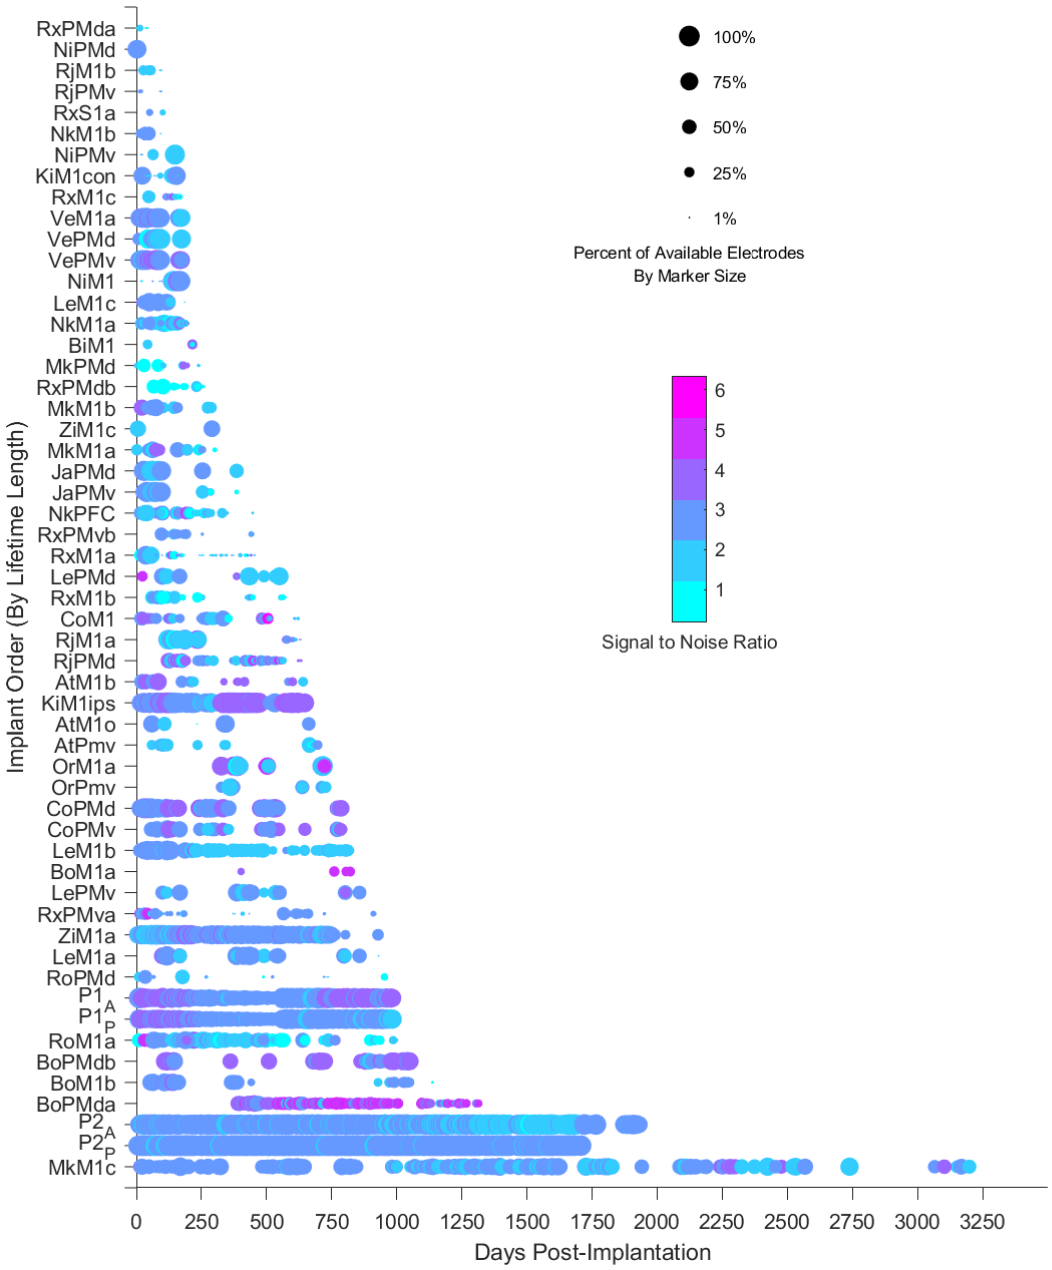


**Figure 6. Summary heat map of signal-to-noise ratio (circle color) and array yield (circle size) over time for all array implants analyzed**. Each colored circle denotes a single recording session. The color of the circle denotes the average SNR of the array for that given recording session. The size of the circle denotes the percentage of electrodes in that array that demonstrated a signal-to-noise ratio above a threshold of 1.5.

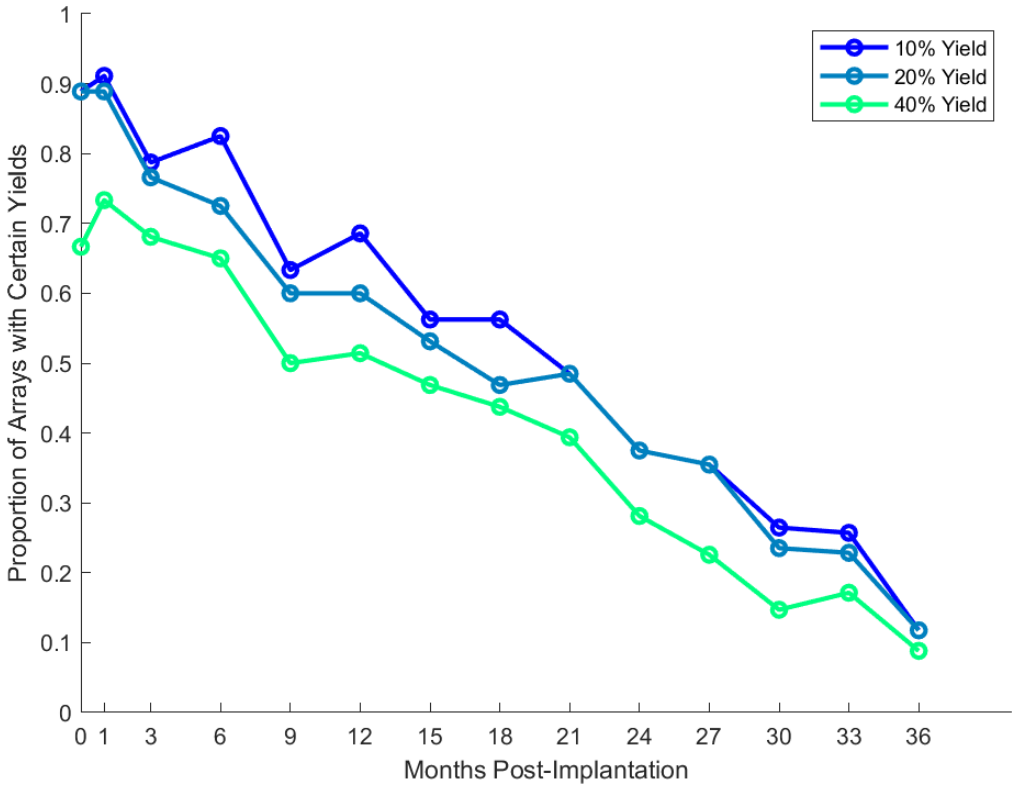


**Figure 7. Proportion of arrays exceeding a certain yield at month-to-month intervals post-implantation**. Each line delineates a different percentage yield threshold. If an array is not recorded from in a given time window, it is not counted in the "proportion" estimate. After their last available recording date, arrays are counted in the proportion measurement, to accurately depict the degradation of arrays on average. However, arrays which ceased recording due to medical, hardware failure, or study end reasons (see Table 1) were not included in the proportion measurement after their last available recording date.

(excluding our human participants) (Santhanam et al. 2007), headstage malfunctions, and sources of electrical noise of unknown origin that could not be eliminated. The two arrays implanted in a human subject (P2) are still implanted as of July 2021 with good recording quality.

### 3.3 Reliability among arrays possessing viable recordings

We next examined the temporal evolution of average yield and SNR among arrays from which recordings continued to be available and did not include arrays from which recordings were terminated (Figure 9). Therefore, the number of arrays contributing to the average decreased with time, and standard error values increased. However, with this measure, we were able to confirm that arrays implanted in NHPs typically maintained their electrode yield and SNR over most of their lifetime (Figure 9a). In contrast to long-term yield, short-term yield displayed a rapid increase over the first 40 days post-implantation in NHP implants (Figure 5b).

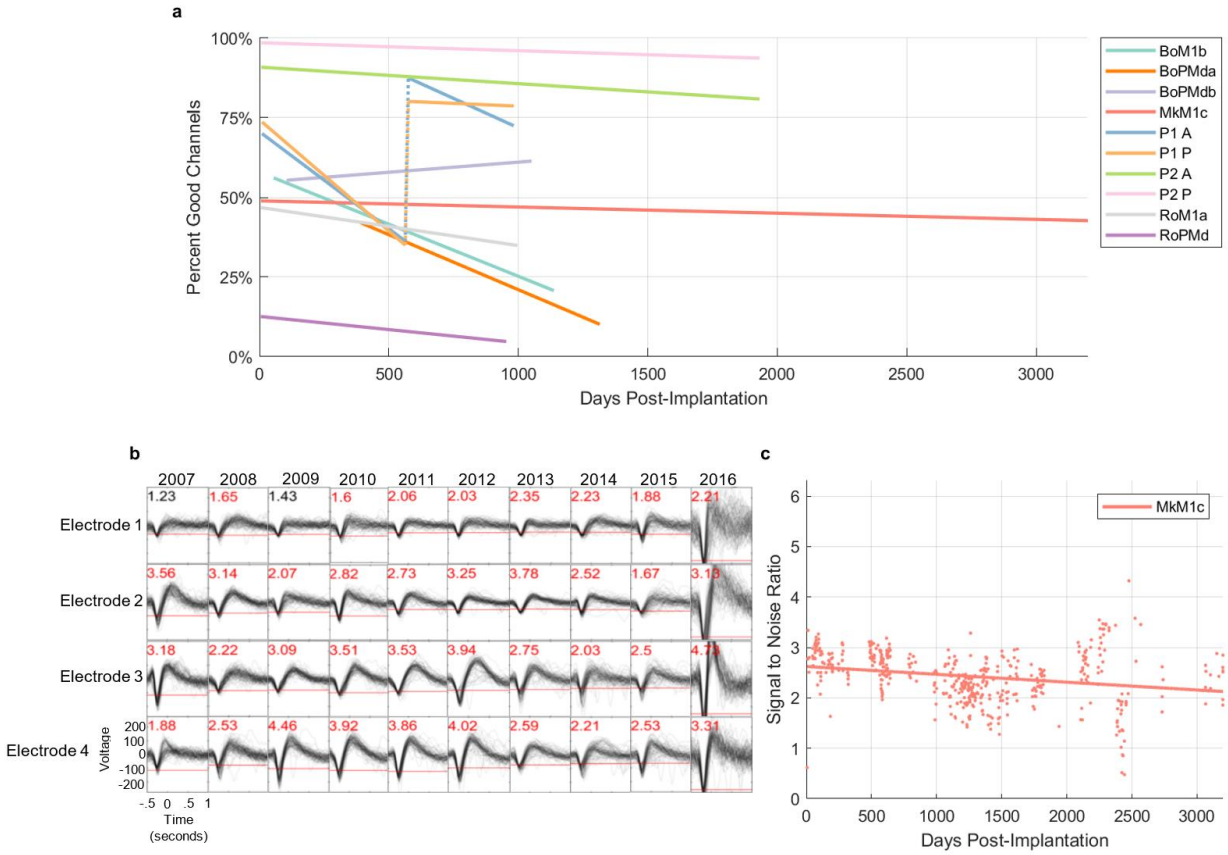


**Figure 8. Extended long-term performance for a subset of arrays**. a. Yield for long-term array implants as a function of days post-implantation. Lines indicate best fit linear regressions to the data for each array. The recording quality metrics exhibit a discontinuity at day 565 for subject P1, due to a change in the spike threshold from −5.25 to −4.5 RMS, indicated by the vertical dashed lines. b. Example spike waveforms from one array implant (MkM1c) at regular intervals over nearly 9 years. SNR values are presented at the top left corner of each panel (red font denotes electrodes that exceeded an SNR threshold of 1.5). c. SNR of Monkey Mk's recordings, over the lifetime of the implant.

Using the same methods used for NHP analyses, we also examined the signal quality over time for UEAs implanted in human participants as part of a brain-machine interface study (Collinger et al. 2013). We found that electrode yield increased slightly and SNR steadily increased in the short-term, saturating to some overall maximum value after approximately four weeks (Figure 9c,d). Due to the consistent and long-term nature of the human BCI study, signal quality metrics were available for human arrays for almost 1000 and 2000 days for subject P1

and P2, respectively. The electrode yield from these two human participants was generally higher and more consistent over time as compared to our NHP results though a slow decline in SNR was noted over time as in the NHP arrays.

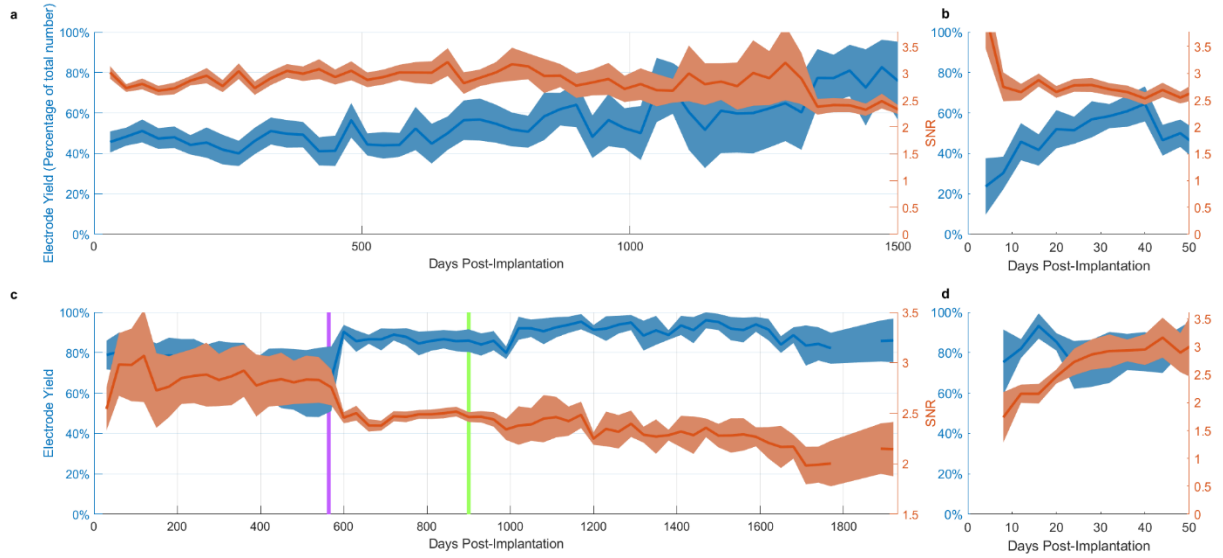


**Figure 9. Reliability of viable chronic recordings over time**. Shaded regions denote standard error of the mean. a. Average yield (blue) and SNR (red) over arrays with viable recordings (i.e. array recordings that were not terminated) versus number of days post-implantation for NHP implants. Full time range not shown (maximum lifetime is in excess of 3000 days). b. Close-up of average yield and SNR over the first 40 days post-implantation. c,d. equivalent statistics for human implants. Further details can be found in Downey et al., 2018 (Downey et al. 2018). The purple vertical line indicates the date at which the R.M.S. spike threshold was set to a different value during data collection (see methods for details). The change in the figure is due to a larger number of recordings becoming possible to average across, reducing variance in the estimate of yield and SNR. The green line indicates the date data collection with P1 ended.

### 3.4 Performance effects of electrode tip metallization and length

Platinum metallization of the electrode tips was adopted in the early fabrication of UEAs but was not well suited for microstimulation of the cortex. The option to use iridium oxide metallization was made available in 2009 which allowed for chronic stimulation due to its improved charge injection capacity, as well as recording (iridium oxide arrays were used for acute stimulation in our lab but were not used for chronic stimulation) (Negi et al. 2010). We

directly compared recording yield and SNR between these two metallization options and found that average array yield was significantly higher for iridium oxide as compared to platinum tips during chronic recordings in the intermediate time range although there was no difference in the short and long term (Figure 10a). In contrast, there was no significant difference in mean SNR between the two metallization materials at any time (Figure 10b).

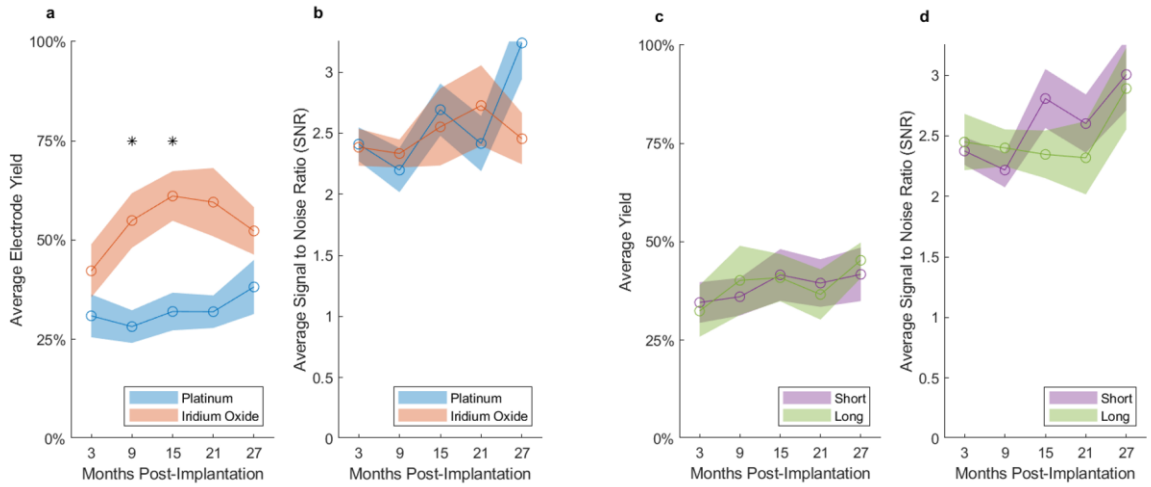


**Figure 10. Effects of electrode tip metallization and electrode length on array performance in NHP implants**. a, c. The average yield for each six-month date window was calculated for each array; an average was then taken across all arrays with the same metallization or electrode length (short-1.0 mm, long-1.5 mm), respectively. Shaded regions for all plots indicate the standard error of the mean, across arrays. b, d. Average SNR based on metallization and electrode length, respectively. Statistical tests were corrected for multiple comparisons using Bonferroni correction. Stars indicate statistically significant differences between groups, Bonferroni corrected. Averages were calculated over recordings and arrays for which data was available for each time window. Only NHP data were used for this figure.

We also compared recording performance between short (1.0 mm) and long (1.5 mm) electrodes. Given our research interests in motor and premotor cortices, our initial bias was to use long electrodes to target as close to layer five as possible in the thicker motor cortex. However, we were also concerned about the known observation that UEAs sink over the long term by compressing superficial layers, and so we also used short electrodes in some cases in order to prevent the tips from sinking into white matter particularly in the thinner premotor

cortex (unpublished histology). We found no difference in average yield or SNR between short and long electrodes (Figure 10c,d).

During our analysis, questions arose about the potential impact of manufacturer changes or surgeon experience over the entire range of dates during which implants were placed in subjects. To address these concerns, we analyzed the reliability of all NHP implants, based on their implantation date. Despite having been implanted by different surgeons, and having undergone largely minor changes in manufacturing process, we observed no significant relationship between array longevity and date of implant ($r = 0.11$, $p = 0.43$).

### 3.5 Maximum performance over array lifetime

We next addressed the maximum possible performance for these arrays over their lifetimes. We calculated the distribution of maximum yield over arrays and the time post-implantation when the maximum yield occurred. Despite the large variance in maximum yield across arrays, we found that 16 (out of 55) arrays exhibited a maximum yield exceeding 90% (Figure 11a). The maximum yield typically occurred within the first 150 days of a given array's recording life (Figure 11b).

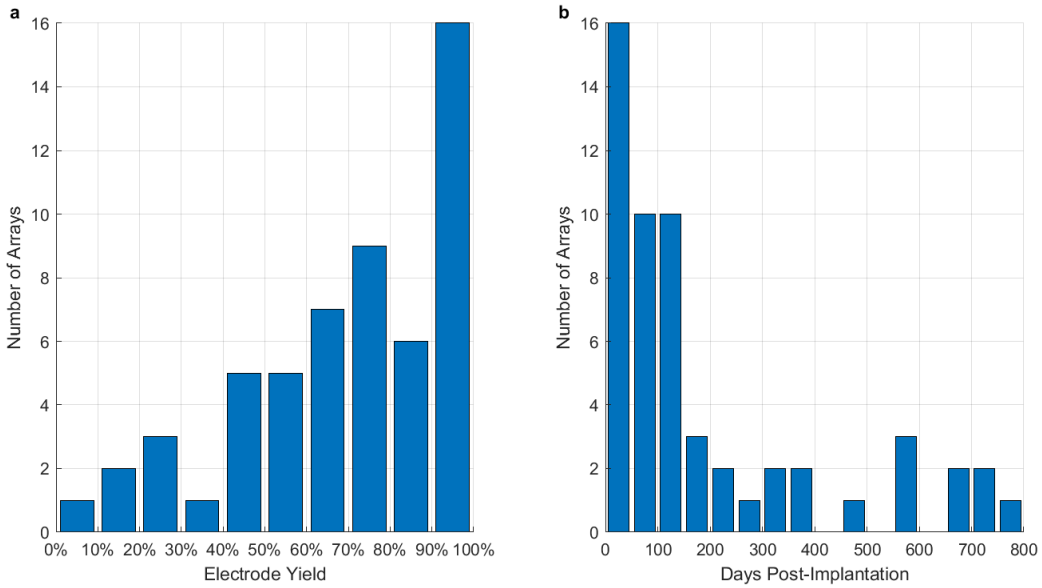


**Figure 11. Maximum possible yield**. a. Distribution of the maximum recorded electrode yield over all arrays, from 0% to 100% of possible electrodes. b. Distribution of dates when the maximum number of recorded channels occurred. Most maximum recording days occurred within the first 150 days post-implantation.

## 4. Discussion

UEAs are now a standard recording device for chronic electrophysiology in animals and is the only chronic intracortical recording technology that is FDA-cleared for investigational purposes in humans. Assessing the longevity and reliability of these devices is important for long-term experiments in NHPs and is of particular significance for chronic use in human BCIs. Most previous publications have examined the signal quality of Utah arrays using data from a limited number of subjects or explicitly focused on physical failure modes of the arrays (however, see (4) for one exception); other papers have examined the longevity and/or reliability of electrophysiological recording techniques, while not specifically focusing on the Utah Array (Rousche and Normann 1998; Suner et al. 2005; Chestek et al. 2011; Downey et al. 2018; Bullard et al. 2020; Fraser and Schwartz 2012; Kozai et al. 2015; Salatino et al. 2017; Buzsáki

2004; Polikov, Tresco, and Reichert 2005a). We found that nearly sixty percent of array implants exceeded 40% electrode yield after 6 months post implantation and nearly fifty percent of array implants exceeded 40% yield at one year. We also observed a subset of arrays (n=12) that exhibited moderate yield and SNR beyond 900 days post-implantation and one implant lasted nearly 9 years. We also characterized an important aspect of this recording technology, namely, a rapid increase in electrode yield within the first 40 days post-implantation in the non-human primate implants, most likely because of acute inflammation and subsequent recovery after surgery (Downey et al. 2018). Early recovery of signal quality post-implantation is likely due to alleviation of acute inflammation and potential small hematomas as swelling and bleeding reduce. The acute immune response to a foreign body is associated with activated microglia and astrocytes which begins at the moment of insertion and lasts approximately 6-8 weeks post-implantation followed by a chronic response associated with glial encapsulation of the foreign body (Polikov, Tresco, and Reichert 2005b). The improved yield in the first 40 days post-implantation may be related to the termination of the acute immune response whereas the long-term decrease in yield may be in part the result of gliotic encapsulation.

We also observed that the yield of human implants was higher than NHP implants. There are many potential reasons for this difference, despite the hardware itself being identical: Utah arrays approved for human implant undergo additional quality control checks prior to implant; surgical procedures are executed by neurosurgeons as opposed to an investigator with a scientific or engineering background; most importantly, human study participants and their caregivers take care of the connector and surrounding skin whereas monkeys can damage their own connector or the skin around the percutaneous connector pedestal leading to infection or failure of the implant.

The results directly reported here are limited to two human participants, but successful long term human implants have been commonly reported in the literature (Bullard et al. 2020)

Our findings also suggest that UEAs with iridium oxide metallized tips result in higher recording yield as compared to platinum metallization (in non-human primates). Iridium oxide also has the added benefit that chronic electrical stimulation is possible. However, these results should be interpreted with caution given the fact that our platinum implants were done earlier than the iridium oxide implants over the 17 years of implants examined in this study. The improved yield of the iridium oxide implants may have been due to improvements in surgical techniques as we learned better implantation methods. There have been a few improvements in the UEA manufacturing process over the 17-year period most of which were relatively minor with the exception of transitioning from array to wafer scale manufacturing in 2009-10 that may explain the improved yield of the iridium oxide implants (Blackrock Microsystems, n.d.). Nevertheless, when considering all our implants regardless of metallization, we found no correlation between array longevity and date of implant.

### 4.1 Species-specific differences in signal quality

Data from our human subjects tended to display higher array yields and SNR than the NHP data. However, the number of human datasets was vastly smaller (only four arrays and two subjects) than the datasets in the NHPs, and so it is difficult to draw any strong conclusions regarding differences in signal quality between humans and NHPs. One possibility may be that the surgical ability with which microelectrode arrays were implanted may have had a significant impact on the eventual efficacy and quality of long-term recordings in each subject. With each human implant, trained neurosurgeons and a team of surgical staff supported the careful

placement of UEAs; with each NHP implant, the process was conducted by PIs, postdocs, and graduate students. Although veterinarian and animal clinic personnel are sometimes available for NHP implantation procedures in some research institutions, the availability of these resources is far from guaranteed. An added factor that may have contributed to the human implants' higher signal quality and lifespan was more diligent care. NHPs have a higher likelihood of implant damage due to their activity and behavior, while human participants are generally more careful with their implants. Moreover, the wound margins surrounding the percutaneous connectors tend to become dirtier and more prone to infections in NHPs versus humans.

Although our ability to comment on the consistency of human implant reliability is limited due to our small sample size, previous work indicates that the duration of human implants can also vary (Bullard et al. 2020). Across eighteen chronic human implants (including the two participants included here), at least nine remained implanted for at least 1 year and seven of these for more than 2 years. Importantly, the authors noted that some of these experiments remain ongoing and that the total reported duration does not indicate that the array failed.

### 4.2 Study Limitations

One of our unique advantages in investigating questions concerning multi-electrode array signal quality is the massive amount of data available to us over the past 17 years. However, our intentions for implanting these arrays were to address scientific questions regarding cortical function and not to specifically address array longevity and reliability. Therefore, one limitation of this study is that we did not sample recordings evenly in time, and we sometimes terminated recording when we were finished with experiments instead of due to array failure.

Another limitation due to the size and scale of the datasets with which we analyzed was the lack of spike sorting to distinguish individual single units within a given electrode. As a result, all our metrics and analyses are on the scale of individual electrodes and arrays, instead of individual single units. Our decision to forgo spike sorting, therefore, underestimates unit yield and SNR because two high SNR single units of different amplitudes on a single electrode, for example, would result in a lower average SNR. It should be reiterated that meaningful information about the brain may be derived from channels which exhibit poorly isolated units.

The Utah Array continues to be a reliable and valuable tool for systems neuroscience and human brain-computer interfaces. We have demonstrated that most non-human primate implants will reliably last over a year, with the potential for arrays to last for the better part of a decade; human implants also display impressive longevity, lasting multiple years. While the Utah array may be acceptable for current animal research, more research and development must be done to create chronic implants that will reliably last into years and decades of use, for human applications and more advanced animal studies.

# B. Supplementary Figures

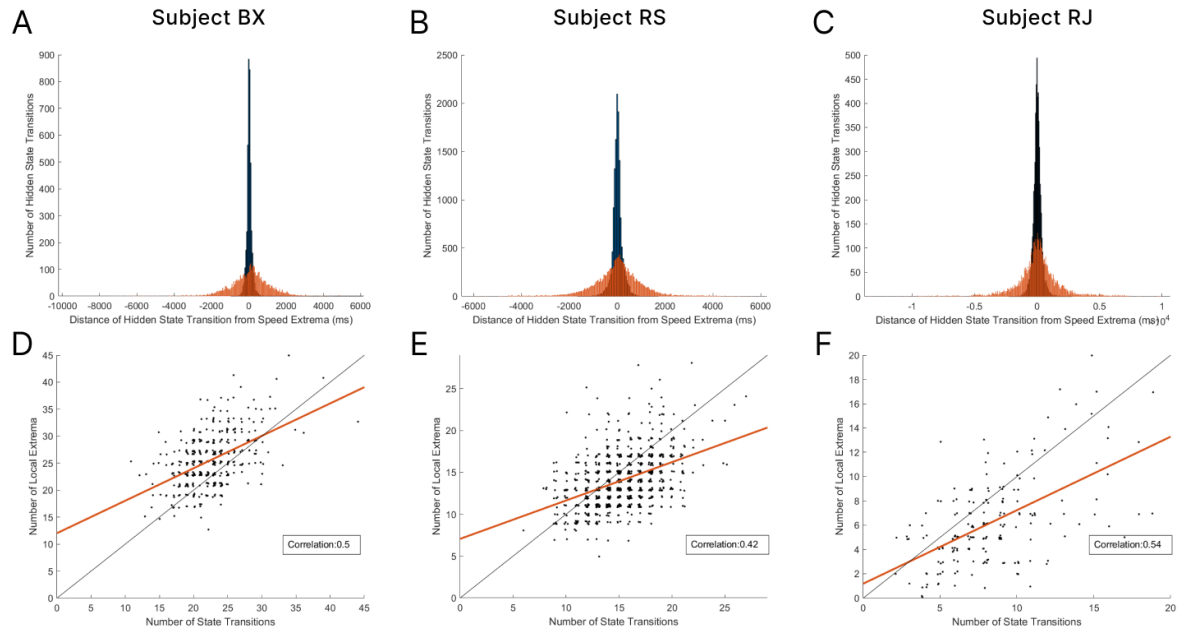


**Figure 12.** Transition and Extrema Comparisons. A,B,C. histograms showing time in ms from each speed extrema to the closest neural transition (black), and the time in ms from each extrema to a randomly sampled timepoint in the trial (orange). D,E,F. Scatter plots showing the number of per-trial speed extrema versus the number of discrete state transitions, with identity line (black) and linear regression line (red) and correlation strength (in text box).

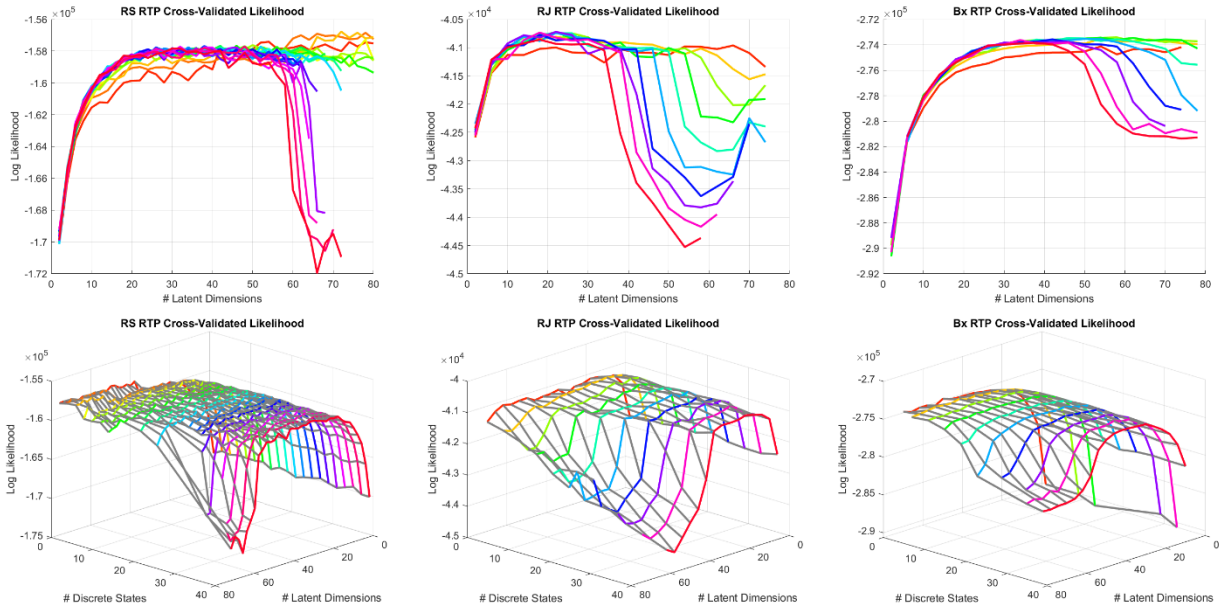


**Figure 13. Grid Search for Optimal Discrete state and Dimensions from rSLDS models.** Three-dimensional and two-dimensional representations of hyper-parameter grid search, plotting cross-validated log-likelihood.