Towards population coding principles in the primate premotor and parietal grasping network

Dissertation

For the awarding of the degree

"Doctor rerum naturalium"

of Georg-August-Universität Göttingen

within the doctoral program *Systems Neuroscience*
of the Georg-August University School of Science (GAUSS)

submitted by

Jonathan A. Michaels

from Toronto, Ontario, Canada

Göttingen, 2017

Doctoral Thesis Committee: Prof. Dr. med Hansjörg Scherberger (First Referee, Supervisor)
Research Group Neurobiology
German Primate Center
Kellnerweg 4
37077 Göttingen

Prof. Dr. Alexander Gail (Second Referee)
Cognitive Neuroscience Laboratory
German Primate Center
Kellnerweg 4
37077 Göttingen

Prof. Dr. Stefan Treue
Cognitive Neuroscience Laboratory
German Primate Center
Kellnerweg 4
37077 Göttingen

External Examiners: Prof. Dr. Siegrid Löwel
Systems Neuroscience Group
Johann-Friedrich-Blumenbach-Institute of Zoology and Anthropology
Bernstein Focus for Neurotechnology
Von-Siebold-Str. 4
37075 Göttingen

Prof. Dr. Florentin Wörgötter
Georg-August University Göttingen
Third Institute of Physics
Biophysics
Friedrich-Hund-Platz 1
37077 Göttingen

Prof. Dr. Tobias Moser
University Medical Center Göttingen
Institute for Auditory Neuroscience & InnerEarLab
Robert-Koch-Str. 40
37075 Göttingen

Date of oral examination: January 12th, 2016

Herewith I declare that I have written this thesis independently and with no other aids and sources other than quoted.

Göttingen, 2nd of January, 2017   Jonathan A. Michaels

For Lois

Acknowledgements

I'd like to thank Hans Scherberger, who brought me into his lab. He provided me with experiments, ideas, support, and constructive criticism when needed. Even more, he encouraged a positive lab atmosphere and empowered me with the creative freedom to find my own scientific way. Thanks also to the other members of my thesis committee, Alexander Gail and Stefan Treue, for their helpful comments over the years.

Thank you to the whole Neurobiology lab. To Natalie Nazarenus and Ricarda Ahlert for tireless animal training, and Matthias Dörge for tireless technical assistance and explanation. To Sebastian Lehmann for teaching me how to break as few electrodes as possible. To Benjamin Dann for countless constructive discussions and debates. To Rijk, Wei-an, Stefan, Jeroen, Roman, Clio, Valeska, Antonino, Rebecca, Philipp, Caio, Lukas, Josey, and 1,3,7-Trimethylpurine-2,6-dione, for continued support both scientific and otherwise.

Thank you to Patty, Selma, Sally, Zara, and Bart for contributing more than anyone else to these projects.

Thank you to my family. To Katharina, who went on this journey with me and continues to do so, bringing her endless support and love along at every step. To Howard, for instilling in me a love for science and knowledge that is still going strong, and to Lois, for reminding me that sometimes there are more important things than science and knowledge, and for supporting and encouraging me no matter what I chose to do.

Table of Contents

1 – GENERAL INTRODUCTION

"Population analyses are necessary in settings in which there may be neural mechanisms that involve coordination of responses across neurons. These mechanisms exist only at the level of the population and not at the level of single neurons, such that single-neuron responses can appear hopelessly confusing or, worse, can mislead the search for the true biological mechanism."

(Cunningham and Yu, 2014)

The cornerstone of psychology research is behavior. Although the ultimate aspirations of psychology are to tackle the inner-workings of the brain, the only sensible recourse for an early psychologist was to carefully measure the inputs and outputs of a system, make conclusions about the underlying computations, and make predictions based on these observations. This "black-box" approach to the nervous system was necessary at a time when the cellular structure of the brain was mysterious, and the technology did not permit more detailed study. Much of the earliest electrophysiological work, such as Sherrington's experiments in the canine reflex system (Sherrington, 1906), were undertaken and initially interpreted under this framework.

It was the birth of neuroscience that squarely focused attention inside the box. The birth of neuroscience did not precipitate out of the work of psychologists, but that of histologists and pathologists. In 1873, Camillo Golgi introduced the Golgi staining procedure (Golgi, 1873), allowing for visualization of individual neuronal bodies, axons, and dendrites. However, at this time the prevailing theory was 'reticular theory,' which posited that the nervous system was composed of an immense web of connectivity, termed a syncytium.

1.1 – Birth of the neuron doctrine

A few years later, Ramon y Cajal began experimenting with Golgi's technique and introduced many refinements, tuning the procedure to the specific tissue being investigated among other factors (Cajal, 1888). Based on his painstaking research and documentation of individual cells, Cajal proposed the neuron doctrine. He believed, although couldn't prove definitively, that the nervous system could be broken down into individual neurons distinct from one another. Furthermore, he postulated that the

function of these cells was to integrate the input of potentially thousands of other neurons and produce a self-regulated output. Although he was not the first person to propose neurons as the basic unit of the nervous system, and the physical gaps between individual neurons could not be proven until later (Porter et al., 1945), his neuron doctrine is still the prevailing theory today. In 1906, Cajal and Golgi shared a Nobel Prize for their many contributions. Together with the notion that individual neurons have specific receptive fields, i.e., areas of the skin or visual field that they respond to exclusively, these discoveries position the neuron as the most basic functional unit of the nervous system.

It is remarkable how closely linked the prevailing theories in the history of neuroscience are linked to the technological history. Indeed, this is a phenomenon that is true of many disciplines and has been pointed out a number of times (Dyson, 2012; Yuste, 2015). Sometimes ideas drive us to develop the technologies necessary to prove them, and sometimes, new technologies reveal basic facets of nature we never conceived.

1.2 – Mapping the brain: neuron by neuron

It was the discoveries of Cajal and Golgi that began neuroscience in earnest, but the explosion of neuroscience as we practice it today owes its state to the discoveries that allowed recording the activity of single neurons *in vivo*.

> *"An electrode has been developed to fill the need for an easily made, study device capable of resolving single-neuron action potentials at least as well as the commonly used micro-pipette."*
> (Hubel, 1957)

In no small part due to Hubel's development of the tungsten microelectrode, recording from single neurons during behavior became more practical and feasible than it ever was before. Implementing their techniques in cat visual cortex, Hubel and Wiesel were able to map out what kinds of stimuli excited individual cells in particular parts of the visual field (1962), and even developed a schematic of how such activity could arise in a network. Imagine identifying a neuron that seems to respond exclusively to a bar moving at a certain angle across a certain spot in an

animal's visual field. Or, imagine finding a neuron that seems to respond only to the presentation of your favorite breakfast cereal. Indeed, finding out what kinds of stimuli or behaviors reliably modulate individual neurons is the cornerstone of neuroscience research up to the present time. Yet, in a behaving animal, how can one be sure that the very same neuron does not participate in many processes? How does it respond to auditory stimuli? How does it behave when an animal imagines a similar stimulus? What is its function when the animal does nothing at all?

Herein lies the crux of the neuron doctrine. By identifying neurons in the brain that form not just functional units, but perceptual units, we very easily begin to pigeonhole individual neurons as players fulfilling singular roles that are imaginable by us. That is, it's easy to assume that because a neuron seems to care whether you stretch your arm or sit still; it must be causing that action and only that action. The discovery of specific receptive fields played perhaps the largest part in cementing this view, since under this framework neurons are presumed to be 'inactive' when their receptive field is not being addressed. Taking this point further, such a perspective additionally only considers the firing rate of a neuron itself, not when the individual spikes occur, or how they may be synchronized with other neurons (Gray et al., 1989; Fries, 2015).

Indeed, recordings in higher level brain areas, as well as areas closer to sensory inputs and motor outputs, have yielded highly heterogeneous and complex single-neuron responses, both across neurons and across experimental conditions. In some cases, single-neuron responses may bear no obvious moment-by-moment relationship with the sensory input or motor output that can be externally measured. Classically, such heterogeneity has been considered to be a result of biological noise or other confounds, and often researchers study only neurons that 'make sense' in terms of externally measurable quantities. However, this single-neuron complexity may be the realization of a coherent and testable neural mechanism that exists only at the level of the population."

(Cunningham and Yu, 2014)

As described above, the danger of viewing neurons as perceptual units is that during experiments, they tend to be selected based on how task-related they appear during recording sessions, thereby biasing recordings and obscuring the true

computational basis of a given area.

To give a more thorough exposition of how receptive field coding can be misleading, let's consider the visual and somatosensory systems. While basic processing such as edge detection, both visual and sensory, has been long considered to be cortical in origin, new research challenges the established hierarchy of processing. In the visual system, studies have found that feature extraction can occur in the retina itself (Gollisch and Meister, 2010), as early as the first step following light detection by the rods and cones (Venkataramani and Taylor, 2010). From another modality it is known that first-order tactile neurons in the periphery have axons that branch in the skin, forming a distributed systems of receptive fields (Paré et al., 2002), the consequences of which have been elusive. A recent study recorded from first-order tactile neurons in human subjects while a large set of point and edge stimuli were passed over the receptive fields of these neurons (Pruszynski and Johansson, 2014). The experiment revealed that the distributed nature of these receptive fields allowed first-order tactile neurons to signal edge orientation, both through magnitude of response and temporal code. Therefore, while categorizing neurons directly by receptive field allows us to simplify our perceptual understanding of the computations being undertaken in the nervous system, these interpretations may underestimate the sophistication of the underlying circuit. Furthermore, in higher-order brain areas in the cortex, where 'receptive fields' cannot be contained to individual areas of visual field or sensory patches, the concept of a clear receptive field becomes almost impossible to ascertain, as it would involve testing an endless number of stimuli and behavioral conditions.

1.3 – Mapping the brain: en masse

Although the arsenal of single electrode techniques has expanded to include tetrodes and laminar probes, the most significant development towards wide-scale recording is the implantable electrode array (Churchland et al., 2007; Kipke et al., 2008). The most widely used of these arrays is the so called Utah array, with 100 channels distributed on a grid (Nordhausen et al., 1996; Maynard et al., 1997). There are three immediate benefits of such a technology. Firstly, while recording with single electrodes requires many repeated recording sessions over the course of months, an array can record hundreds of individual neurons in a single session, meaning that data can be collected

more quickly, and therefore more experiments can be undertaken. Secondly, since the arrays are chronically implanted, it is often possible to obtain stable recording of single neuron for many hours, days, weeks, or in some cases months (Carmena et al., 2005; Chestek et al., 2007; Dickey et al., 2009; Chestek et al., 2011). Thirdly and crucially, neural activity during repeated trials of identical stimuli often yield varying neural and behavioral results (Churchland et al., 2006a; Ko et al., 2011; Churchland and Abbott, 2012), and arrays permit the capture of population dynamics during single trials. How do neurons covariate during a trial? How is the timing of spikes coordinated across a population (Okun et al., 2015)? These are questions that can only be properly answered using parallel recording. Along with the advances in implantable array technology, much new research employs advanced 3D imaging techniques to visualize large populations of neurons simultaneously (Schrödel et al., 2013; Prevedel et al., 2014), techniques which will no doubt become more and more relevant, especially in the pursuit of relating functional connectivity to anatomical connectivity.

It is also worthwhile to note how the development of these new technologies also biases us towards particular types of analysis in the same manner as single electrode recordings have biased us for many decades. While arrays tilt the scales in favor of population-based analysis, they also bias us away from other analyses. For example, understanding the layer structure of cortical areas in vivo requires detailed acquisition of depth information, for example using laminar probes with many contacts all down the shaft. The kinds of mass arrays that are in use tend to make ascertaining the identity of different layers of cortex virtually impossible and is therefore not emphasized or even attempted in most experiments employing array recordings. Furthermore, distinguishing different cell types in array recordings has been used controversially in recent studies. Some believe that excitatory neurons and inhibitory interneurons can be distinguished based on their spike-waveform shape (Kaufman et al., 2010; 2013), but proving a direct connection between spike waveform and cell class is an ongoing quest (Vigneswaran et al., 2011).

1.4 – Dimensionality reduction: rotating the perspective

Implantable arrays provide the necessary data complexity for elucidating how neural networks operate in real-time. However, they also produce an overabundance of data

to interpret (Sejnowski et al., 2014). Therefore, to go along with advanced recording techniques we require sophisticated analysis techniques (Stevenson and Kording, 2011). Dimensionality reduction methods fill this requirement by summarizing and transforming high-dimensional signals into a realm that can be interpreted, analyzed, or utilized as a control signal. The many types of dimensionality reduction used in neuroscience research, and when they should be employed, has been reviewed in detail (Cunningham and Yu, 2014).

Dimensionality reduction is useful any time one suspects that the number of available signals is much greater than the number of informative latent dimensions, i.e., when redundancy exists between independently collected signals. Neural data often fits this description as neurons in the brain, especially within one area, are not expected to be completely independent given the connectivity between them. Therefore, dimensionality reduction techniques seek to find a lower-dimensional representation of high-dimensional data that captures certain aspects of the data. The specific objective varies from method to method and determines what parts of the signal are conserved in this low-dimensional representation. For example, data compression techniques by definition employ dimensionality reduction. In image compression, neighboring pixels in a natural image are unlikely to have wildly differing colors or intensities. Therefore, image compression finds a low-dimensional transformation of an image that preserves the unpredictable or visually relevant changes in an image while discarding others, taking up less digital storage space. The image can then be transmitted to another user and reconstructed from this reduced representation (note: in the case of image and sound compression the basis functions for reconstructing a signal are predetermined, and it is how these dimensions are combined that is computed in the compression https://en.wikipedia.org/wiki/Discrete_cosine_transform, accessed 18.11.2015).

As mentioned before, the primary difference between dimensionality reduction methods is the objective. The most common techniques, principal component analysis (PCA) and factor analysis (FA), have the objective of explaining covariance between dimensions. In the case of PCA, orthogonal dimensions are found (eigenvectors of covariance matrix) that can be linearly combined to optimally reconstruct the high-dimensional signal. These methods only consider overall covariance, not temporal dynamics. Other methods, such as Gaussian-process factor analysis (GPFA) (Yu et al., 2009) and the widely used Kalman filter (Harvey, 1990),

take into account how signals change from time point to time point, thereby emphasizing the dynamics of a signal. Recently, toolboxes have become available making many types of dimensionality reduction accessible to a wide audience of scientists (Cowley et al., 2013).

From a basic array de-noising perspective, PCA can be used for artifact elimination (Musial et al., 2002). Since arrays have many channels distributed over a large area of brain tissue, it is unlikely that a biological signal would appear simultaneously on all channels. Therefore, PCA dimensions that represent signals present on all channels with no time difference are likely artifacts and can be subtracted from each channel prior to spike sorting to improve signal quality.

Dimensionality reduction that takes dynamics into account is commonly used to reduce a population of simultaneously recorded units into a kinematic output for a neuroprosthetic application, i.e., controlling a cursor on a screen or a robotic arm. Many studies have successfully developed and applied these methods in a research and clinical setting (Velliste et al., 2008; Collinger et al., 2012; Hochberg et al., 2012; Aflalo et al., 2015). However, these types of dimensionality reduction have limitations with respect to basic research. In some cases these methods are required to assume certain types of tuning within the model (Georgopoulos et al., 1986), and in all cases they map neural activity directly onto movement parameters. The drawback of projecting onto movement parameters is the assumption that an internal population of neurons follows the dynamics of movement variables as opposed to an abstract representation determined by the local circuit (Churchland et al., 2007). Therefore, to reveal the underlying processes in neural populations, unsupervised dimensionality reduction that extracts latent dimensions based on internal covariance, such as PCA, is required.

The prospect of recording from large neural populations also raises an important practical concern. How many neurons do we need to record in order to understand neural circuit dynamics, as well as the resulting cognition and behavior? One of the most universal findings in neuroscience experiments employing dimensionality reduction is that the number of latent dimensions is far less than the number of recorded neurons (Gao and Ganguli, 2015). Do these underlying latent dimensions paint an accurate picture of the internal processing going on during a particular task, or would the number of dimensions increase when recording a larger pool of neurons? Some effort has been made to develop a theory of neuronal task

complexity for predicting the interplay between number of recorded neurons, quality of recording signal, and behavioral demands of a task. Using this theory, Gao and Ganguli (2015) were able to show that under some task designs, increasing the number of recorded neurons cannot provide a richer picture of internal processes, while on the other hand, a sufficiently rich task design does not necessarily require the acquisition of many more neural signals than is currently being collected. Such theories of data collection will be essential in the coming years as we hone experimental design to very specific questions, and analysis methods towards single-trial investigation.

1.5 – Dimensionality reduction: insights and pitfalls

We have discussed the advantages of large populations of neurons being projected down onto a much smaller amount of latent dimensions during certain behaviors, but do these dimensions have an intrinsic meaning? That is to say, can the apparent functional interpretation of these dimensions elucidate the underlying computations of a particular brain region or circuit, or are they measurement artifacts of the mathematical methods used to extract them?

As a cautionary tale, let us examine the classical genetics paper of Menozzi et al. (1978), published in the journal *Science*. In this and other studies, the authors generated allele-frequency maps of specific genetic variants over large geographical regions by collecting genetic samples from many individuals. Then, to make this data palatable, they reduced the dimensionality of this data using PCA to form synthetic spatial maps. These synthetic maps revealed extremely interesting patterns of apparent human migration. However, as was pointed out many years later (Novembre and Stephens, 2008), these patterns arise even if control data is simulated that contains no heterogeneous migration patterns. In other words, purely mathematical artifacts could explain the observed patterns. All types of dimensionality reduction techniques have inherent pitfalls requiring proper controls.

However, it would be inappropriate to describe dimensionality reduction pitfalls without providing positive examples. A prime example of this comes from the neuroprosthetic experiment of Sadtler et al. (2014). In this experiment, monkeys controlled a cursor on a screen in two dimensions using recordings from a large population of neurons in primary motor cortex. First, population activity was linearly

combined and reduced to only 10 dimensions using a method similar to FA, which finds dimensions that capture the inherent covariance between neurons, similar to PCA. Then, these dimensions were reduced again to produce two-dimensional kinematics that controlled the cursor. Monkeys were able to use this interface to move the cursor to targets laid out in a circle on the screen with a very high degree of accuracy. Following this, perturbations were made to the transformation matrices, either between the full population and the latent set of 10 dimensions, or between those and the kinematics. When the perturbation was made between these 10 dimensions and the kinematics, monkey performance dropped, but very quickly returned to perfect performance. However, when the perturbation was made in the transformation initially extracted using FA, monkeys were unable to improve their performance within a reasonable time scale. This study shows that the mechanisms underlying plasticity during closed-loop prosthetic control are readily able to recombine the latent dimensions extracted in the first step in order to produce meaningful kinematics. In other words, neurons that tend to covariate together are unlikely to de-couple their firing patterns in the short term, but as a group the magnitude of their firing can be adjusted by internal plasticity mechanisms to reduce output error (Martinez and Wang, 2015).

Another interesting experiment asked the question of how activity in primary and premotor cortex, which have direct connections to motor neurons in the spinal cord, could produce muscle activation during movement, but not during preparation for movement (Kaufman et al., 2014). It has been postulated that a non-linear gating mechanism could inhibit the effect of this output on muscle during preparation (Benjamin et al., 2010), as has been found for eye movements (Evinger et al., 1982), but no such mechanism could be detected for arm movements so far. In this study, it was found that the latent dimensions explaining neural population activity in premotor and primary motor cortex during movement could be linearly combined to produce the latent dimensions explaining muscle activation patterns during reaching. This linear transformation is meant to approximate how cortical output might synapse onto motor units in the spinal cord. Interestingly, they then showed that the neural activity during preparation for movement fell into the 'null-space' of this output transformation, suggesting that neural activity during preparation may 'avoid' population states that might produce unwanted muscle activity prematurely. This explanation was able to explain a large amount of variance in preparatory activity and

provides an explanation of how movement gating can be achieved in a linear fashion. Furthermore, they found evidence that the same principles may be at play between premotor and primary motor cortex, suggesting that connected pools of neurons could potentially selectively avoid or inhabit output-potent and output-null dimensions to selectively transfer or gate information.

Another area of research where parallel recording and dimensionality reduction is becoming essential is decision making.

"If the neural activity is not a direct function of externally measurable or controllable variables (for example, if activity is more a reflection of internal processing than stimulus drive or measurable behavior), the time course of neural responses may differ substantially on nominally identical trials."

(Cunningham and Yu, 2014)

As one progresses from large sensory driven areas to high-level association areas, consistency of processing between trials may decrease. Under these circumstances, averaging activity over many repeated trials of the same behavior can be misleading. Examples of this issue can be seen in the constant debate between certain decision processes represented as either accumulated evidence or spontaneous state shifts (Hanks et al., 2015). In an attempt to observe the decision making process in real-time, Kaufman et al. (2015) trained monkeys to reach to different targets on a touch screen while recording from primary motor and premotor cortex using multiple arrays. On some trials, multiple reach opportunities were presented to the monkey at different times, and with varying difficulty, occasionally giving the animal reason to change its mind. Using the GPFA dimensionality reduction technique, they were able to extract smooth neural trajectories on single-trials, and subsequently using the trajectories of these trials during deliberation to make inference about the internal decision making process of the animals. They found that monkey's decisions were often visibly reflected in the neural trajectory, both under forced conditions and spontaneously, revealing for the first time moment-by-moment correlates of the internal decision process. However, since behavior normally lies along a continuum of possibilities, distinguishing changes in internal processing from random fluctuations can be very difficult. The ultimate goal of such research should be not only interpretation of single-trial intention, which can be a subjective matter, but

direct manipulation of output behavior through inactivation (Erlich et al., 2015) or stimulation (Inoue et al., 2015).

As a final example, there exist some more targeted types of dimensionality reduction that combine unsupervised techniques with *a priori* information. Demixed principal component analysis (dPCA) seeks to explain existing covariance between neurons while also separating information about different task dimensions into components (Machens, 2010; Machens et al., 2010; Brendel et al., 2011; Kobak et al., 2014). Using dPCA, one can separate stimulus effects from choice representation in large populations of neurons, as well as information about how signals change over time, and compare how the population of neurons contribute to each of these components. The extracted task-specific dimensions form the interpretable perceptual units of the network, much in the same fashion as single neurons were treated as perceptual units in the classical studies mentioned earlier. Such techniques also raise questions of how individual neurons contribute to the representation of many different task features, even across modalities, known as mixed selectivity. Mixed selectivity in individual neurons is something which appears to be essential and widespread in prefrontal (Mante et al., 2013; Rigotti et al., 2013) and parietal (Raposo et al., 2014) cortex, and is part of the concept that individual neurons participate in many different distributed networks for processing a multitude of stimuli and actions.

Even Mountcastle, who proposed the cortical column as the functional unit of cortex, and Hebb (1949) before him, recognized that the appearance of distinct functional units could be deceiving (Hawkins and Blakeslee, 2007). Single cells cannot be assigned perceptual identities based on how they respond to a single set of stimuli. A strong response to one stimulus does not preclude the participation of that cell in many distributed and variable processes.

"It is obvious that the total number of distributed systems within the brain is much larger than had once been thought, and perhaps by several orders of magnitude. Thus major entities are parts of many distributed systems, contributing to each a property determined for the entity by those connections common to all of its modular subsets and by the particular quality of their intrinsic processing. Even a single module of such an entity may be a member of several (though not many) distributed systems."

(Edelman and Mountcastle, 1978)

1.6 – Modeling neural networks

Understanding how a specific neuron contributes to an underlying computation or what kinds of computations a specific circuit can carry out is difficult to study directly *in vivo*. One would require knowledge of the whole neural network involved, full information regarding anatomical connectivity, and the ability to manipulate the system. Although some new techniques permit some of these aspects (Peron et al., 2015), modeling is an essential tool to bridge the gap between theory and empirical data.

One of the most relevant methods for modeling the dynamics of a neural circuit in recent years is the recurrent neural network (RNN) (Sussillo, 2014). RNNs generally simulate individual neurons quite simply, modeling only firing rate and not individual spikes. Each neuron has an internal activation, which is related to its firing rate by some non-linear function, and has a set of inputs and outputs to and from other neurons in the network. The power of RNNs is that they allow complex temporal interactions between neurons, producing complex output even in the absence of input, similar to the brain. On the other hand, feed-forward neural networks, which are used to solve many complex problems, produce no output when no input is present.

There are a number of essential similarities between RNNs and biological systems. Firstly, neurons are non-linear, i.e., internal activation is non-linearly related to output firing rate (related through an inverse tangent function, for example). Secondly, neurons have strong feedback connections, a hallmark of cortical circuits. Thirdly, neurons themselves are modeled very simply, so completing complex tasks requires coordination over a large population. RNNs can approximate any dynamical system given enough internal units (Doya, 1993), so their power is not limited to biological modeling.

However, producing an endless stream of complex models does not directly allow us to understand the principles of neural computation. One of the most powerful aspects of modeling studies is to compare the space of all possible models that successfully reproduce a given empirical observation and those solutions observed in nature. For example, a study of a particular model of the yeast cell-cycle (Li et al., 2004) showed that although many models could be generated that 'fit' the data, only 3% of those models outperformed the empirical results when it came to robustness (Lau et al., 2007). This result provides strong evidence that the type of model that

evolved over time in nature seems to follow principles of robustness that could not have been identified without exploring the space of potential models.

In this way, we can sift through the endless pool of possible models and find those that are most likely informative about brain circuits and processes. For the next example, let's take the field of center-out reaching, which has been very widely studied in neuroscience. Many classical studies have found that individual neurons in primary motor cortex were cosine-tuned for specific reach directions (Schwartz et al., 1988; Caminiti et al., 1990), and that in the population these preferred directions were generally uniformly distributed throughout the field of possible movements. The robustness and uniformly distributed nature of these preferred directions has been used to conclude that motor cortex is coding high-level features of reaching, such as intended movement direction, rather than muscle activity, and has been leveraged in many decoding efforts such as the population vector (Georgopoulos et al., 1986). However, it has been shown that during two dimensional planar reaching, where movements can be more easily related to specific muscles, that the preferred directions of primary motor cortex neurons is not uniformly distributed, but biased based on the biomechanics of the limbs (Scott, 2000; Scott et al., 2001). Going on from this point, Lillicrap and Scott (2013) used modeling of this experiment to find the key ingredients that allow a model to perform like empirical data. In this study, they created a feed-forward neural network with state feedback in closed-loop to generate muscle activation pattern that could control an arm in two dimensions to complete center-out reaching movements. This model does not fit into the class of RNNs described above, since the network was feed-forward, but the state feedback provided a large degree of temporal interaction between read-outs and future states of the network. They showed that such a neural model reproduced the behavior as well as the non-uniform distribution of preferred directions within the internal neurons of the network, replicating the empirical findings. Crucially, they were able to directly test which aspects of their model were necessary to reproduce empirical results. Specifically, removing all bi-articulate muscles, i.e., those that span more than one joint (in the model with postural loads), produced a uniform preferred direction distribution. Therefore, they were able to step-wise test which ingredients of the model were essential. These types of studies are becoming more common in recent years and are vital to the field of motor control theory (Suminski et al., 2015), as many of these concepts still remain controversial (Naselaris et al., 2006;

Georgopoulos et al., 2007; Kurtzer and Herter, 2007; Tanaka, 2015), to the extent that cross-referencing between different laboratories has been minimized (Georgopoulos and Carpenter, 2015).

The study just described required not only solving a control problem with a neural network, but looking into the strategy of the network itself and making inferences. The same way the mind was treated as an impenetrable black-box for much of psychological history, so have RNNs. Although RNNs have been used to solve all manner of problems, the internal strategy and key ingredients of these solutions have not been analyzed in great detail. New techniques have made opening this black-box possible (Sussillo and Barak, 2013).

In the work of Mante et al. (2013), neural activity was recorded from the frontal eye field while monkeys performed a context-dependent motion or color discrimination task. Monkeys were presented with a set of moving dot patterns and had to selectively discriminate color or motion from the same visual stimuli based on a contextual cue. They then reported their response with a saccade, and were able to do this successfully. Neurons in this area show a wide range of mixed selectivity to the color and motion stimuli, as well as the performed saccade. So the question arose if the area in question could selectively process either the color or motion information depending on the context, without having to receive differing visual input? To test this, they simulated an RNN to execute the same task, i.e., select a saccade direction based on the sensory stimuli and the context. They found that like the recorded data, color and motion information were represented in parallel and orthogonally in the population of RNN units, while choice was represented independent of modality. Furthermore, by analyzing the fixed points of the system, they found that the temporal dynamics could be described by a simple combination of line attractors and selection vectors, mechanisms that arise directly from the learned connectivity within the RNN, and revealing a mechanism that is likely employed in the brain during context-dependent computations.

Another recent study showed how an RNN could produce the temporal dynamics matching muscle activation during reaching (Sussillo et al., 2015). While the ability of RNNs to produce complex temporal patterns similar to muscle activation has been shown before (Hennequin et al., 2014), this study went further by comparing the activity in the internal neurons in the RNN to neural activity recorded from the brain. They showed that when internal activity was constrained to progress in a

biologically plausible manner; it matched the patterns of activity seen in empirical data very well. Furthermore, strong rotational dynamics underlay both the real data and the model, but not a similar model with complex internal activity, i.e., activity that was not constrained in a biologically plausible manner. Perhaps most interestingly, when perturbations were introduced, only the biological RNN was robust to these perturbations, while the complex model quickly degraded. The power of RNNs is tightly coupled to the simplicity of the model, which allows mapping properties such as robustness directly to certain aspects of the model. Although the simplicity of the individual neurons modeled in RNNs are appealing, much research is still focused around finding the limits of modeling individual neurons in spiking neural networks (Hawkins and Ahmad, 2015; Markram et al., 2015). While these spiking neuron models are essential in the pursuit of whole-brain modeling, using simplified neural models allows us to examine the key ingredients of network dynamics without making too many assumptions (O'Leary et al., 2015), and in many cases the simple models capture the essential aspects of neural computation (Churchland and Abbott, 2012; Litwin-Kumar and Doiron, 2012),

1.7 – The grasping circuit

Up to this point, we've spoken generally about the brain and the technologies that allow us to peer inside. Let us bring our previous discussions to bear on an essential aspect of primate behavior, grasping. From the moment I see my cup of coffee, to the moment I drink the hot coffee out of the mug, lie a multitude of complex processing steps and actions. The shape of the cup, its orientation, its location in space relative to my eyes and my hand, the hand shape required to grasp it, the muscle activation required to move the arm and hand, as well as the holding of the cup and motion towards the mouth, all must be determined and regulated by the brain. Furthermore, all of this must be undertaken in the right context, depending on whether the cup itself is hot or cold, full or half-full. As primates, reaching out, grasping, and manipulating objects are some of our most essential and unique skills.

A few specific parts of the primate cortex are particularly remarkable in the quest to understand grasping. The anterior intraparietal area in the parietal lobe, also know as AIP, and the hand area of the ventral premotor cortex (PMv) in the frontal lobe, also know as F5, together form an essential circuit for the preparation and

execution of grasping movements in macaque monkeys and humans. These areas share extensive anatomical connections (Luppino et al., 1999), forming a long-range and very specialized circuit. Lesion studies in monkeys have allowed us to get an initial impression of how areas participate in actions. Reversible inactivation of AIP results in a selective deficit in pre-shaping the hand during grasping (Gallese et al., 1994), and the same is true for F5 (Fogassi et al., 2001). From many studies over the years (Janssen and Scherberger, 2015) we know that AIP and F5 are both involved during visual fixation of an object and manipulation of objects, even when no movement is required (Murata et al., 1997; 2000). Furthermore, in both areas, activity persists during delays between object presentation and movement, even when the object must be remembered and grasped in the dark (Murata et al., 1996; Baumann et al., 2009; Fluet et al., 2010). While F5 is thought to coordinate grasping as it happens, the role of AIP in online control is still controversial (Tunik et al., 2005; Schettino et al., 2015). Despite decades of research into both areas, differentiating their functional underpinnings has been elusive. As mentioned earlier, primary motor and premotor cortex have been sites for prosthetic application, but AIP has also been considered as a potential site for prosthetics (Menz et al., 2015; Schaffelhofer et al., 2015) and nearby areas of parietal cortex have recently been exploited in human patients (Aflalo et al., 2015).

   In order to determine the computations of this circuit, we can look to current working models of the AIP/F5 circuit. The most comprehensive model is that of Fagg and Arbib (1998), termed the FARS (Fagg/Arbib/Rizzolatti/Sakata) model. In this model, the authors simulated simple cell units consisting of leaky integrator neurons and linear threshold units that are triggered by certain visual stimuli. In this model only AIP and F5 are explicitly modeled, so visual information and context information arrives completely processed. Although this model is rather simple, it produces firing rates of individual units that were very similar to the empirical data existing at the time. In addition to matching experimental data, the model included a number of hypotheses. The primary hypothesis was that the main function of AIP is to extract "affordances," i.e., provide information about the features of an object relevant for grasping, probably from connections to inferotemporal cortex (Webster et al., 1994; Borra et al., 2008). On the other hand, F5 then selects only the desired grasp plan, based on information from prefrontal areas regarding the context of the action. Other hypotheses proposed that AIP maintains an active memory of the possible

affordances during execution, and that both areas should contain mixed information about objects and grasping. This second hypothesis has gathered substantial evidence, since neurons in AIP retain information about the orientation of an object often for an entire trial and do not change their preferences between cuing and movement, suggesting that they are maintaining a working memory of the object properties (Baumann et al., 2009). In addition to this, the idea that both visual and motor type activity exists in AIP is supported by a number of studies. It appears that object identity information likely is communicated to the posterior portion of AIP, while the anterior portion of AIP is more connected to F5, suggesting that a very large amount of the visuo-motor transformation process likely takes place within the gradient of AIP itself (Baumann et al., 2009; Michaels et al., 2015; Premereur et al., 2015).

Along with these hypotheses, a number of predictions were made as well. One prediction is that if an object could be grasped in two ways, but the specific way to be used is instructed after a delay, both areas should first represent both plans and then reduce to a single plan after the second instruction (Cisek and Kalaska, 2002). Indeed, if the object is presented first, but the particular grip not indicated, both options seem to become active in AIP, suggesting that both affordances are prepared. On the other hand, if grip information is presented before the object to be grasped is made visible, almost no differential tuning exists (Baumann et al., 2009), suggesting that without any visible object features, the context cues have no meaning in AIP.

Another prediction is that if two visually different objects are presented that are grasped identically, parts of AIP should distinguish these cases, and parts not, while F5 should mostly consider these objects identical. Furthermore, AIP should show the largest differences for aspects of objects that determine how they are grasped, regardless of the magnitude of visual differences. In a recent experiment, monkeys were presented with up to 50 different objects to be grasped in the dark, while activity was recorded from arrays in AIP and F5 along with kinematic recording from 27 joint angles in the shoulder, arm, and hand (Schaffelhofer, 2014). Interestingly, the similarity between the neural representation of these objects in the population of neurons in F5 closely matched the kinematic similarity between the hand shapes used to grasp the objects, implicating coding in a grasp-relevant manner. On the other hand, the population of neurons in AIP more closely represented the visual similarities the objects. As a test of the prediction above, a set of unique objects was specifically designed that was grasped identically regardless of difference in

appearance. Neurons in AIP much more readily indicated differences in these objects, while this was not the case in F5. Together, these findings lend support to the FARS model, which had predicted that AIP should generate responses that matched the affordances of each object; while F5 should represent only the selected grasp itself.

If AIP is encoding grasp affordances, then why is tuning for an instructed grasp type, in cases where the object to be grasped is visually identical, so high in AIP (Baumann et al., 2009; Lehmann and Scherberger, 2013)? There are two main possibilities. The first is that AIP is more connected directly to prefrontal cortex than originally thought (Petrides and Pandya, 1984; Borra et al., 2008), and therefore selection of appropriate grasp could potentially be undertaken in AIP itself, sending on to F5 only the grasp selected (Rizzolatti and Luppino, 2001). However, recent evidence suggests that when the monkey is free to choose which grasp to use, this information is first represented in F5 (Dann and Scherberger, unpublished data). The second possibility is that since the monkeys are typically trained in these types of experiments over tens of thousands of trials, the cues indicating the type of grasp to use may have been incorporated into the representation of the grasping object itself, thereby determining the object affordances at an early stage in visual processing (Bonini et al., 2012).

Although possible models have been discussed in detail in the last decade (Rizzolatti and Luppino, 2001; Cisek, 2007; Cisek and Kalaska, 2010; Grafton, 2010), the FARS model is still the most comprehensive computational model to date that directly simulates firing rates of cell populations in AIP and F5. Why has the AIP/F5 circuit resisted more sophisticated modeling? The visual circuit for object identification has been modeled extensively using many methods including deep neural networks (DiCarlo et al., 2012; Pagan et al., 2013; Kriegeskorte, 2015). These models have been able to shed light on the processing steps likely taking place in each sub-area along this stream. One of the likely reasons why this has not been the case in the AIP/F5 circuit is that the inputs and output of the circuit are not clear-cut. Both areas receive input from many areas and participate in the planning and execution of movements, making reducing each area to singular processing functions unreasonable.

However, there is another reason why modeling of these areas will likely always be a challenge. Both these areas are essential participants in action understanding. Understanding the actions and intentions of others is an essential

component of human social behavior. Deficits in social cognition have been associated with prevalent mental disorders such as autism (Frith and Frith, 2010). A recently discovered neural substrate that might be crucial for understanding the action of others is the mirror neuron system (MNS) (Gallese et al., 1996; Rizzolatti et al., 1996a). Mirror neurons in macaque monkey premotor cortex fire during the execution of specific goal-oriented behaviors such as object grasping, holding, and manipulating, as well as the observation of similar actions by humans or other monkeys. Mirror neurons even respond if part of their preferred action is visually occluded (Umiltà et al., 2001), or a sound related to the action is heard (Kohler et al., 2002). A number of human studies employing non-invasive brain imaging techniques have found evidence suggesting that a similar network exists in humans (Grafton et al., 1996; Rizzolatti et al., 1996b; Binkofski et al., 1999; Buccino et al., 2001; Grèzes et al., 2003). Mirror neurons with motor properties have been identified in area F5 along with areas in and around the intraparietal sulcus, especially areas PFG and AIP (Rizzolatti and Sinigaglia, 2010). The MNS has been linked to imitation, motor learning, empathy, and deficits caused by autism (Iacoboni and Dapretto, 2006). The circuit not only permits the automatic preparation of potential grasping movements, but also likely underlies our understanding of others and is a substrate for learning. Therefore, convincing modeling of these areas requires not only the simulation of a successful grasp, but also many of the most complex behaviors we undertake as humans.

"Rather than focusing on the question of 'what is represented' by a particular area or set of neurons, it may be more productive to address this issue by understanding the causal role of the activity of these neurons: how the activity is decoded or read out by the downstream areas and eventually used for behavior."

(Murakami and Mainen, 2015)

2 – ORIGINAL ARTICLES AND MANUSCRIPTS

This chapter contains the following research articles and manuscripts:

2.1 Predicting reaction time from the neural state space of the premotor and parietal grasping network

Michaels JA, Dann B, Intveld RW, Scherberger H (2015) Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. *Journal of Neuroscience 35*:11415–11432. doi:10.1523/JNEUROSCI.1714-15.2015. (Michaels et al., 2015)

Author contributions: J.A.M., B.D., R.W.I., and H.S. designed and performed research; J.A.M. and B.D. analyzed the data; J.A.M. wrote the paper. All authors revised the manuscript.

2.2 Probing the continuum of immediate to withheld grasping movements in the macaque fronto-parietal network

Michaels JA, Dann B, Intveld RW, Scherberger H (in preparation).

Author contributions: J.A.M., B.D., R.W.I., and H.S. designed and performed research; J.A.M. analyzed the data; J.A.M. wrote the paper. All authors revised the manuscript.

2.3 Neural population dynamics during reaching are better explained by a dynamical system than representational tuning

Michaels JA, Dann B, Scherberger H (2016). Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLoS Computational Biology*, *12*(11), e1005175. http://doi.org/10.1371/journal.pcbi.1005175. (Michaels et al., 2016)

Author contributions: J.A.M., B.D., and H.S. designed research; J.A.M. performed research. J.A.M. and B.D. analyzed the data. J.A.M wrote the paper. All authors revised the manuscript.

2.4 Hand-invariant to hand-specific representation of grasping movements in the macaque fronto-parietal network

Michaels JA, Scherberger H (in preparation).

Author contributions: J.A.M. and H.S. designed and performed research; J.A.M. analyzed the data; J.A.M. wrote the paper. All authors revised the manuscript.

2.1 – Predicting reaction time from the neural state space of the premotor and parietal grasping network

Authors:

Jonathan A. Michaels
Benjamin Dann
Rijk W. Intveld
Hansjörg Scherberger

"In 1991, Leroy Burrell set a world record for the 100 m dash with a spectacular time of 9.90s, stunning the prerace favorite Carl Lewis, who finished second with a time of 9.93s. It was later noted, however, that Burrell was not the faster runner. Rather, his reaction time to the gun that marked the start of the race was much shorter than Lewis's: a hair-trigger 117 ms against a relatively lethargic 166 ms. Without this difference, Lewis would have won handily."

(Afshar et al., 2011)

Systems/Circuits

Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network

Jonathan A. Michaels,[1] Benjamin Dann,[1] Rijk W. Intveld,[1] and Hansjörg Scherberger[1,2]

[1]German Primate Center, D-37077 Göttingen, Germany, and [2]Faculty of Biology, Georg August University Göttingen, D-37073 Göttingen, Germany

Neural networks of the brain involved in the planning and execution of grasping movements are not fully understood. The network formed by macaque anterior intraparietal area (AIP) and hand area (F5) of the ventral premotor cortex is implicated strongly in the generation of grasping movements. However, the differential role of each area in this frontoparietal network is unclear. We recorded spiking activity from many electrodes in parallel in AIP and F5 while three macaque monkeys (*Macaca mulatta*) performed a delayed grasping task. By analyzing neural population activity during action preparation, we found that state space analysis of simultaneously recorded units is significantly more predictive of subsequent reaction times (RTs) than traditional methods. Furthermore, because we observed a wide variety of individual unit characteristics, we developed the sign-corrected average rate (SCAR) method of neural population averaging. The SCAR method was able to explain at least as much variance in RT overall as state space methods. Overall, F5 activity predicted RT (18% variance explained) significantly better than AIP (6%). The SCAR methods provides a straightforward interpretation of population activity, although other state space methods could provide richer descriptions of population dynamics. Together, these results lend support to the differential role of the parietal and frontal cortices in preparation for grasping, suggesting that variability in preparatory activity in F5 has a more potent effect on trial-to-trial RT variability than AIP.

Key words: grasping; nonhuman primate; parietal; premotor; single unit recording

---

Significance Statement

Grasping movements are planned before they are executed, but how is the preparatory activity in a population of neurons related to the subsequent reaction time (RT)? A population analysis of the activity of many neurons recorded in parallel in macaque premotor (F5) and parietal (AIP) cortices during a delayed grasping task revealed that preparatory activity in F5 could explain a threefold larger fraction of variability in trial-to-trial RT than AIP. These striking differences lend additional support to a differential role of the parietal and premotor cortices in grasp movement preparation, suggesting that F5 has a more direct influence on trial-to-trial variability and movement timing, whereas AIP might be more closely linked to overall movement intentions.

Introduction

In the sport of fencing, rapid actions are required on the millisecond scale. Small rotations of the wrist can make the difference between a hit and a complete miss. The response of athletes to various attacks is highly variable, despite the rigorously trained nature of their skill set. What are the factors that contribute to the variability of such complex actions? It is known that voluntary movements are prepared before they are executed (Kutas and

Donchin, 1974; Wise, 1985; Ghez et al., 1997). A benefit of longer preparation is a reduction in reaction times (RTs), which is the time between a go signal and the initiation of a movement (Rosenbaum, 1980; Riehle and Requin, 1989). Nevertheless, RT varies even for similar amounts of preparation.

The ideal design for studying motor preparation is the delayed reaching task, in which a movement must be planned and withheld for a certain time. Studies have shown that population activity of neurons in the dorsal premotor cortex (PMd) of the primate brain, recorded either sequentially (Riehle and Requin, 1993) or in parallel (Churchland et al., 2006c; Afshar et al., 2011; Churchland, 2015), can explain a large portion of the variability in reach RT and reach velocity (Churchland et al., 2006a,b). Similar results have been obtained using sequential recordings in the parietal reach region (Snyder et al., 2006) and lateral intraparietal area (Janssen and Shadlen, 2005). However, a comparative study of the frontoparietal network has not been undertaken.

To analyze RT variability, an understanding of preparatory activity is vital. A number of models have been proposed to explain how preparation of movement is undertaken. Initial models related the preparatory activity of single neurons to behavior by suggesting that subpopulations of neurons may hold activity near a movement threshold that is crossed to initiate movement (Riehle and Requin, 1993; Hanes and Schall, 1996; Erlhagen and Schöner, 2002), whereas more recent models use a state space framework of population activity. In the latter framework, the firing of each neuron represents a dimension in a high-dimensional space of all neurons. Hence, the firing of all neurons at a particular time represents a single point in the state space, de-emphasizing the importance of tuning properties of individual neurons (Fetz, 1992). The "optimal subspace" hypothesis posits that a preparatory state is achieved during planning and that deviations from this state may delay movement (Churchland et al., 2006c; Churchland and Shenoy, 2007a). The "initial condition hypothesis" augmented this view by further stating that trial-to-trial fluctuations in the neural trajectory are correlated with RT (Afshar et al., 2011).

To elucidate the role of the frontoparietal network in preparation, the established hand grasping circuit (Luppino et al., 1999) consisting of the hand area (F5) of the ventral premotor cortex (PMv) and the anterior intraparietal area (AIP) were investigated using a delayed grasping task. Neural activity in these areas is modulated strongly by visual object properties (Murata et al., 1997, 2000), extrinsic goals (Kakei et al., 2001), performed grip types (Baumann et al., 2009; Fluet et al., 2010), and preparatory activity in these areas can be used to decode the visual properties of objects and complex hand shapes required to grasp a diverse range of objects (Carpaneto et al., 2011; Townsend et al., 2011; Schaffelhofer et al., 2015).

In the current study, we analyzed population activity in a delayed grasping task with multiple grip types to evaluate how population activity of simultaneously recorded units in F5 and AIP might inform subsequent behavior. Preparatory activity in F5 could explain up to 18% of the variability in trial-to-trial RT, a significant finding, whereas AIP could explain only up to 6%. By demonstrating a significant advantage of F5 over AIP in RT prediction, our results support the concept that the encoding of RT is represented primarily in the frontal and not the parietal lobe, at least when grasping in the dark.

Materials and Methods

Basic procedures. Neural activity was recorded simultaneously from area F5 and area AIP in one male and two female rhesus macaque monkeys (*Macaca mulatta*, animals B, S, and Z; body weight, 11.2, 9.7, and 7.0 kg, respectively). Animal care and all experimental procedures were conducted in accordance with German and European law and were in agreement with the *Guidelines for the Care and Use of Mammals in Neuroscience and Behavioral Research* (National Research Council, 2003).

Basic experimental methods have been described previously (Townsend et al., 2011; Schaffelhofer et al., 2015). We trained animals to perform a delayed grasping task. They were seated in a primate chair and trained to grasp a handle with the left hand (animals B and Z) or the right hand (animal S; Fig. 1D). This handle was placed in front of the monkey at chest level and in the vertical position at a distance of ~26 cm, i.e., the monkeys had to reach a distance of 26 cm to grasp the handle. The handle could be grasped either with a power grip (opposition of fingers and palm) or precision grip (opposition of index finger and thumb; Fig. 1E). Two clearly visible recessions on either side of the handle contained touch sensors that detected thumb and forefinger contact during precision grips, whereas power grips were detected using an infrared light barrier inside the handle aperture. The monkey was instructed which grip type to make by means of two colored LED-like light dots projected

from a thin-film transistor (TFT) screen (CTF846-A; screen size, 8 inches, digital; resolution, 800 × 600; refresh rate, 75 Hz) onto the center of the handle via a half mirror positioned between the animal's eyes and the target. A mask preventing a direct view of the image was placed in front of the TFT screen and two spotlights placed on either side could illuminate the handle. Apart from these light sources, the experimental room was completely dark. In addition, one or two capacitive touch sensors (model EC3016NPAPL; Carlo Gavazzi) were placed at the level of the animals' midtorso and functioned as hand-rest buttons. The nonacting arm of animals B and Z were placed in a long tube, preventing it from interacting with the handle. Monkey S was trained to keep its nonacting hand on an additional hand-rest button.

Eye movements were measured using an infrared optical eye tracker (model AA-ETL-200; ISCAN) via a heat mirror directly in front of the monkey's head. To adjust the gain and offset, red calibration dots were shown at different locations at the beginning of each session for 25 trials that the animal fixated for at least 2 s.

Eye tracking and the behavioral task were controlled by custom-written software implemented in LabView Realtime (National Instruments) with a time resolution of 1 ms. An infrared camera was used to monitor behavior continuously throughout the entire experiment.

Behavioral paradigm. Animals B and S performed Task 1 (Fig. 1E), whereas animal Z performed Task 2 (Fig. 1F). The following is an explanation of the trial course of Task 1. Trials started after the monkey placed the acting hand on the resting position and fixated a red dot (fixation period). The animal was required to keep the acting hand, or both hands (animal S), completely still on the resting position until after the go cue. After 400–700 ms, two flashlights illuminated the handle for 300 ms, followed by 600 ms of additional fixation. In the cue period, a second light dot was then shown next to the red one to instruct the monkey about the grip type for this trial (grip cue). Either a green or white dot appeared for 300 ms, indicating a power or a precision grip, respectively. After that, the monkey had to memorize the instruction for a variable memory period. This memory period lasted for 0–1300 ms (i.e., the go cue could appear simultaneously with the grip cue), in discrete memory period bins of 0, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000, or 1300 ms, which were pseudorandomly sampled with an equal number of trials from each condition. Regardless of memory period length, the grip cue was always shown for 300 ms. Switching off the fixation light then cued the monkey to reach and grasp the target (movement period) to receive a liquid reward. Animals were required to hold the appropriate grip for 300 ms. Additionally, catch trials were interleaved randomly (~8% of trials), in which a go cue was never shown and the animal only received a reward if it maintained fixation and the hands on the hand rests for 2000 ms after the grip cue. All trials were interleaved randomly and in total darkness.

The differences between Task 1 and Task 2 are as follows. In Task 2, there was only one fixation period that lasted for 600–1000 ms. In Task 2, the illumination of the handle took place at the time of grip cue. In the instructed version of Task 2, the grip cues were identical to Task 1. In the free-choice version, both a green and white dot appeared simultaneously, indicating that the monkey was free to choose between the two grip types. This was followed by a memory period lasting 400–600 ms, and then either the green or white dot reappeared for 300 ms in 50% of all free-choice trials, which turned the free-choice task into a delayed-instructed task and was followed by a second memory period (duration, 400–600 ms). In all other trials (instructed or free choice), only the red fixation dot was shown during the second cue period, making it impossible to distinguish the first and second memory periods. The hold period in Task 2 was 200 ms as opposed to 300 ms in Task 1. Importantly, during free-choice trials, the reward was reduced every time the monkey repeatedly chose the same grip type.

Surgical procedures and imaging. After completion of behavioral training, each animal received an MRI scan to locate anatomical landmarks for subsequent chronic implantation of microelectrode arrays. Each monkey was sedated (e.g., 10 mg/kg ketamine and 0.5 mg/kg xylazine, i.m.) and placed in the scanner (GE Healthcare 1.5T or Siemens Trio 3T) in a prone position. T1-weighted volumetric images of the brain and skull were obtained as described previously (Baumann et al., 2009). We measured the stereotaxic location of the arcuate and intraparietal sulci to guide placement of the electrode arrays.
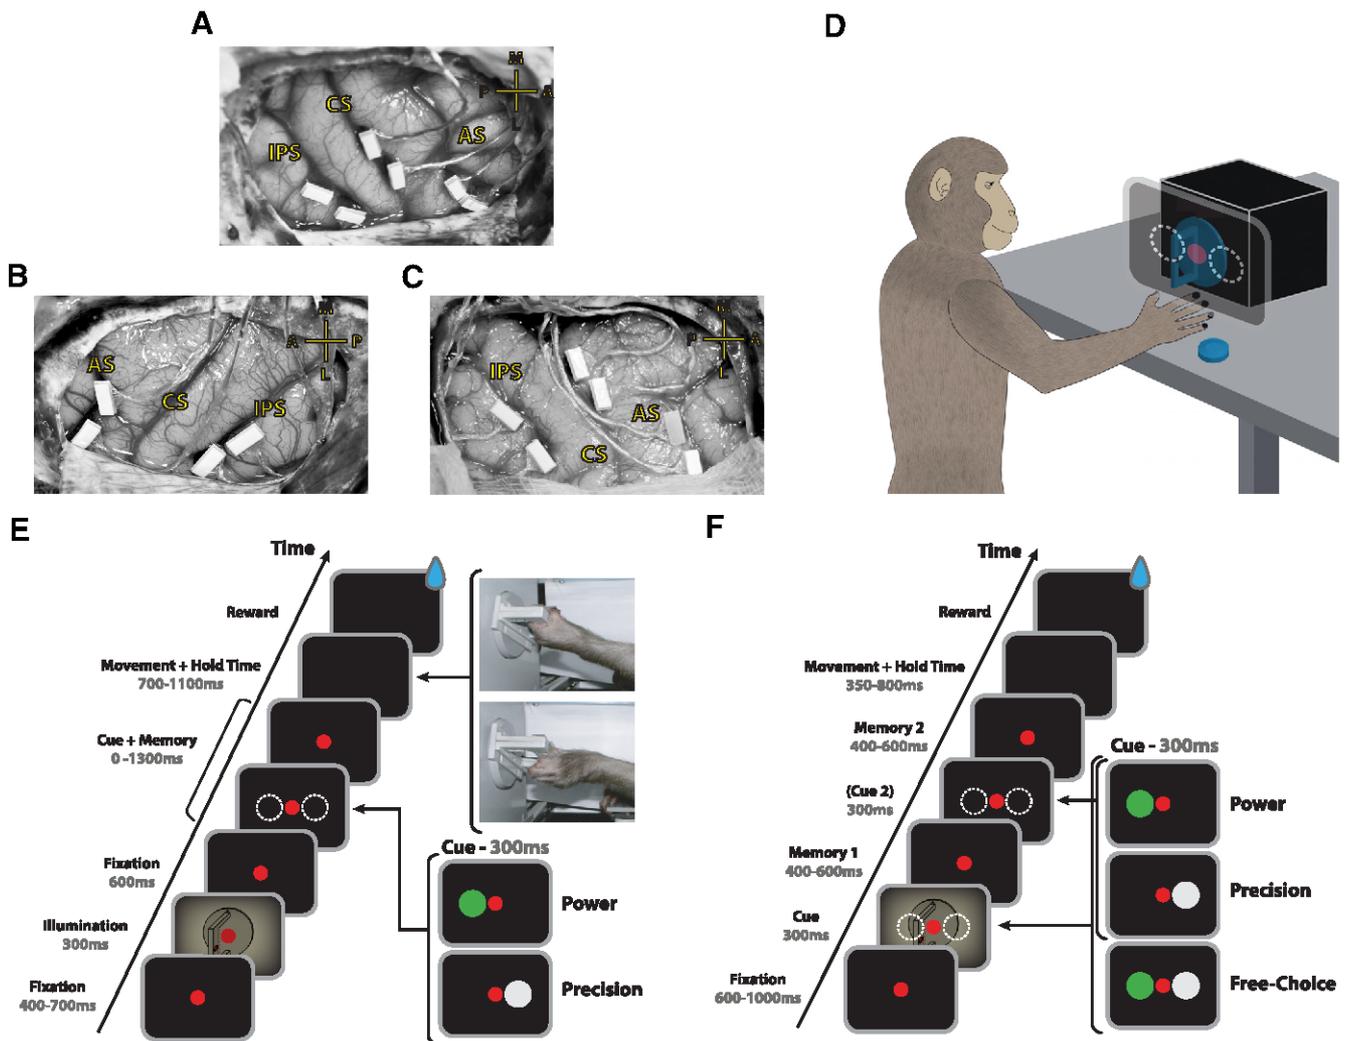
Figure 1. FMA implantation and task design. ***A–C***, Array locations for animals B, S, and Z, respectively. Two arrays were placed in F5 on the bank of the arcuate sulcus (AS). Two additional arrays were placed in AIP toward the lateral end of the intraparietal sulcus (IPS). In animals B and Z, two more arrays were placed on the bank of the central sulcus (CS). The cross shows medial (M), lateral (L), anterior (A), and posterior (P) directions. Note that animal S was implanted in the left hemisphere and animals B and Z in the right hemisphere. ***D***, Sketch of an animal in the experimental setup. The cues were presented on a monitor projected onto a mirror, making the light dots appear superimposed onto the grasping handle. ***E***, Delayed grasping task with two grip types (Task 1). An example of each grip type can be seen during the movement epoch (top, power grip; bottom, precision grip). The handle was rotated to a supine orientation for demonstration purposes only. ***F***, Delayed grasping task with two grip types and three decision conditions (Task 2). Free-choice trials were presented twice as often as each of the other conditions. Delayed-instructed trials contained a second grip cue turning a free-choice trial into a delayed-instructed trial. Trials were presented in a pseudorandom order.

Chronic electrode implantation. An initial surgery was performed to implant a head post (titanium cylinder; diameter, 18 mm). After recovery from this procedure and subsequent training of the task in the head-fixed condition, each animal was implanted with floating microelectrode arrays (FMAs; MicroProbes for Life Science) in a separate procedure. Animal S was implanted with 32 electrode FMAs and received two arrays in each area (Fig. 1*B*). The arcuate sulcus of animal S did not present a spur, but in the MRI a small indentation was visible in the posterior bank, ∼2 mm medial to the knee, which we treated as the spur. We placed both anterior FMAs lateral to that mark. Animals B and Z were implanted with six electrode arrays in the right hemisphere, each with 32 electrodes (Fig. 1*A*, *C*). Two such arrays were implanted in area F5, two in area AIP, and two in the primary motor cortex (M1). FMAs consisted of nonmoveable monopolar platinum–iridium electrodes with initial impedances ranging from 300 to 600 kΩ at 1 kHz measured before implantation. Post-implantation measurements in the first months after implantation confirmed these values *in vivo*. Lengths of electrodes were 1.5–7.1 mm.

All surgical procedures were performed under sterile conditions and general anesthesia (e.g., induction with 10 mg/kg ketamine, i.m., and 0.05 mg/kg atropine, s.c., followed by intubation, 1–2% isoflurane,

and analgesia with 0.01 mg/kg buprenorphene, s.c.). Heart and respiration rates, electrocardiogram, oxygen saturation, and body temperature were monitored continuously, and systemic antibiotics and analgesics were administered for several days after each surgery. To prevent brain swelling while the dura was open, the animal was hyperventilated mildly (end-tidal CO_2, ∼30 mmHg), and mannitol was kept at hand. Animals were allowed to recover fully (∼2 weeks) before behavioral training or recording experiments commenced.

Neural recordings and spike sorting. Signals from the implanted arrays were amplified and stored digitally using a 128 channel recording system (sampling rate, 30 kS/s; 0.6–7500 Hz hardware filter; Cerebus; Blackrock Microsystems). Data were first filtered using a median filter (window length, 3 ms), and the result was subtracted from the raw signal. Afterward, the signal was low-pass filtered with a causal Butterworth filter (5000 Hz; fourth order). To eliminate movement noise (i.e., common component induced by reference and ground), principal component analysis (PCA) artifact cancellation was applied for all electrodes of each array (as described by Musial et al., 2002). To ensure that no individual channels were eliminated, PCA dimensions with any coefficient >0.36 (with respect to normalized data) were retained. Spike waveforms were

extracted and semiautomatically sorted using a modified version of the offline spike sorter *Wave_clus* (Quiroga et al., 2004; Kraskov et al., 2009).

Unit isolation was evaluated using four criteria: (1) the absence of short (1–2 ms) intervals in the interspike interval histogram for single units, (2) the degree of homogeneity of the detected spike waveforms, (3) the separation of waveform clusters in the projection of the first 17 features detected by *Wave_clus*, and (4) the uniqueness of the shape of the interspike interval distribution.

After the semiautomatic sorting process, redetection of the average waveforms (templates) was done to detect overlaid waveforms (Gozani and Miller, 1994). Filtered signals were convolved with the templates starting with the biggest waveform. Independently for each template, redetection and resorting was run automatically using a linear classifier function (MATLAB function *classify*). After the identification of the target template, the shift-corrected template (achieved by up and down sampling) was subtracted from the filtered signal of the corresponding channel to reduce artifacts for detection of the next template. This procedure allowed a detection of templates up to an overlap of 0.2 ms. As a control, unit isolation was evaluated again as described previously to determine the final classification of all units into single units or multiunits. In case of ambiguity, a unit was not classified as single. Stationarity of firing rate was checked for all units, and, in case the firing rate was not stable over the entire recording period (>30% change in firing rate between the first 10 min and the last 10 min of recording), the unit was excluded from additional analyses (<3% of all single units).

Data preprocessing. In all datasets trials with outlying RTs, >700 ms in Task 1 and >500 ms in Task 2 and <200 ms in either task were excluded. In animals B and S, these trials comprised <1% of the data and <3% in animal Z. Clearly, all animals were careful to wait for the appropriate go cue and did not act preemptively. We used this conservative check on outlier RTs to safely exclude the possibility that animals were acting in anticipation of the go cue.

Crucially, for all analyses of Task 1, trials with memory periods <500 ms were excluded from analysis. These short memory period trials were removed to ensure that animals had sufficient time to fully plan the movement before acting. Such an exclusion criteria was not used in Task 2, because the animal never had <700 ms to plan (delayed-instructed condition) and was trained for many months to acquire this timing scheme.

All recorded units (single unit and multiunit) were used in our main analyses. After spike sorting, spike events were binned in overlapping 100 ms windows and sampled every millisecond to produce a continuous firing rate signal (1 kHz). This means that firing rates at the time of the go cue considered spikes occurring 50 ms before to 50 ms after the go cue. Because it is unlikely that (sensory) responses to the go cue would be represented in AIP or F5 already at 50 ms after presentation, we believe this binning does not bias the predictive power of RTs. In fact, our conclusions do not change when using a binning that does not extend beyond the go cue (data not shown).

Dimensionality reduction. Dimensionality reduction was performed for the purposes of visualization only. All quantitative analyses relied on the full dimensionality of the data. Gaussian-process factor analysis (GPFA) was performed on the neural data from cue presentation to movement onset (Yu et al., 2009). This method performs smoothing of spike trains and dimensionality reduction simultaneously within a common probabilistic framework. It assumes that the activity of each unit is a linear function (plus noise) of a low-dimensional neural state whose evolution in time is well described by a Gaussian process. This methods allows for better visualization on the single-trial level than other published methods (Yu et al., 2009). The data were reduced to 12 dimensions (the optimal number of latent dimensions in the data as determined by cross-validation) using 20 ms nonoverlapping spike bins to produce the trajectories in Figure 3A. In this reduction, the three displayed dimensions explain 63% of the total variance. In this figure, a rotation of the first three latent dimensions is shown (equivalent to a linear combination of the three dimensions explaining the most variance overall).

Similarly, neural trajectories in Figure 9 were generated by performing PCA on the peristimulus time histograms of all units for each grasp condition separately. All individual trials were then transformed into the two principal components explaining the most variance and binned into

slow, medium, and fast RTs. All trials were aligned to the go cue and plotted from 350 ms before to 280 ms after the go cue.

Projection methods. As can be seen in Figure 3A, trials of the same condition tend to follow a stereotypical trajectory through neural space. Following the study by Afshar et al. (2011), we reasoned that the farther the neural state had advanced along the mean neural path at the time of the go cue would be predictive of subsequent performance. To test this hypothesis, we projected neural activity of individual trials at the go cue on the mean neural trajectory of similar trials (excluding the tested trial of the same condition). The projection is denoted in Figure 3B with the symbol α. The vector formed between the mean firing rate at the go cue and the firing rate at the go cue of an individual trial is projected onto the vector between the mean firing at go and the mean firing at go $\pm$ some Δt. The data were tested empirically to determine the optimal Δt values over all datasets. Selected Δt values ranged from 300 ms before to 300 ms after the go cue.

Additionally, as depicted in Figure 3D, the instantaneous velocity of the neural data, $[t_{go} - (t_{go} - 20)]$, in the high-dimensional neural space of individual trials was projected onto the mean neural trajectory. Similar to the projection method, the velocity projection method hypothesized that trials in which the neural space is changing in the direction of the mean trajectory will have shorter RTs. Importantly, trials were segregated into 100–200 ms bins based on the length of the memory period to minimize the effect of memory period length on neural position, i.e., the mean trajectory used as a reference for each trial was calculated solely on other trials within the same memory period bin.

Euclidian distance method. The Euclidian distance method was performed also equivalently to the study by Afshar et al. (2011). Single-trial RT was correlated with the Euclidian distance between the high-dimensional firing rate at the go cue on the single trial and the mean high-dimensional firing rate of all other trials of the same condition at some time offset, Δt, as depicted in Figure 3C.

The optimal subspace method, as originally reported by Churchland et al. (2006c), was also performed. It is equivalent to the Euclidian distance method with a time offset of $\Delta t = 0$ ms. Both of these methods are based on the hypothesis that trials in which firing rates are close to the mean rates observed for similar trials have shorter RTs.

Average rate method. The average rate (AR) method is based on the simple hypothesis that trials during which particular units have higher firing rates will be associated with shorter RTs. This method posits that neural activity increases during preparation and crosses a movement threshold to initiate a movement, also known as the rise-to-threshold hypothesis (Erlhagen and Schöner, 2002). Under the assumptions of this method, higher preparatory activity would always be associated with shorter RTs. Four implementations of this method were tested initially. The trial-by-trial RT was correlated with the following: (1) the signed difference between firing rate at go cue and at cue onset (i.e., an approximation of baseline firing), averaged across all units; (2) the same method but using the unsigned difference (absolute value); (3) the average firing rate at the go cue across all units; and (4) the average firing rate at the go cue across all units for their preferred grip type only. The third version, which does not rely on baseline firing rate or unit preferences, was the best performing (data not shown) and was therefore the one used for additional analysis. For clarity, we opted to name our implementation of the rise-to-threshold hypothesis as the AR method.

Sign-corrected average rate method. As hypothesized by the AR method, if units that increase their activity (relative to the mean) during movement preparation are associated with trials having short RTs, then they are negatively correlated with RT. However, if the activity of some units were in fact reduced (relative to the mean) for trials with short RT, this would result in a positive correlation. If many of each of these types exists in the same population, which is averaged to produce an RT prediction, these two inverted populations would be in conflict and cancel out each other, thereby causing poor RT prediction.

To overcome this obstacle, we introduced the sign-corrected average rate (SCAR) method. It is identical to the simple AR method as described in the previous section; however, the signal of all units was first multiplied with a sign-correction vector. That is, units that were correlated positively with RT were inverted to produce a negative correlation. To
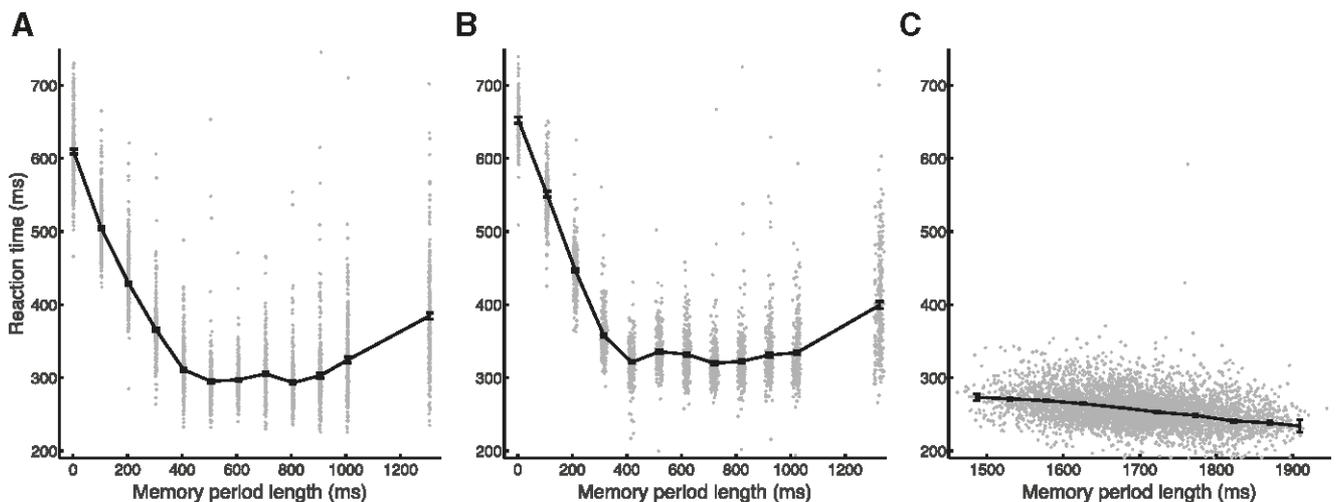
Figure 2. Scatter plots of RT versus memory period length. *A–C*, The RT of animals B, S, and Z, respectively, as a function of memory period length for all task conditions and datasets. The solid line represents the mean, and error bars indicate SEM within nonoverlapping 50 ms bins.

decide which units were to receive a negative value in the sign-correction vector, (1) the firing rates at the go cue of individual units were correlated with RT (twofold cross-validated) over all conditions. (2) Units received a -1 value in the sign-correction vector if they possessed on average (over all conditions) an r value >0. All other units received a value of 1 in the correction vector. This method preserves the absolute magnitude of the mean firing rate across trials because no normalization is performed. A number of inversion criteria were tested; however, we found that it was sufficient to invert only units with a positive r value (data not shown). It is important to note that units were not tested separately for each condition (grip type/decision condition), i.e., a unit that was inverted for a precision grip would also be inverted during a power grip. Testing on each condition separately would have increased RT prediction further.

RT correlations. When correlating single-trial neural metrics, i.e., the previously described methods excluding the AR and SCAR methods, with RT, we did not include the neural data from that trial in the calculation of the mean neural trajectory used for that prediction, as in the study by Afshar et al. (2011). The predicted and observed RTs were then correlated with each other. This technique, termed leave-one-out cross-validation, ensured that predictions of the RT of each trial were not based on movement activity from that trial.

Whenever average RT variance explained was calculated across an average of datasets, each average was weighted by the number of trials in each dataset.

Partial RT correlations. In our tasks, memory period length was correlated highly with RT (Fig. 2). To disentangle the relationship of memory period length to RT and the relationship of our neural prediction metrics on RT, partial correlation, which bares much similarity to multiple regression, was performed (Cramér, 1946). Partial correlation is a method for determining the correlation between two variables while controlling for one or more other variables. The partial correlation between two variables, while controlling for a single other variable, is described by

$$\rho_{NB \times M} = \frac{r_{NB} - r_{NM}r_{MB}}{\sqrt{1 - r_{NM}^2}\sqrt{1 - r_{MB}^2}},$$

where ρ is the partial correlation of a neural prediction metric (N) with RT (B), while controlling for the effect of the length of the memory period (M). r_{XY} is the standard Pearson's correlation between vectors X and Y.

Cross-validation. The results of all methods were twofold cross-validated. All trials of each dataset were first randomly segregated into two sets of equal size and methods performed separately on each set. Furthermore, the SCAR method required the preevaluation of preparatory correlations with RT to determine which units should have their

firing rates inverted. To avoid false-positive results, SCAR was first trained on a training set of trials and always tested on trials that were not used for training. All analyses were twofold cross-validated by flipping the role of both sets. Segregating the data into more than two cross-validation folds would severely reduce the number of test trials in each condition and therefore the reliability of prediction.

Multiple linear regression. To determine whether a combination of the tested methods could improve the amount of variance explained in RT, a number of regressions was performed. Multiple regression was performed using the leave-one-out technique, in which regressing on all other trials generated the prediction for each trial, and this process was repeated for each individual trial. First, the same model as described by Afshar et al. (2011) was used, which consisted of the projection method on both the pre-go and post-go cue axes, as well as the velocity projection method on both the pre-go and post-go cue axes. Alternatively, a number of simpler combinations were tested, although most are not presented here because they yielded poor results.

To test whether or not a multivariate model could explain significantly more variance than a simpler model, the F test was used. The F test is ideally suited to compare models (regressions) that use nested predictors, that is, models that use a subset of predictors of a more complex model. However, because we wanted to compare models over a number of conditions (each with varying degrees of freedom), we had to generate a nonstandard F distribution for testing. Therefore, the F statistic comparing each pair of models was calculated separately and then summed. Additionally, because each cross-validation fold contained different trials, each fold was considered as a separate condition for a total of four conditions in Task 1 (two behavioral, two folds) and 12 conditions (six behavioral, two folds) in Task 2. To generate a testing distribution the probability density functions (pdfs) of each corresponding F statistic were convolved with each other to form a new distribution. We then calculated the likelihood of observing the calculated sum of F statistics and from there derived the p value.

To extend this test over all datasets and reach general conclusions, the sum of F statistics was summed across all conditions and datasets and tested on an F distribution of convolved pdfs over all conditions and datasets.

Chance-level calculation. Many individual correlations were computed in the current study. To ensure that all relevant methods were truly identifying relationships between neural data and RT, all correlations were tested against a chance distribution. For each method and condition, chance distributions were generated by correlating the prediction of each method with a corresponding vector of randomly shuffled RTs (1000 repetitions). We could then calculate the probability of observing the empirical R^2 given our shuffled distributions and use this as a p value.

In the case in which a significance calculation was required over multiple conditions and datasets, the generated chance distributions were convolved with each other to form a new distribution, precisely as with the F statistic for testing multiple linear regression. The p value for significance was fixed at 0.01.

For the SCAR method, an additional control was performed. The SCAR method involved the inversion of the firing rates of some units. To ensure that this inversion did not artificially produce our results, the following control was performed. A random sample of units of the same size as in the real data was inverted and the method performed as normal (1000 repetitions, permutation test). The resulting chance-level distributions could be tested against the empirical results as was done for the other chance-level calculations.

Variance selection. All recorded units were included in the main analyses. To determine whether one could select a subset of units that would perform equally or better than the entire population, a variance selection was performed. The units with higher variances in spike count (at the go cue) across trials were preferentially included first. In addition, a random unit selection was performed alongside the first analysis with the same number of units per test. The random selection of units was performed 1000 times per percentage value. Data were interpolated to the range of 0–100% to facilitate averaging between datasets.

Significance testing was performed by summing the R^2 over all datasets and testing the likelihood of obtaining this value against the distribution of convolved pdfs over all datasets as generated by the random unit selection, precisely as was done with the F statistic for testing multiple linear regression. The significance level was set at 0.05 and Bonferroni's corrected for the number of percentages tested (100).

Bayesian information criterion. The Bayesian information criterion (BIC) is a well known model selection criterion (McQuarrie and Tsai, 1998). It is described by the following:

$$\text{BIC} = -2\ln\acute{L} + p\ln N,$$

where $\acute{L}$ is the posterior likelihood of the data given the best-fit model, p is the number of parameters used to generate the model, and N is the number of observations used. A smaller BIC is associated with a better explanatory model. BICs were calculated for single conditions and averaged either over conditions or over conditions and datasets.

Results

Behavior
All three animals performed the task successfully. After initiating trials to the point of obtaining task information, i.e., receiving a grip cue, animals B, S, and Z successfully completed those trials 96, 98, and 95% of the time, respectively. Catch trials in which the animal was required to withhold movement were included in Task 1. Animals B and S completed these catch trials successfully 95 and 98% of the time, respectively. Figure 2 plots the RTs of all animals as a function of memory period length. The memory period in Task 1 lasted 0–1300 ms, whereas the memory period in Task 2 was relatively longer (a minimum of 1400 ms in the instructed condition including the grip cue) to facilitate a second cue period in the delayed-instructed condition. RTs were reduced during longer memory periods, consistent with the established hypothesis that motor preparation improves over time (Rosenbaum, 1980; Riehle and Requin, 1989). The exception to this was the 1300 ms memory condition in Task 1, in which RT slightly increased, likely because of the expectation of a catch trial, which appeared periodically and lasted 2000 ms. For animals B, S, and Z, the correlation coefficients over all datasets between memory period length and RT were −0.55, −0.57, and −0.33, respectively. Similar experiments have shown that saturation of RT, i.e., the minimum length of memory period after which RT does not significantly improve, is at least 100–200 ms (Churchland et al., 2006c) in a reaching task. In Task 1, we observed RT saturation, but we did not observe this in Task 2.

In contrast, there was no significant correlation between memory period length and movement time, which is the time between the hand leaving the hand-rest button and making contact with the handle, indicating that animals only initiated a movement when the movement was fully prepared. The only exceptions are the movement times of animal B, which were slightly longer in the 1300 ms memory period condition than in shorter memory periods, potentially an effect of decreased attention for long memory periods. For all animals, the hands remained completely stationary on the hand-rest buttons before the go cue. The experiments from which these data were collected were originally designed to assess hypotheses that are not presented here and will be addressed elsewhere.

Neural recordings
The analyzed datasets include a collection of 18 recording sessions, six from each animal. In animal B, an average ± SD of 63 ± 17 and 28 ± 18 units were recorded in F5 and AIP, respectively, as well as in animal S (mean ± SD, 132 ± 15 and 131 ± 26) and animal Z (85 ± 18 and 81 ± 24). An average of 483 trials per dataset met the inclusion criteria, as described in Materials and Methods. This corresponded to an average of 77 trials per condition and cross-validation fold overall. In animals S and Z, there was no significant difference between the two brain regions in the number of units recorded per dataset ($p = 1$ and $p = 0.56$, Wilcoxon's signed-rank test). However, in animal B, significantly more units were obtained in area F5 ($p = 0.03$), which may have affected the quality of RT decoding in area AIP. The majority of units obtained in all animals were identified as multiunits (52% in animal B, 60% in animal S, 70% in animal Z). All recorded single units and multiunits were included in additional analyses.

Although the response characteristics of each individual unit are not analyzed here in detail, it is worth noting that the overall tuning characteristics of units in F5 and AIP were very similar regardless of the task design used (Task 1 or 2). Furthermore, both tasks were able to elicit strong grip type tuning in both F5 and AIP. An average of 32% of recorded units were significantly tuned for grip type during the late memory period in F5, whereas 26% were tuned in AIP ($p < 0.05$, two-sample t test), which did not differ between areas ($p = 0.09$, Kruskal–Wallis ANOVA), although differences were seen between animals ($p = 0.002$, Kruskal–Wallis ANOVA), with animal B showing slightly less tuning overall (24% in F5 and 18% in AIP). This finding is particularly robust when considering that there are no visual cues present in the memory period, and, therefore, grip type tuning tends to reach a minimum during this epoch.

Low-dimensional visualization of single-trial trajectories
To visualize how neural data evolves on single trials, a low-dimensional representation of the full neural space of both brain areas combined is shown in Figure 3A for an exemplar dataset (instructed precision grip, dataset Z120829). Dimensionality reduction was performed using GPFA, as described in Materials and Methods. Single trials tended to evolve from cue onset to a preparation state and further to a movement state after the go cue. Conversely, it did not appear that variability between trials decreased in a systematic way when comparing the size of the confidence ellipses at cue onset, go cue, and movement onset. To determine whether the trajectory of an individual trial could be related to RT, three methods were formulated, as depicted in Figure 3B–D. These three methods, the projection method, Euclidean distance method, and velocity projection method, are presented here virtually identical to how they were performed by
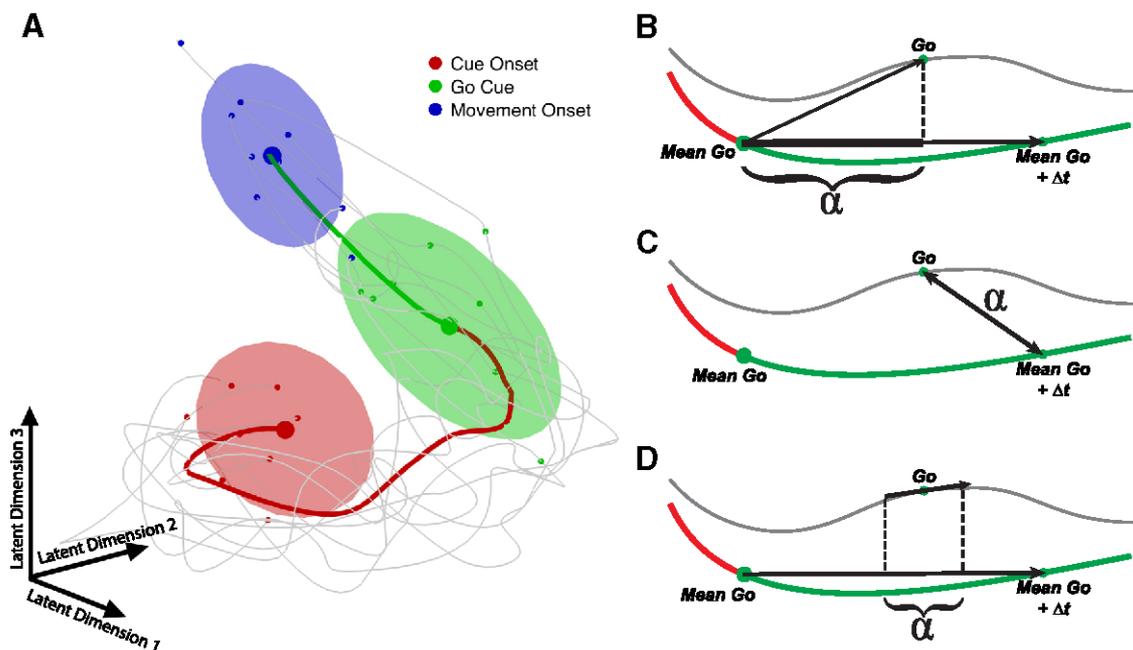
Figure 3. Neural data and RT prediction methods visualized as low-dimensional trajectories. ***A***, Neural data of both areas of an exemplar condition reduced to a low-dimensional representation of the trial course (determined by GPFA). Thick trace represents the mean of trials for one condition (instructed precision grip, dataset Z120829). Thin gray traces represent 10 random single trials. Shaded ellipses (90% confidence) represent the state of all selected single trials at the start of each epoch. ***B–D***, High-dimensional RT prediction methods in a two-dimensional illustration. Thick red and green traces represents the mean of trials. Thin gray trace represents a single exemplar trial. α denotes the component used to predict RT for the projection method (***B***), Euclidean distance method (***C***), and velocity projection method (***D***).

[Afshar et al. (2011)](#) and are described in detail in Materials and Methods. Although GPFA aids the visualization of single trials, in the following section, these RT prediction methods are tested in the high-dimensional state space of all recorded units.

Finding optimal reference points for trajectory-based methods

As can be seen in [Figure 3B–D](#), each of the three high-dimensional state space methods relies on a reference time point, or Δt, on the mean trajectory. To find the optimal reference point, the time domain from 500 ms before to 500 ms after the go cue was tested. For purpose of determining optimal Δt values, this time window was further limited to ± 300 ms, because reference points become more unreliable between conditions and datasets at large offsets. Each method was performed with this range of Δt values, and the predictions correlated with RT, as seen in [Figure 4](#). The mean of all conditions is shown with a thick trace, and the used offsets before and after the go cue are marked with open circles (limited within ± 300 ms). Inset histograms show all individual correlation coefficients [datasets (6) $\times$ conditions (2–6) $\times$ cross-validation folds (2)] before squaring and averaging and for each animal separately. The darker bars indicate correlations that are statistically significant ($p < 0.05$, Pearson's correlation). The results from animals B and Z are very similar and use identical offsets, whereas the results of animal S differed significantly.

As seen in [Figure 4A](#), the projection method using references both before and after the go cue have correlation distributions with nonzero median in F5: one distribution is shifted to the negative and one to the positive. This finding is consistent with our hypothesis, because trials that are farther along a mean trajectory going forward in time should lead to shorter RTs and therefore an overall negative correlation between our neural predictor and RT. In AIP the projection method also performed

significantly, although the resulting R^2 is much lower than in F5 ([Fig. 4B](#)). Based on this analysis, the best positive and negative Δt values, which were then used in all subsequent analysis, were -290 and 60 ms in F5 for animals B and Z and -170 and 260 ms for animal S. In AIP, values of -210 and 200 ms were used for animals B and Z, and values of -40 and 60 ms were used for animal S.

The Euclidian distance method performed similarly to the projection method but explained overall less variance in RT ([Fig. 4C,D](#)). In both F5 and AIP, reference points generally produced correlation histograms that were shifted significantly from zero. In most cases the pre-go distribution was shifted to the negative direction and the post-go to the positive direction, again consistent with the hypothesis that trials that are closer to the state of the network after the go cue will have shorter RTs, with the notable exception of animal B on the pre-go axis, a point that is returned to later. Additionally, when using a time offset of 0 ms, identical to the so-called optimal subspace method ([Churchland et al., 2006c](#)), the correlation distribution tended to be only marginally significantly shifted from zero in F5 and AIP. Despite this, the optimal subspace method was not used in additional analyses, because the Euclidean distance method outperformed it in every case. Based on this analysis, the Δt values that were used in additional analysis were -300 and 170 ms in F5 for animals B and Z and -270 and 270 ms for animal S. In AIP, values of -90 and 300 ms were used for animals B and Z and values of -100 and 300 ms for animal S.

The velocity projection method performed poorly overall, accounting for $<1\%$ of the variance in RT ([Fig. 4E,F](#)). Only rarely do reference points in F5 or AIP have correlation distributions significantly shifted from zero. Furthermore, accounting for the effect of memory period length on RT using partial correlation completely eliminates this effect (data not shown). Therefore, for most of our additional analyses, the velocity projection method
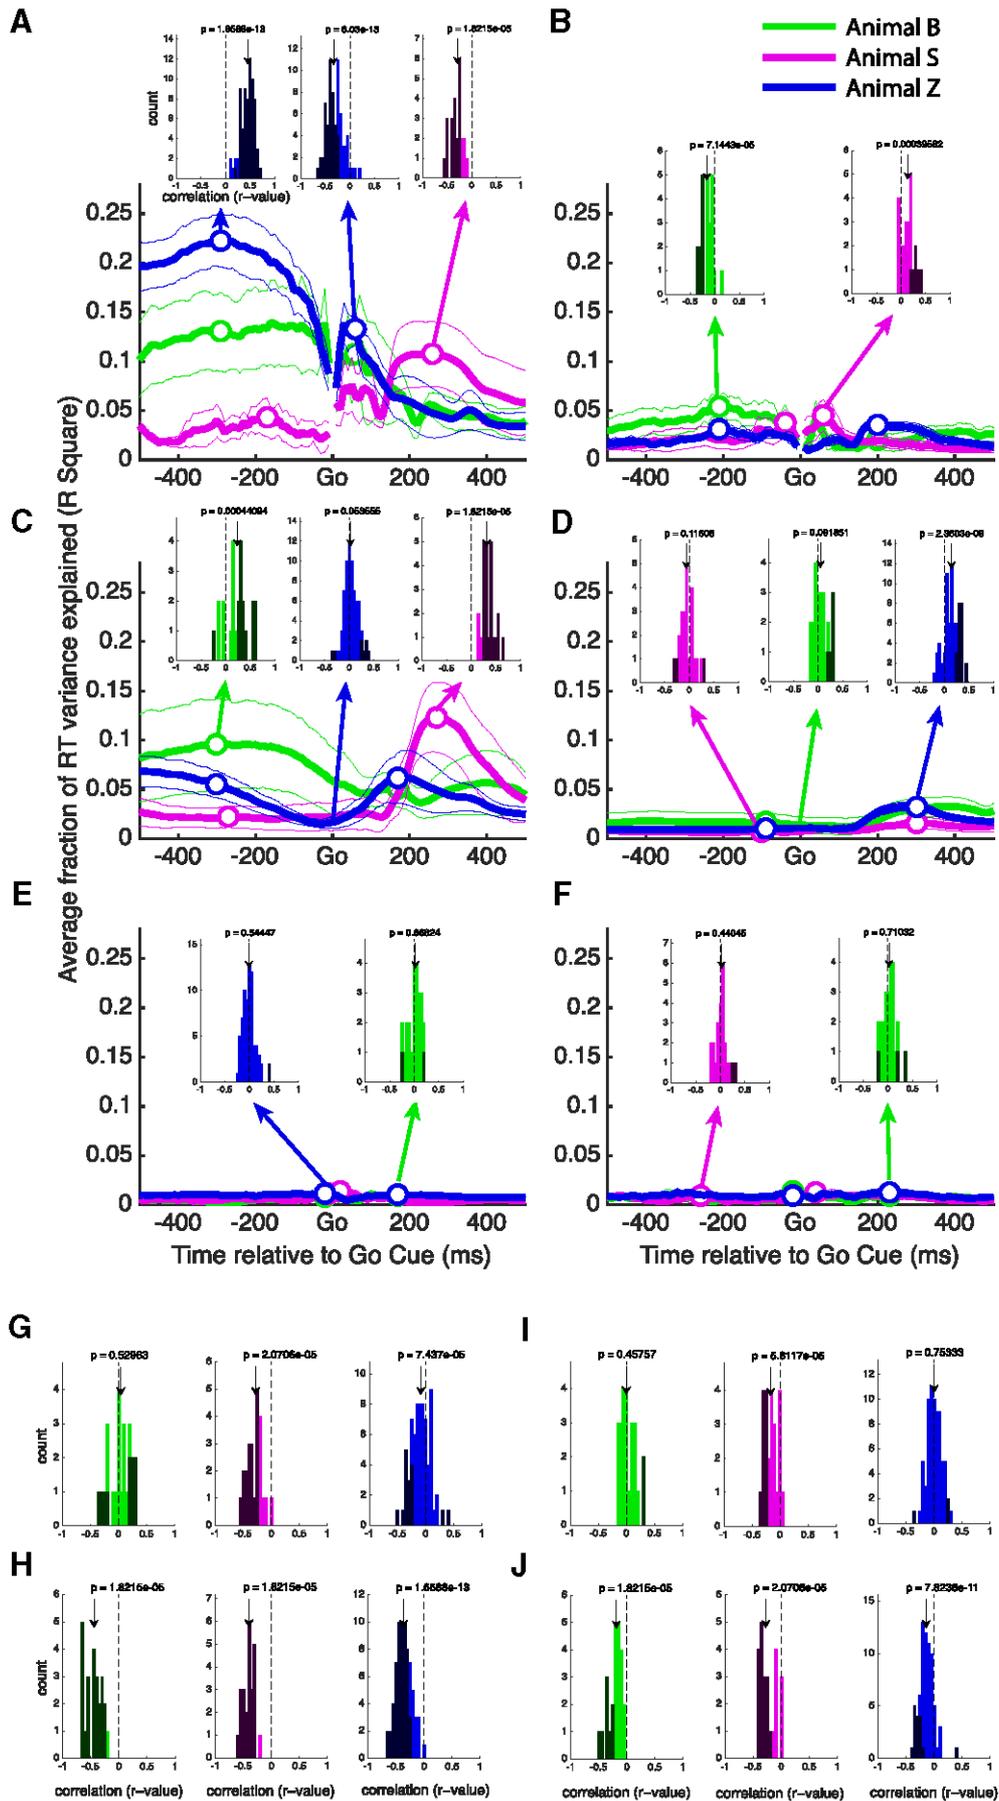
was excluded. It should be noted that trials were segregated into bins based on memory period length, as described in Materials and Methods. However, when all trials are pooled together, the resulting predictions of RT are still nonsignificant (data not shown).

Population firing rate-based methods

In addition to our high-dimensional trajectory-based methods, we also tested simpler methods based on averaging the activity of all units around the go cue. Such methods still depend on simultaneously recorded units, because they require an estimation of the population neural state for each trial. The first method we tested is the AR at go, which is our implementation of the rise-to-threshold hypothesis, as described by Afshar et al. (2011). The correlation histograms obtained by the AR method are shown in Figure 4, *G* for F5 and *I* for AIP. The median of the correlation distribution is significantly shifted negatively in two of the three animals in F5 (Wilcoxon's signed-rank test), suggesting that higher firing rates around the go cue led to shorter RTs. However, in AIP, the distribution was only shifted for one of the three animals. However, in all cases in which a significant shift was present, this shift was in the negative direction, suggesting that higher firing rate tended to be related to shorter RTs.

The AR method relies on averaging. Therefore, if some units in the population are correlated negatively with RT whereas others are correlated positively, these effects could cancel out at the population level. To deal with this issue, we first correlated the firing rate at the go cue of each unit with RT on a set of training trials. Then, as described in Materials and Methods, on a set of testing trials we inverted the firing rates of units that had a positive correlation in the pretesting (twofold cross-validated). Briefly, the process consists of multiplying the firing of all units by a sign-correction vector (see Materials and Methods). This new method was termed the SCAR method. The correlation histograms of the SCAR method are shown in Figure 4, *H* and *J*, for F5 and AIP, respectively. In both areas the median of the correlation distribution was shifted strongly into the negative domain (three of three animals in both areas, Wilcoxon's signed-rank test). Over all datasets, the average number of units whose activity was inverted was 38% in F5 and 42% in AIP, a large portion of the total unit count. The number of units inverted was less for animal S, in which the performance of the AR method was already considerably high.

Pooling of multiunits and single units does not bias RT prediction

To ensure that the previous results were not attributable to the sole contribution of either multiunits or single units, we repeated the analysis using only multiunits or single units. Results are presented as a performance ratio of average fraction of RT vari-

←

Figure 4. Determination of the optimal reference time (Δ*t*) relative to go cue on the mean trajectory. *A*, *B*, Results of the projection method in areas F5 and AIP, respectively. *C*, *D*, Results of the Euclidean distance method in areas F5 and AIP, respectively. *E*, *F*, Results of the velocity projection method in areas F5 and AIP, respectively. Thick traces are the mean of all conditions and datasets of each animal, thin traces are the SEM, and white circles are the optimal Δ*t* used in all subsequent analysis. Insets in *A–F* show histograms of correlation coefficients between each neural predictor and RT over all conditions (2–6), datasets (6), and cross-validation folds (2). Black bars denote correlations with a *p* value <0.05. Arrows show the median together with the *p* value of significant difference from zero (Wilcoxon's signed-rank test). *G*, *H*, Correlation coefficient histograms of the AR at go method and the SCAR method, respectively, in F5. *I*, *J*, Same as *G* and *H*, but for neural data from AIP.

ance explained using only single units or multiunits versus the pool of all units (where 1 represents identical performance). In F5, when including only single units, the pre-go and post-go projection methods produced ratios of 0.81 and 0.84, respectively. In AIP, the pre-go and post-go projection methods produced ratios of 0.76 and 0.91, respectively. The pre-go and post-go Euclidean distance methods produced ratios of 0.65 and 0.95 for F5 and 0.72 and 0.81 for AIP single units.

When including only multiunits, in F5, the pre-go and post-go projection methods produced ratios of 0.64 and 0.69 in F5 and 0.77 and 0.69 in AIP. The Euclidean distance methods produced ratios of 0.87 and 0.87 in F5 and 1.13 and 0.59 in AIP.

The same analysis was done for the population-based methods (AR and SCAR). The AR method had performance ratios of 0.74 and 0.79 when using single units only in F5 and AIP, respectively, whereas multiunit only performance was 1.25 and 1.07. The SCAR method had performance ratios of 0.76 and 0.87 when using single units only in F5 and AIP, respectively, whereas multiunit only performance was 0.66 and 0.63.

In almost every case, including only multiunits or single units in the analysis reduced the overall performance. Using only single units caused a performance reduction of 9–36%. Similarly, using multiunits caused a reduction of 13–41%, with the exception of the Euclidean distance method (before go cue) in AIP and the AR method in both areas, which increased slightly. Overall, the prediction of RT cannot be explained based solely on the contribution of either single units or multiunits. However, the AR method seems to perform best using multiunits, suggesting that compounding single-unit responses stabilizes the performance of this method.

Neural activity predicts trial-by-trial RT

Together, the results of all methods based on optimal Δ*t* selection are shown in Figure 5. The mean R^2 is plotted for all 18 datasets in both F5 and AIP. Open bars mark methods that did not perform above chance level, as described in Materials and Methods. The average R^2 over all methods is significantly higher in F5 than AIP ($p < 0.001$, Kruskal–Wallis ANOVA), suggesting that the population activity in F5 better encodes the variability in grasping plans. However, it is important to note that the current tasks include a large reaching component, which is also represented strongly in F5 and AIP (Lehmann and Scherberger, 2013), although more so in AIP. Therefore, some similarities between the behavioral and neural results of the current study and previous reaching studies are expected.

Not all methods achieved the same level of performance overall ($p < 0.001$, Kruskal–Wallis ANOVA). The best method on average, SCAR, was able to explain 18% of the variance in RT in F5 and 6% in AIP. The SCAR method and the projection (before go cue) method performed best in F5 for animals B and Z, explaining 18 and 16% of the variance in RT, respectively. In animal S, this pattern differed in F5, because the best performing methods were SCAR and Euclidean distance (after go cue), explaining 17 and 13% of variance in RT, respectively. The mean RT prediction for each animal is summarized in Figure 6A. There was no effect of grip type ($p = 0.69$, Kruskal–Wallis ANOVA) in all animals, suggesting that RT could be predicted equally well regardless of grip. Additionally, there was no effect of cross-validation fold ($p = 0.93$, Kruskal–Wallis ANOVA), confirming that segregating the data into training and testing trials did not introduce inconsistencies into the results.

As described previously, Task 2 contained different task types (instructed, free choice, and delayed instructed). There was a

Figure 5. Average fraction of RT variance explained for all methods and datasets (averaged across conditions and cross-validation folds). **A**, Average fraction explained by F5 data. **B**, Average fraction for AIP. Note the clear advantage of area F5 over AIP. Chance-level calculation is based on shuffling neural data with respect to RTs repeatedly. The observed R^2 values are then compared against the shuffled distributions to assess significance. Significant results are illustrated as solid bars, whereas the open bars show results that can be explained by chance ($p \geq 0.01$).

significant effect of task type on RT prediction ($p < 0.001$, Kruskal–Wallis ANOVA) over all methods, although the effect size was very small (effect size, $\eta^2 = 0.018$). The worst performing decision condition was the delayed-instructed condition, in which a second cue was presented later in the memory period. This small, but significant, effect on RT prediction is likely attributable to the disruptive effect of a second cue close to the end of the memory period. Interestingly, there was no difference in RT prediction between the instructed condition and the free-choice condition ($p = 0.80$, Wilcoxon's rank-sum test), suggesting that the way in which a motor plan is selected does not affect the relationship between preparatory activity and RT.

To summarize the number of individual correlations that have significant p values ($p < 0.05$; equivalent to the black bars in the histograms of Fig. 4), the total fraction of significant correlations is plotted in Figure 6C. In F5, between 21 and 96% of the correlations were significant for each method, whereas this range was between 4 and 67% in AIP, therefore confirming the overall better predictability of RT in F5.

Given the success of the SCAR method, an interesting question arises. If it is effective to predict RT by calculating a weighted mean of all units, in which the weights are either exactly −1 or 1, would performance improve if weights were not restricted in any way? This idea can be tested directly by using linear regression to

41
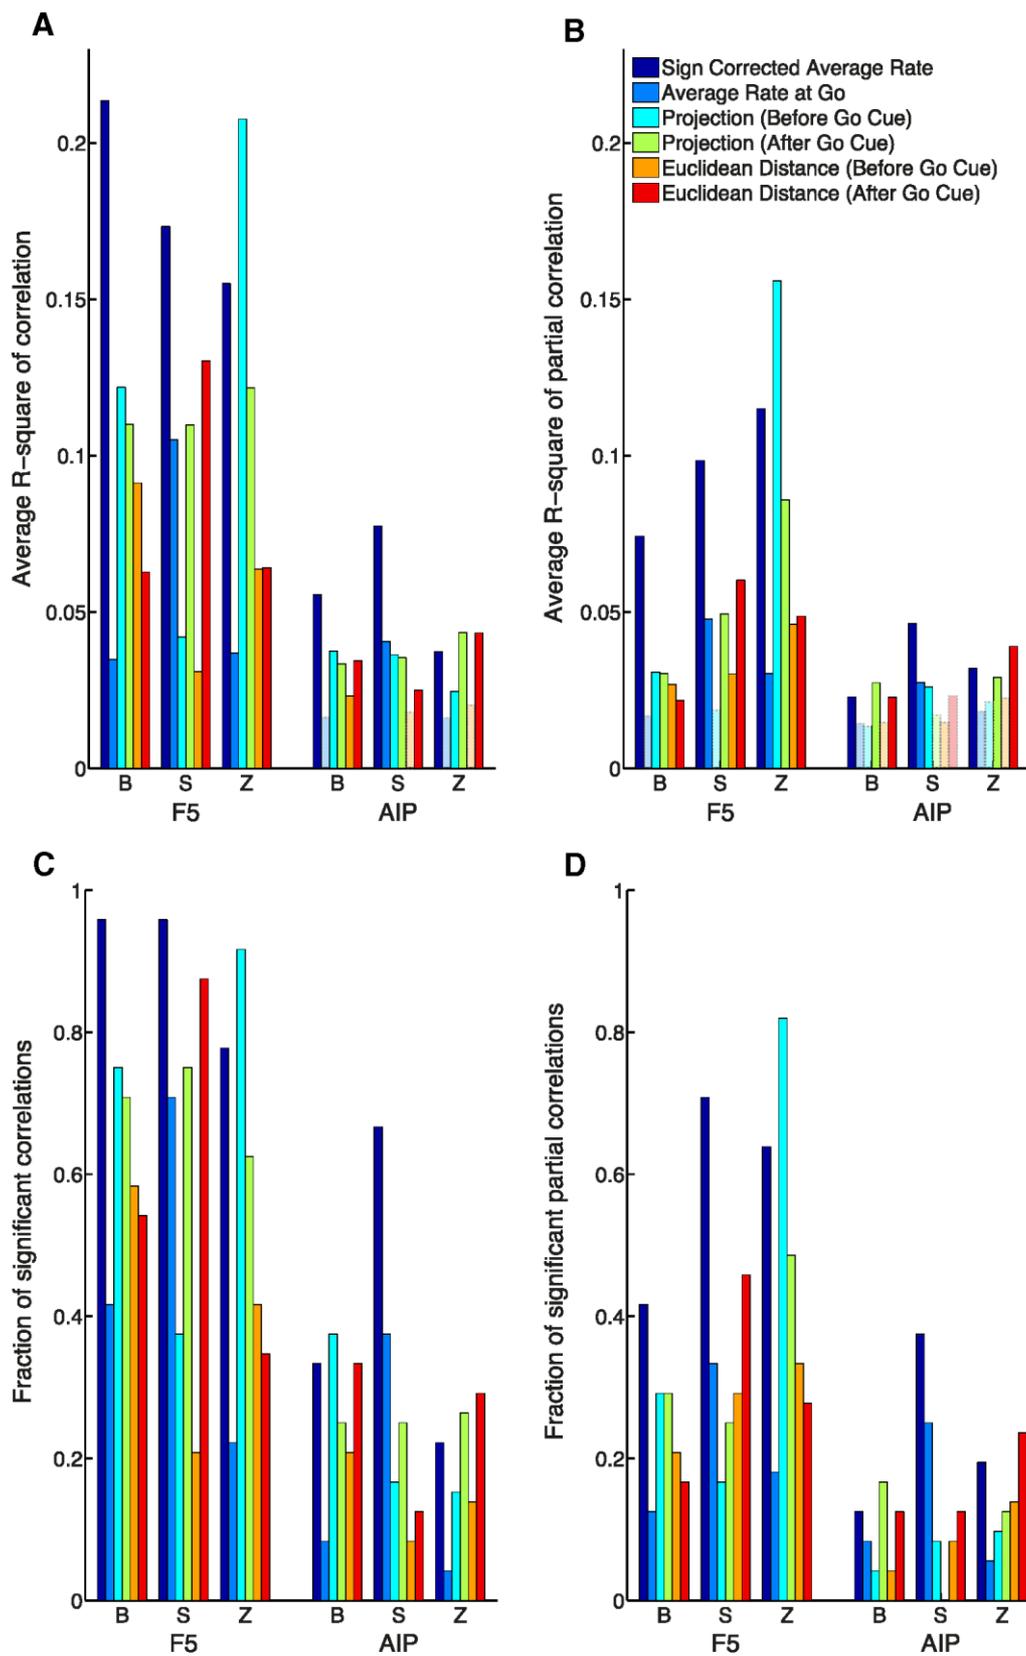
Figure 6. Comparison of prediction performance and fraction of significant full/partial correlations between predictors and RT over all task conditions, datasets, and brain areas. **A**, Average fraction of RT variance explained by correlation. **B**, Average fraction of RT variance explained by partial correlation. Significant results are illustrated as solid bars, whereas the open bars show results that can be explained by chance ($p \geq 0.01$). **C**, Fraction of conditions with significant correlations ($p < 0.05$). **D**, Fraction of conditions with significant partial correlations ($p < 0.05$).

fit a set of weights to all units (MATLAB function *regress*). The results of this analysis, when cross-validated identically to the main analysis (twofold), show that a linear regression over all individual units can explain 3.9% of variance in RT in F5 and 2.2% in AIP, with a significant advantage of F5 over AIP ($p <$ 0.001, Kruskal–Wallis ANOVA). However, this performance is only one-quarter of the projection or SCAR methods overall. Because the number of available units frequently outnumbers the number of available trials, coefficients cannot be ideally identified. For this reason, the regression often excluded up to 15% of the units in each dataset by assigning them a coefficient of zero.

To deal with the small number of trials available, it is also possible to use stepwise linear regression to add or remove units based on how their inclusion affects the model (MATLAB function *stepwisefit*). To produce an optimal solution, the model was initialized with only a constant term and units were subsequently added if they significantly improved the model (F statistic, $p <$ 0.05). The results of this analysis show that a stepwise linear regression over all individual units can explain 12.1% of variance in RT in F5 and 3.4% in AIP, with a significant advantage of F5 over AIP ($p <$ 0.001, Kruskal–Wallis ANOVA). However, in this analysis, between 77 and 97% of units were excluded from the model to produce an optimal fit. Together, the linear regression results are consistent with previous analyses showing an advantage of F5 over AIP and are similar to the results obtained by selecting units by variance over trials (see Fig. 8). However, their usefulness is limited, at least in datasets with a restricted number of trials.

Removing the effect of the memory period does not eliminate RT prediction

The length of the memory period was strongly negatively correlated with RT in all tasks and animals (Fig. 2). To ensure that a straightforward encoding of the memory period in the firing rates of individual units was not responsible for our findings, all methods were retested using partial correlation. As described in Materials and Methods, partial correlation allows for the correlation of two variables while controlling for the linear effects of one or more additional variables. Here we controlled for the effect of memory period length on RT. Figure 6B shows the mean R^2 over all datasets while controlling for the effect of memory period length. Partial correlation reduces the performance of all methods, but almost all methods remain above chance level in F5. In AIP, all methods are reduced to chance level in at least one animal, with the exception of the SCAR method. The largest reduction in performance caused by partial correlation was 66% over all methods in animal B, suggesting a strong reliance on the memory period length and consistent with the unexpected direction of the shift in the correlation coefficient distribution of animal B in Figure 4C. The smallest reduction in performance was 25% in animal Z. In AIP, results of each animal were never reduced by >38%. For comparison, the mean R^2 using the standard correlation metric is shown in Figure 6A. Similarly, the number of significant correlations was reduced when using partial correlation as illustrated in Figure 6D.

Anterior AIP outperforms posterior AIP

A number of recent studies have highlighted that the anterior (aAIP) and posterior (pAIP) subdivisions of AIP differentially encode visual task parameters (Baumann et al., 2009; Romero and Janssen, 2014) and differ drastically in their effective connectivity (Premereur et al., 2015). Because it is not well understood how these two areas differ in their contribution to preparatory activity for grasping, we further segregated our units into aAIP and pAIP corresponding to the ante-

rior and posterior implanted arrays, respectively, and repeated the main analyses.

Unlike the comparison between F5 and AIP, the number of units recorded on each array within AIP differed significantly for all animals ($p < 0.05$, Wilcoxon's signed-rank test). Therefore, for each dataset, units were discarded randomly from the larger set until an equal number of units were present from each subarea (stratification).

If the same RT prediction methods used in the main analysis are applied to subdivisions of AIP, there is a small, but significant, advantage of aAIP over pAIP ($p = 0.021$, Kruskal–Wallis ANOVA). Most of this advantage comes from the projection (pre-go) method, with an average R^2 of 0.031 in aAIP and 0.019 in pAIP ($p < 0.01$, Wilcoxon's signed-rank test). In agreement with the main results, there was no significant difference in RT prediction between grip conditions or cross-validation folds ($p = 0.36$ and $p = 0.86$, Kruskal–Wallis ANOVA). These findings are in line with the emerging view that a gradient of visual to motor processing exists between pAIP and aAIP.

Multiple regression does not improve RT prediction

By combining multiple prediction methods in a multiple regression, it is possible to capitalize on the potential orthogonality between different predictors. To test whether a multiple regression could increase overall prediction of RT, we first replicated the regression described by Afshar et al. (2011), which consists of a regression of the pre-go and post-go cue versions of the projection and the velocity projection methods. Because the velocity projection method performed poorly in our analysis, it was not expected for this regression to significantly improve RT prediction. In fact, this four-factor multiple regression only outperformed simpler unimodal and bimodal regressions consisting of subsets of these factors in 16.7% of all datasets in F5 and 11.1% in AIP (F test). Furthermore, this regression never achieved a lower BIC score than more parsimonious regressions in any dataset or brain area, suggesting that combining these four factors in a regression is not justified in our dataset.

A number of other regressions were tested, but in no case were >50% of datasets in F5 and 16.7% of datasets in AIP able to significantly outperform simpler regressions (F test). Furthermore, none of these multiple factor regressions achieved a lower BIC in >11.1% of datasets in F5 and in none of the datasets in AIP.

Because multiple regression performs best when individual variables are independent, it would be unlikely to explain significantly more variance in RT if our predictors are highly correlated. In fact, most methods are highly correlated with one another in our dataset (minimum $R^2 > 0.14$), with the exception of the velocity projection method ($R^2 < 0.03$), which performed poorly in the main analysis.

No alternative reference point can outperform SCAR

The SCAR method relies on first correlating the firing rate of each unit with RT and then inverting based on the resulting correlation coefficient. Because this method relies on cross-validation, it would be preferable to perform a method that does not rely on previous information. To ensure that this alternative was not possible, a control was performed. The mean firing rate at multiple time points (up to 2 s) before the go cue was subtracted from the firing rate of each single trial, and the absolute value of the resulting signal was taken. Subsequently, the firing rate on each trial was averaged over units and correlated with RT. This method has the effect of inverting the activity of each unit relative to the mean firing rate at some previous time point. In no case
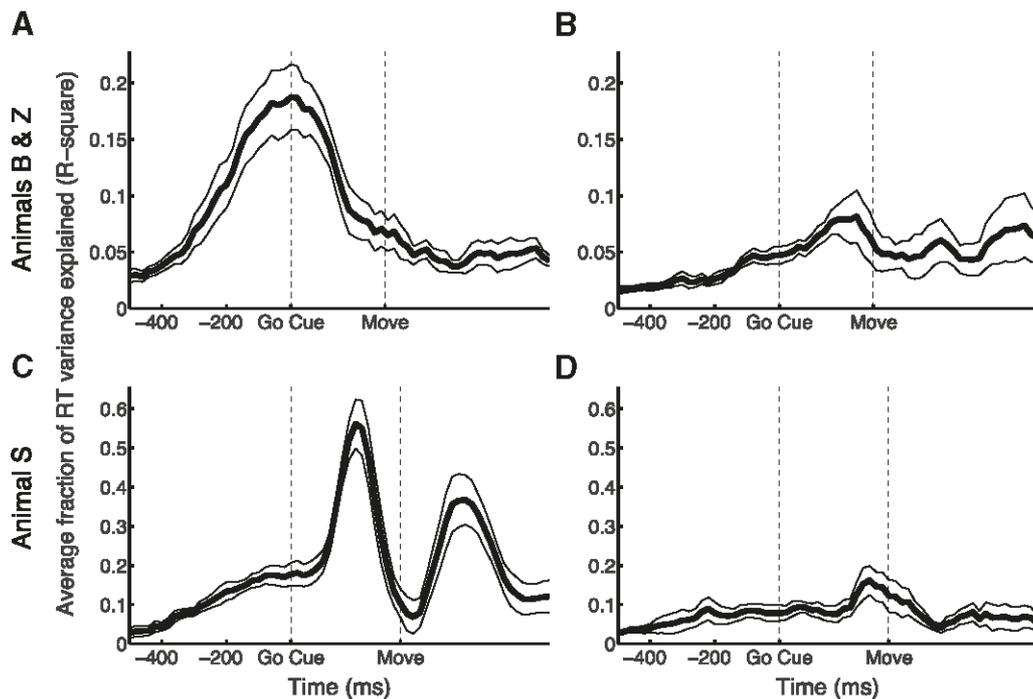
Figure 7. Stability of the sign-correction vector determined at the go cue by the SCAR method. RT prediction is calculated using sign-corrected neural activity around each time point. *A*, *B*, SCAR method as a function of time for animals B and Z in areas F5 (*A*) and AIP (*B*). *C*, *D*, SCAR method as a function of time for animal S in areas F5 (*C*) and AIP (*D*). Dashed lines indicate the go cue and the median movement onset (Move). Note the difference in peak RT prediction in F5 between animals B and Z and animal S.

was such a method able to explain more or equal variance in RT than the SCAR method, suggesting that there exists no trivial alternative to pretesting each individual unit with respect to measured RT. The same was true if grip cue-aligned activity was used for reference.

Consistency of the sign-correction vector during movement
Because the SCAR method relies on previous information of the relationship between firing rate and RT, we wondered whether the learned sign-correction vector, which is used to invert the activity of positively-correlated units, revealed a natural alignment of firing rates across time or whether it simply fit the data well at the time of the go cue. To test this, we used the sign-correction vectors determined in the main analysis to repeat the SCAR method using neural activity not just from the go cue but also at variable time points from 500 ms before to 800 ms after the go cue. The results of this analysis are depicted in Figure 7. As can be seen in Figure 7C for animal S in F5, maximal RT prediction is achieved shortly before median movement onset ($R^2 = 0.56$). If neural activities of many units peak shortly before movement onset, then the SCAR method should perform best at this time. Such a result would suggest that trials on which activity drifted toward the firing rate observed at movement onset were more likely to be trials with a short RT, in line with the rise-to-threshold hypothesis (Erlhagen and Schöner, 2002).

However, as can be seen in Figure 7A, the results of animals B and Z differed significantly in F5. In this case, the peak RT prediction occurs precisely around the time of go cue. In contrast to the results of animal S, a peak at the go cue suggests that, although our sign correction was able to properly align the firing of each unit at the go cue, it does not necessarily represent a consistent pattern in the firing of the underlying units.

In AIP, peak RT prediction was achieved in all animals shortly before median movement onset, i.e., a higher (sign-corrected)

firing rate on single trials tended to lead to shorter RTs. Maximal RT prediction before movement onset can be explained by the idea that activity either rises during the memory period to achieve peak activity during the movement or that activity is decreased during the memory period to reach a minimum during movement. This result in AIP is consistent with a study showing significant RT prediction from activity in AIP shortly before movement onset (Verhoef et al., 2015).

Variance selection allows high performance with a subset of units
Because all recorded units were included in the previous analysis, we were curious whether a subset of units could be selected that performed equally well or better than the entire population. To test this, a variance selection of units was performed. Units were discarded from the analysis in order of increasing variance in spike count (at the go cue) across trials. This way, units with higher variances were preferentially included. For the two best performing methods, SCAR and projection (before go cue), the variance selection performed significantly better than chance ($p < 0.05$, Bonferroni's corrected) in F5 for all animals (Fig. 8A,C). In AIP, only variance selection using the SCAR method outperformed chance (Fig. 8B,D).

In all cases, selecting units by variance did not improve maximal performance, as expected. In fact, when comparing performance using all units to a smaller subset in F5, using a variance-selected subset of only 32 or 18% of recorded units, for the SCAR and projection (before go cue) methods, respectively, suffered only a 5% decrease in performance. For the SCAR method, it was only necessary to use a subset of 23% of the available units in F5 to attain 95% of maximal performance. Together, these results suggest that, when units are selected by variance at the go cue, only relatively small subsets of the recorded units are required to attain virtually maximal performance. More impor-
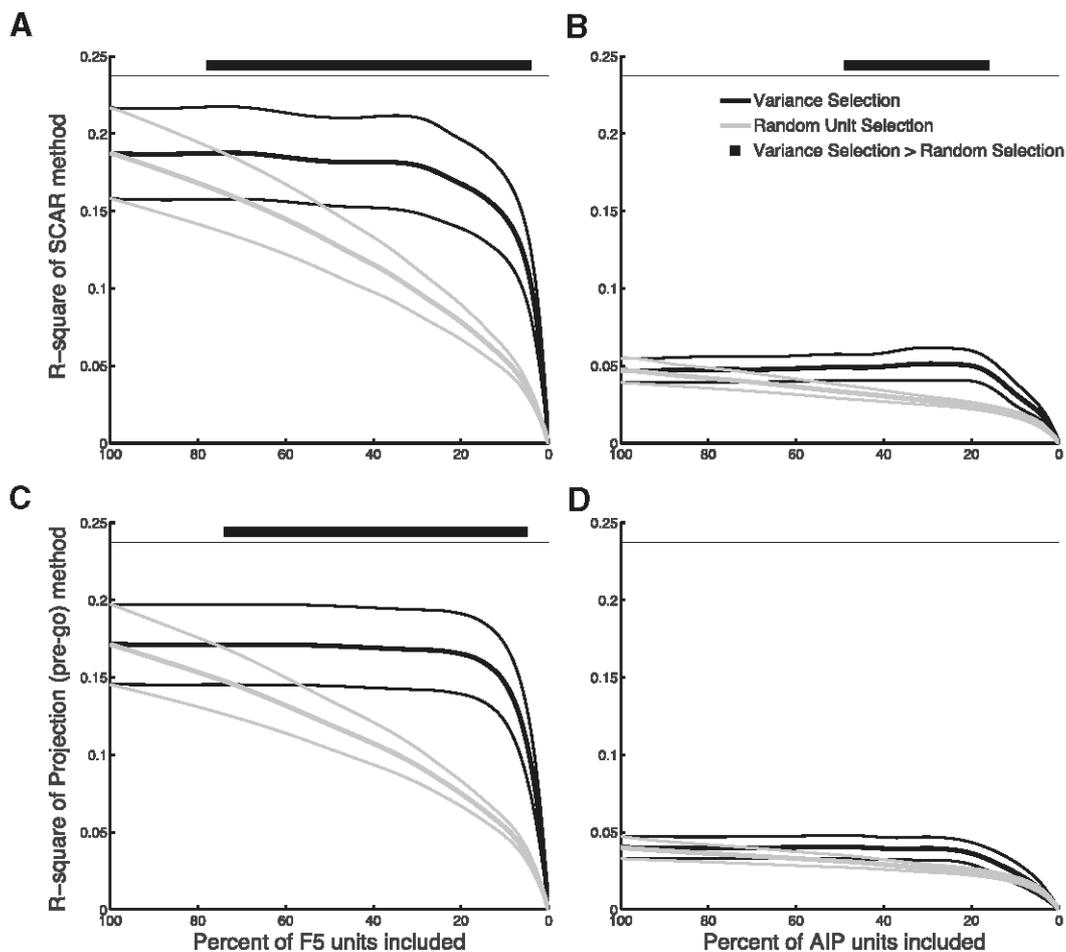
Figure 8. Selection of units by firing rate variance at the go cue for the two best performing methods in animals B and Z. *A*, *B*, Variance selection of units versus random selection using the SCAR method in F5 (*A*) and AIP (*B*). *C*, *D*, Variance selection of units versus random selection using the projection (before go cue) method in F5 (*C*) and AIP (*D*). Horizontal black bars on top represent unit percentages in which the variance selection performed significantly better than random selection ($p = 0.05$, Bonferroni's corrected, permutation test).

tantly, including all units in the population does not appear to add noise to these methods, because maximal performance is achieved when including all units, suggesting that they properly describe the relationship between preparatory activity and RT at the population or network level.

Variability of RT axis from day to day and animal to animal
When considering each dataset separately, it became clear that the day-to-day variability in RT prediction for each method is relatively high (Fig. 5). To elucidate how neural trajectories, which are presumably very similar over sessions, could explain very different amounts of variance in RT, we visualized a few individual sessions using PCA. This second dimensionality reduction method was introduced, in addition to the GPFA used in Figure 2A, to visualize average trajectories as opposed to single trials. In contrast to GPFA, which applies many different and sometimes large smoothing kernels, PCA allows more direct control over the amount of smoothing over time. In Figure 9, the first two principal components of individual conditions of individual recording sessions are shown. The mean trajectory over all trials is depicted along with the mean trajectory of trials binned into slow, medium, and fast RT trials. In every subplot, a visualization of the projection (after go cue) method is presented from the data. In this visualization, the position of single trials along the dashed projection axis would determine our measurement for

how far along the mean neural trajectory this trial is. The subsequent length of the projection of each single trial onto this axis would then be used to predict RT. In Figure 9A, trajectories of a power grip condition are shown from dataset B140509. It appears that the fast and slow RT trials are located distantly to each other along the projection axis, suggesting that this axis would be valuable in explaining trial-to-trial RT variability. This was in fact the case, because the projection (after go cue) method was able to explain 27% of the variance in RT in the main analysis of this dataset. However, note that the position of the fastest RT trials is less far along the mean trajectory than slow trials, directly contradicting the predictions of our hypothesis. As we noted in Figure 6B, much of the RT prediction obtained in animal B was eliminated by controlling for the effect of memory period length. Based on the trajectory in Figure 9A, it seems that trials with longer memory periods tended to continue along the projection axis instead of lingering near the mean trajectory. Because longer memory periods led to slower RTs for the most extreme memory period lengths (1300 ms), trials that have progressed farther along the mean indicated slower RT trials.

Plotted in Figure 9B is the mean trajectory of the precision grip on the very next dataset (B140515). The mean trajectory for this condition is very similar to that of Figure 9A. However, the orientation of the projection axis is approximately orthogonal to that of an axis running through the slow and fast RT trials, sug-
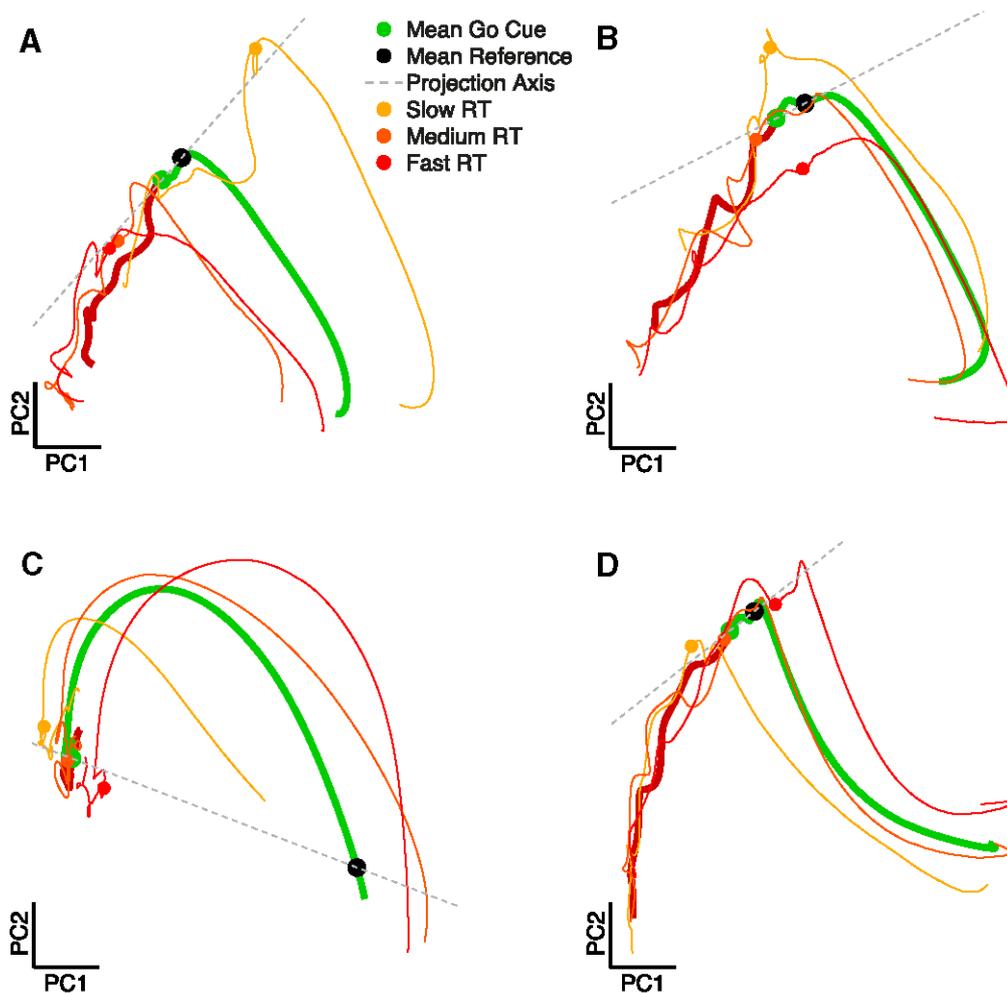
Figure 9.　Low-dimensional representation (PCA) of neural state space illustrating day-to-day and animal-to-animal variability. Trajectories are plotted in the two principal components of area F5 explaining the most variance. *A*, Trajectory of a power grip from dataset B140509 in which RT is most variable along the projection axis. *B*, Trajectory of a precision grip from dataset B140515 in which RT is most variable orthogonal to the projection axis. *C*, Trajectory of a precision grip from dataset S120913 in which RT is most variable on the shared distance and projection axis. *D*, Trajectory of a power grip from dataset Z120921 in which RT is most variable along the projection axis.

gesting that this axis would explain only a small amount of variability in trial-to-trial RT. The projection (after go cue) method performed at chance level for this dataset, only explaining 3% of the variance in RT. Therefore, it seems plausible that, even when trajectories are similar, it is possible for RT variance to be rotated relative to the mean trajectory, suggesting that the mean trajectory may not always be an ideal reference. In this case, the Euclidean distance (after go cue) method performed significantly better because trials with shorter RT were located closer to the movement onset state in the state space.

The trajectory in Figure 9C shows a precision grip from dataset S1209013. It seems to differ substantially from the other trajectories. In this case, there is no clear progression of the preparatory trajectory near the go cue, and there is also no abrupt change in the directionality of the trajectory after the go cue. Therefore, it is not surprising that both projection methods performed quite poorly on this dataset. Only by selecting a Δt for the projection (after go cue) method that was quite large (210 ms) could improve RT prediction. In this case, the projection (after go cue) and Euclidean distance (after go cue) methods performed similarly, which is not surprising because projection and distance become mathematically similar for large Δt values.

Figure 9D represents an ideal trajectory of a power grip from dataset Z120921. In this case, trials that are farther along the projection axis correspond to trials with shorter RTs, in line with the predictions of the projection method.

Discussion

Using simultaneous neural recordings from three animals, we have shown that preparatory activity in both premotor and parietal cortices is correlated with trial-to-trial variability in RT. However, the activity in F5 is far more predictive of RT than in AIP. Although the length of the memory period facilitated RT predictability, our findings cannot be explained purely based on this relationship. The use of a state space framework, made possible by the parallel recording of many units, represents a major step forward in understanding the relationship between preparatory activity and behavioral parameters.

Trial-to-trial RT prediction

Although response characteristics and tuning properties of AIP and F5 neurons can be very similar (Baumann et al., 2009; Fluet et al., 2010), we have shown that their trial-to-trial relationship with RT differs greatly (Fig. 5). The current result is not trivial, because F5 and

AIP are densely and reciprocally connected (Luppino et al., 1999; Borra et al., 2008). However, the level of RT predictability found in the preparatory activity of AIP units is consistent with previous studies in nearby areas (Janssen and Shadlen, 2005; Snyder et al., 2006).

Together, the relative advantage of F5 over AIP is not altogether surprising given the fact that F5, which has projections to the spinal cord (He et al., 1993; Borra et al., 2010) and a facilitation effect on M1 (Shimazu et al., 2004), must naturally be involved in the transition between preparation and movement execution.

Comparing prediction methods

High performance of the projection method, matched only by the SCAR method, is consistent with the "initial conditions" hypothesis formulated by Afshar et al. (2011) from activity in PMd. Even after controlling for the effect of memory period length, the projection method still performs above chance level; however, the SCAR method can outperform the projection method, especially in Task 1. SCAR offers an alternative explanation for the relationship between single-unit firing and RT. In this framework, most individual units have a consistent relationship with RT, i.e., higher or lower firing rates before the go cue are associated with shorter RT. Controlling for the sign of this relationship was able to increase RT prediction up to four times and follows well from the observation that the preparatory activity in PMd is both positively and negatively correlated with RT (Riehle and Requin, 1993), as in F5.

Could subpopulations of these units explain prediction of RT? If this were the case, we would expect units that fire more during the delay would continue to rise during the movement. In two of three animals, the SCAR method peaked in RT prediction at the go cue, suggesting that this sign correction was a local property and not a consistent property of each unit (Fig. 7A,C). Indeed, activity is often higher in the delay period than during the movement (Crammond and Kalaska, 2000), suggesting that the relationship between firing during preparation and movement is complex (Churchland and Shenoy, 2007b).

Interestingly, maximal performance is always achieved for the projection and SCAR methods when including all recorded units and not a variance-selected subset in F5 (Fig. 8), supporting the conclusion that both methods accurately describe population-level features and are not simply dominated by specific subpopulations of units.

Previously, the best performing method was a multiple regression of projection and velocity projection components (Afshar et al., 2011). We did not find significant performance of the velocity projection method or any multiple regression. When examining our neural trajectories, it seems that in many cases the speed of change in neural signal remained high or even increased during the memory period, especially in Task 1 (our unpublished data). This may represent an interesting quality of F5 activity that differs from activity in PMd. Such memory-related activity could mask relationships between trial-to-trial neural velocity and RT, especially after factoring in the length of the memory period.

Differences between PMv and PMd

To our knowledge, the preparatory activity recorded in F5 in our study explains more trial-to-trial variance in RT than any other published study. However, the results obtained in nearby PMd are quantitatively comparable (Afshar et al., 2011). A number of studies have systematically contrasted PMv and PMd (for review, see Hoshi and Tanji, 2007). It may be that F5 is more involved in the specific timing and execution of reaching movements than PMd, as evidenced by chemical inactivation (Kurata and Hoff-man, 1994). Nevertheless, both PMv and PMd are essential for grasping movements (Raos et al., 2004, 2006). Electrical microstimulation in PMd during preparation (Churchland and Shenoy, 2007a), and potentially F5 (Gerits et al., 2012), delays movement onset.

Although PMd and PMv are part of relatively distinct parieto-frontal networks, they both have an important effect on behavioral timing. Additionally, both PMd and PMv project to similar locations within M1 and lack a clear hierarchy (Dum and Strick, 2005), suggesting that their roles are complementary and not sequential.

Limitations

To rule out premature muscle contractions as an explanation for RT prediction obtained during a delay, electromyographic recording of relevant muscles has been used in the past (Churchland et al., 2006c; Afshar et al., 2011). Such recordings were not undertaken in the current study; however, we do not believe that premature muscle contractions are a likely cause of the RT prediction observed here for three reasons. First, the hands of all animals remained completely still on the hand rest buttons until after the go cue had been given, as confirmed by infrared monitoring. Second, the RTs of all animals were well above 200 ms in at least 97% of trials, suggesting that they appropriately awaited the go cue. Third, animals successfully withheld movement during the catch trials, suggesting that they were properly awaiting the go signal.

Although the primary interest of the current study was grasping actions, all movements included a large reaching component as well. It remains a possibility that the relative advantage of F5 over AIP could in part be attributable to a larger role of F5 in reaching than AIP. However, previous studies dissociating reaching and grasping have shown that PMv is greatly involved in the representation of grasping without a reach component (Hepp-Reymond et al., 1994) and is potentially even less involved in reach encoding than AIP (Lehmann and Scherberger, 2013). Therefore, finding higher RT prediction accuracy in F5 rather suggests a larger influence of the grasping component in the neural signal.

Implications for models of motor preparation

It is clear that the most dominant factor in the neural trajectories of animal B is the length of the memory period itself (Fig. 9A), which seems to act counter to the notion of an optimal subspace, because trials do not congregate within an area of low variability. It has been shown that variability is decreased by external stimuli, which was observed in PMd (Churchland et al., 2006c) and a number of other cortical areas (Churchland et al., 2010). If F5 neurons were multiplexing many factors in addition to a motor plan such as anticipation of the go cue, similar to hazard rate (Janssen and Shadlen, 2005), or variability in attention over longer periods of time, trial-to-trial variability might be increased at go cue. Furthermore, encoding of the length of the memory period clearly increased RT predictability in F5 and AIP, as evidenced by the decrease in predictability when using partial correlation. Additional work is needed to determine the extent to which F5 and AIP encode cue anticipation or attention-related factors.

Alternatively, it could be that the subspace required to successfully complete the grasping movement is sufficiently large to allow trajectories to lie in a relatively wide space. The absence of a static prepare-and-hold state is consistent with the augmented view of the initial conditions hypothesis posited by Ames et al.

(2014), who showed that the memory state is bypassed in PMd when time to prepare an action is lacking. In this view, the subspace required to successfully complete an action, i.e., with no penalty in movement generation, but a possible penalty in RT, could be quite broad. However, it is clear that F5 firing rates do not necessarily congregate in a specific part of the state space given enough time, as would be predicted by an attractor model of preparatory dynamics. The interesting question of determining whether such a prepare-and-hold state is necessary in F5 or AIP, along with whether the observed preparatory processes set the initial conditions of a dynamical system, as they do in PMd and M1 (Churchland et al., 2012; for review, see Shenoy et al., 2013), are left to future works.

Recently, the ability to record activity from many neurons simultaneously has opened up new possibilities in the investigation of the motor and premotor cortices (for review, see Churchland et al., 2007). The current study explores the relationship between preparatory activity in large populations of neurons and subsequent behavior, shedding light on the differential role of parietal and frontal cortices in this process.

References

Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV (2011) Single-trial neural correlates of arm movement preparation. Neuron 71: 555–564. CrossRef Medline

Ames KC, Ryu SI, Shenoy KV (2014) Neural dynamics of reaching following incorrect or absent motor preparation. Neuron 81:438–451. CrossRef Medline

Baumann MA, Fluet MC, Scherberger H (2009) Context-specific grasp movement representation in the macaque anterior intraparietal area. J Neurosci 29:6436–6448. CrossRef Medline

Borra E, Belmalih A, Calzavara R, Gerbella M, Murata A, Rozzi S, Luppino G (2008) Cortical connections of the macaque anterior intraparietal (AIP) area. Cereb Cortex 18:1094–1111. CrossRef Medline

Borra E, Belmalih A, Gerbella M, Rozzi S, Luppino G (2010) Projections of the hand field of the macaque ventral premotor area F5 to the brainstem and spinal cord. J Comp Neurol 518:2570–2591. CrossRef Medline

Carpaneto J, Umiltà MA, Fogassi L, Murata A, Gallese V, Micera S, Raos V (2011) Decoding the activity of grasping neurons recorded from the ventral premotor area F5 of the macaque monkey. Neuroscience 188:80–94. CrossRef Medline

Churchland MM (2015) Using the precision of the primate to study the origins of movement variability. Neuroscience 296:92–100. CrossRef Medline

Churchland MM, Shenoy KV (2007a) Delay of movement caused by disruption of cortical preparatory activity. J Neurophysiol 97:348–359. CrossRef Medline

Churchland MM, Shenoy KV (2007b) Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. J Neurophysiol 97:4235–4257. CrossRef Medline

Churchland MM, Afshar A, Shenoy KV (2006a) A central source of movement variability. Neuron 52:1085–1096. CrossRef Medline

Churchland MM, Santhanam G, Shenoy KV (2006b) Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. J Neurophysiol 96:3130–3146. CrossRef Medline

Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV (2006c) Neural variability in premotor cortex provides a signature of motor preparation. J Neurosci 26:3697–3712. CrossRef Medline

Churchland MM, Yu BM, Sahani M, Shenoy KV (2007) Techniques for extracting single-trial activity patterns from large-scale neural recordings. Curr Opin Neurobiol 17:609–618. CrossRef Medline

Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, Finn IM, Ferster D, Ryu SI, Santhanam G, Sahani M, Shenoy KV (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat Neurosci 13:369–378. CrossRef Medline

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P,

Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. Nature 487:51–56. CrossRef Medline

Cramér H (1946) Mathematical methods of statistics. Princeton: Princeton UP.

Crammond DJ, Kalaska JF (2000) Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. J Neurophysiol 84:986–1005. Medline

Dum RP, Strick PL (2005) Frontal lobe inputs to the digit representations of the motor areas on the lateral surface of the hemisphere. J Neurosci 25: 1375–1386. CrossRef Medline

Erlhagen W, Schöner G (2002) Dynamic field theory of movement preparation. Psychol Rev 109:545–572. CrossRef Medline

Fetz EE (1992) Are movement parameters recognizably coded in the activity of single neurons? Behav Brain Sci 15:679–690.

Fluet MC, Baumann MA, Scherberger H (2010) Context-specific grasp movement representation in the macaque ventral premotor cortex. J Neurosci 30:15175–15184. CrossRef Medline

Gerits A, Farivar R, Rosen BR, Wald LL, Boyden ES, Vanduffel W (2012) Optogenetically induced behavioral and functional network changes in primates. Curr Biol 22:1722–1726. CrossRef Medline

Ghez C, Favilla M, Ghilardi MF, Gordon J, Bermejo R, Pullman S (1997) Discrete and continuous planning of hand movements and isometric force trajectories. Exp Brain Res 115:217–233. CrossRef Medline

Gozani SN, Miller JP (1994) Optimal discrimination and classification of neuronal action potential waveforms from multiunit, multichannel recordings using software-based linear filters. IEEE Trans Biomed Eng 41: 358–372. CrossRef Medline

Hanes DP, Schall JD (1996) Neural control of voluntary movement initiation. Science 274:427–430. CrossRef Medline

He SQ, Dum RP, Strick PL (1993) Topographic organization of corticospinal projections from the frontal lobe: motor areas on the lateral surface of the hemisphere. J Neurosci 13:952–980. Medline

Hepp-Reymond MC, Hüsler EJ, Maier MA, Ql HX (1994) Force-related neuronal activity in two regions of the primate ventral premotor cortex. Can J Physiol Pharmacol 72:571–579. CrossRef Medline

Hoshi E, Tanji J (2007) Distinctions between dorsal and ventral premotor areas: anatomical connectivity and functional properties. Curr Opin Neurobiol 17:234–242. CrossRef Medline

Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. Nat Neurosci 8:234–241. CrossRef Medline

Kakei S, Hoffman DS, Strick PL (2001) Direction of action is represented in the ventral premotor cortex. Nat Neurosci 4:1020–1025. CrossRef Medline

Kraskov A, Dancause N, Quallo MM, Shepherd S, Lemon RN (2009) Corticospinal neurons in macaque ventral premotor cortex with mirror properties: a potential mechanism for action suppression? Neuron 64: 922–930. CrossRef Medline

Kurata K, Hoffman DS (1994) Differential effects of muscimol microinjection into dorsal and ventral aspects of the premotor cortex of monkeys. J Neurophysiol 71:1151–1164. Medline

Kutas M, Donchin E (1974) Studies of squeezing: handedness, responding hand, response force, and asymmetry of readiness potential. Science 186: 545–548. CrossRef Medline

Lehmann SJ, Scherberger H (2013) Reach and gaze representations in macaque parietal and premotor grasp areas. J Neurosci 33:7038–7049. CrossRef Medline

Luppino G, Murata A, Govoni P, Matelli M (1999) Largely segregated parietofrontal connections linking rostral intraparietal cortex (areas AIP and VIP) and the ventral premotor cortex (areas F5 and F4). Exp Brain Res 128:181–187. CrossRef Medline

McQuarrie AD, Tsai CL (1998) Regression and time series model selection. Hackensack, NJ: World Scientific.

Murata A, Fadiga L, Fogassi L, Gallese V, Raos V, Rizzolatti G (1997) Object representation in the ventral premotor cortex (area F5) of the monkey. J Neurophysiol 78:2226–2230. Medline

Murata A, Gallese V, Luppino G, Kaseda M, Sakata H (2000) Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP. J Neurophysiol 83:2580–2601. Medline

Musial PG, Baker SN, Gerstein GL, King EA, Keating JG (2002) Signal-to-noise ratio improvement in multiple electrode recording. J Neurosci Methods 115:29–43. CrossRef Medline

National Research Council (2003) Guidelines for the care and use of mam-

mals in neuroscience and behavioral research. Washington, DC: National Academies Press.

Premereur E, Van Dromme IC, Romero MC, Vanduffel W, Janssen P (2015) Effective connectivity of depth-structure-selective patches in the lateral bank of the macaque intraparietal sulcus. PLos Biol 13:e1002072. CrossRef Medline

Quiroga RQ, Nadasdy Z, Ben-Shaul Y (2004) Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. Neural Comput 16:1661–1687. CrossRef Medline

Raos V, Umiltá MA, Gallese V, Fogassi L (2004) Functional properties of grasping-related neurons in the dorsal premotor area F2 of the macaque monkey. J Neurophysiol 92:1990–2002. CrossRef Medline

Raos V, Umiltá MA, Murata A, Fogassi L, Gallese V (2006) Functional properties of grasping-related neurons in the ventral premotor area F5 of the macaque monkey. J Neurophysiol 95:709–729. Medline

Riehle A, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. J Neurophysiol 61:534–549. Medline

Riehle A, Requin J (1993) The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex. Behav Brain Res 53:35–49. CrossRef Medline

Romero MC, Pani P, Janssen P (2014) Coding of shape features in the macaque anterior intraparietal area. J Neurosci 34:4006–4021. CrossRef Medline

Rosenbaum DA (1980) Human movement initiation: specification of arm, direction, and extent. J Exp Psychol Gen 109:444–474. CrossRef Medline

Schaffelhofer S, Agudelo-Toro A, Scherberger H (2015) Decoding a wide range of hand configurations from macaque motor, premotor, and parietal cortices. J Neurosci 35:1068–1081. CrossRef Medline

Shenoy KV, Sahani M, Churchland MM (2013) Cortical control of arm movements: a dynamical systems perspective. Annu Rev Neurosci 36: 337–359. CrossRef Medline

Shimazu H, Maier MA, Cerri G, Kirkwood PA, Lemon RN (2004) Macaque ventral premotor cortex exerts powerful facilitation of motor cortex outputs to upper limb motoneurons. J Neurosci 24:1200–1211. CrossRef Medline

Snyder LH, Dickinson AR, Calton JL (2006) Preparatory delay activity in the monkey parietal reach region predicts reach reaction times. J Neurosci 26:10091–10099. CrossRef Medline

Townsend BR, Subasi E, Scherberger H (2011) Grasp movement decoding from premotor and parietal cortex. J Neurosci 31:14386–14398. CrossRef Medline

Verhoef BE, Michelet P, Vogels R, Janssen P (2015) Choice-related activity in the anterior intraparietal area during 3-D structure categorization. J Cogn Neurosci 27:1104–1115. CrossRef Medline

Wise SP (1985) The primate premotor cortex: past, present, and preparatory. Annu Rev Neurosci 8:1–19. CrossRef Medline

Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M (2009) Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. J Neurophysiol 102:614–635. CrossRef Medline

49

2.2 – Probing the continuum of immediate to withheld grasping movements in the macaque fronto-parietal network

Authors:

Jonathan A. Michaels
Benjamin Dann
Rijk W. Intveld
Hansjörg Scherberger

"… 'motor preparation' may be more accurately defined as the engagement of a specific set of preparatory dynamics, rather than the achievement of a particular neural state. The set of states that are produced by these dynamics serve as initial conditions that are sufficient to generate a correct reach. "

(Ames et al., 2014)

Abstract

Our actions lie on a continuum between acting immediately and waiting for the right moment. Studying the interplay between planning and movement requires systematically varying preparation time. Two macaque monkeys performed a grasping task with a short instruction followed by variable time to go cue (0-1300 ms) while we recorded from many neurons in parallel from the ventral premotor cortex (F5) and the anterior intraparietal area (AIP), areas essential for grasp generation. Initial population responses passed through a fixed neural space, unique to each grip type, reflecting essential motor preparation. After this, AIP stabilized in a unique memory state while activity in F5 continued to evolve, providing a decodable signature of time. Intriguingly, in both areas activity during movement initiation separated into two groups corresponding to movements 'as fast as possible' and movements from memory, suggesting that withholding a movement causes a network-wide shift whose trace lasts throughout movement initiation.

Introduction

Some actions, such as quickly stopping to spill a cup of coffee, demand an immediate response. Others, such as waiting before a traffic light, require withholding our actions for the right moment. Most of our actions lie on the continuum between the two, and although many actions are carefully planned before they are executed (Kutas and Donchin, 1974; Ghez et al., 1997), we are often required to act with little or no warning. Various studies have examined how movements are planned and held in memory in the primate brain (Tanji and Evarts, 1976; for a review see Wise, 1985; Riehle and Requin, 1989; Murata et al., 1996; Cisek and Kalaska, 2002), but only a few have contrasted well planned movements with situations where little to no preparation is possible (Wise and Kurata, 1989; Crammond and Kalaska, 2000; Ames et al., 2014). None, to our knowledge, have systematically probed the transition between immediate and planned grasping movements in the behaving primate.

Understanding how specific brain areas contribute to movement planning requires being able to dissociate the neural preparation that occurs before a movement and the movement activity itself. Delayed movement paradigms in which actions must be withheld before they are executed have shown that preparatory activity in premotor and parietal cortex can be used to decode object properties and hand shapes (Townsend et al., 2011; Schaffelhofer et al., 2015), as well as arm and hand kinematics during movement itself (Menz et al., 2015), implicating them in reach and grasp generation. Furthermore, preparatory activity in the premotor cortex (Churchland et al., 2006b; Afshar et al., 2011) and parietal cortex (Michaels et al., 2015) is correlated with reach and grasp reaction time (RT), and perturbing this preparation state in premotor cortex delays subsequent movement (Day et al., 1989; Churchland and Shenoy, 2007; Gerits et al., 2012), a clear indication of a functional contribution to action planning.

While relating the responses of single neurons to behavior has been vital in the past, a neuron-by-neuron characterization cannot reveal the dynamics of whole brain regions, or how they interact with one another (for a review see Yuste, 2015). More and more studies show how task features are distributed over many neurons of a network (Raposo et al., 2014), which have been made possible by the increasing implementation of large-scale parallel recording and employing a state space framework of population activity (for a review see Cunningham and Yu, 2014). Under this framework, the firing of each neuron represents a dimension in a high-

dimensional space of all neurons where the firing of all neurons at a particular time represents a single point in the space of all potential network states. A recent study revealed that when reaches are cued immediately, the neural population in dorsal premotor cortex (PMd) does not need to achieve the specific state attained during fully planned movements (Ames et al., 2014), suggesting that successful reach preparation may be achieved through multiple neural trajectories. While contrasting immediate and fully planned reaches reveals important population features, to address the full continuum of preparation, as well as the complex interaction between planning and movement, a large range of preparation times must be investigated. Additionally, it is unclear how neural populations stabilize or maintain the plan of movements that must be executed from memory when visual cue information is only presented briefly. Crucially, to understand how the motor system encodes and executes movements, multiple distributed brain regions must be investigated, thereby providing a picture of how diverse areas differentially contribute to action.

To investigate how parts of the frontal and parietal lobes differentially encode the continuum of grasp preparation, we recorded large neural populations from the grasping circuit (Luppino et al., 1999) consisting of the hand area (F5) of the ventral premotor cortex (PMv) in the frontal lobe and the anterior intraparietal area (AIP) in the parietal lobe. While neural activity was recorded, two macaque monkeys performed a delayed grasping task, with a memory component, in which the amount of preparation time was systematically varied using 12 discrete delays (0-1300 ms). We found that during short delays, population activity takes shortcuts to movement onset, bypassing the states achieved during fully prepared grasps. However, initial trajectory (first 300 ms) was the same regardless of delay, but specific to each grip type, suggesting that this activity is required for successful movement. Furthermore, the entire trajectory of the short delays (0-200 ms) was the same throughout the entire trial, suggesting varying the presentation of the go cue within this range resulted in the same grasp execution.

While single unit activity was similar between the areas, there was a dramatic difference in population dynamics for fully planned grasps. Activity in AIP stabilized after 600 ms for long delays, but activity in F5 was highly dynamic and continued to evolve for the entire memory period. The signature of time could be decoded from activity of either area, but time decoding was significantly better in F5 specifically during the memory period, suggesting that F5 and AIP play different roles

during memory. Activity was broadly distributed in the state space at the time of go cue, but rapidly converged in the 200 ms preceding movement onset with F5 leading AIP by 60 ms. Interestingly, activity in both areas approached movement onset in two distinct trajectory clusters composed of delays shorter or greater than 500-600 ms, suggesting that a network-wide shift occurs when movements no longer occur immediately, but instead must be withheld and executed from memory.

Results

Task and behavior

We trained two macaque monkeys to perform a delayed grasping task, with a memory component, in which the amount of preparation time was systematically varied between non-delayed (0 ms) and a long delay (1300 ms) in 12 distinct increments. Monkeys B and S fixated a central point, received a grip cue (300 ms) corresponding to either precision or power grip, and were cued to perform this grip following a variable delay when the central fixation point turned off (Figure 1a-b). The grip cue appeared for 300 ms regardless of delay length, so for longer delays monkeys had to remember which grasp had been cued at the beginning of the trial. The performance of both monkeys was high, correctly completing trials after receiving grip information 95% and 98% of the time for monkeys B and S, respectively (Table 1). In addition to the normal task, we also randomly inserted no-movement trials to ensure that the monkeys always waited for the go cue before acting. In these trials monkeys were given a grip cue as normal, but were required to continue fixating and withholding movement for 2 seconds (8% of all trials), after which a reward was given for withholding movement. Both monkeys completed these added trials successfully (monkey B: 100%; monkey S: 97.7%).

	Trial Count	Correct Performance	Units Recorded in F5	Units Recorded in AIP
B1	485	91%	65	29
B2	685	96%	88	35
B3	586	96%	43	25
B4	814	96%	64	28
B5	775	96%	46	19
B6	745	97%	72	33
Average:	682	95.3%	63.0	28.2
S1	502	98%	124	134
S2	514	97%	136	148
S3	571	97%	142	137
S4	658	99%	121	97
S5	590	99%	115	104
S6	546	98%	156	165
Average:	564	98.0%	132.3	130.8

Table 1. Table of trial counts, performance, and number of units recorded for all data sets.

In addition to number of correctly executed trials, examining the reaction times (RTs) and movement times (MTs) of the monkeys can provide useful insight into the difficulty of the task and how reliably monkeys could perform. RT decreased steadily with increasing amounts of preparation (Rosenbaum, 1980), approaching a minimum after approximately 400 ms of preparation (Figure 1c)(Churchland et al., 2006b). RT tended to increase slightly for the longest delay, which is expected since it is hard to maintain focus for long periods of time. For monkey S, MT did not correlate with length of the delay period (Figure 1d), suggesting that although RT was slower for short delays, once movements were initiated, the movements were identical regardless of delay. In monkey B there was a small positive correlation between delay and MT ($r = 0.18$, Figure 1 – Figure Supplement 1), indicating that movements after longer delays were slightly slower. The number of errors showed no clear relationship to the length of the delay period (Figure 1 – Figure Supplement 2), and the number of errors was extremely low, providing evidence that the monkeys could complete all conditions equally well.
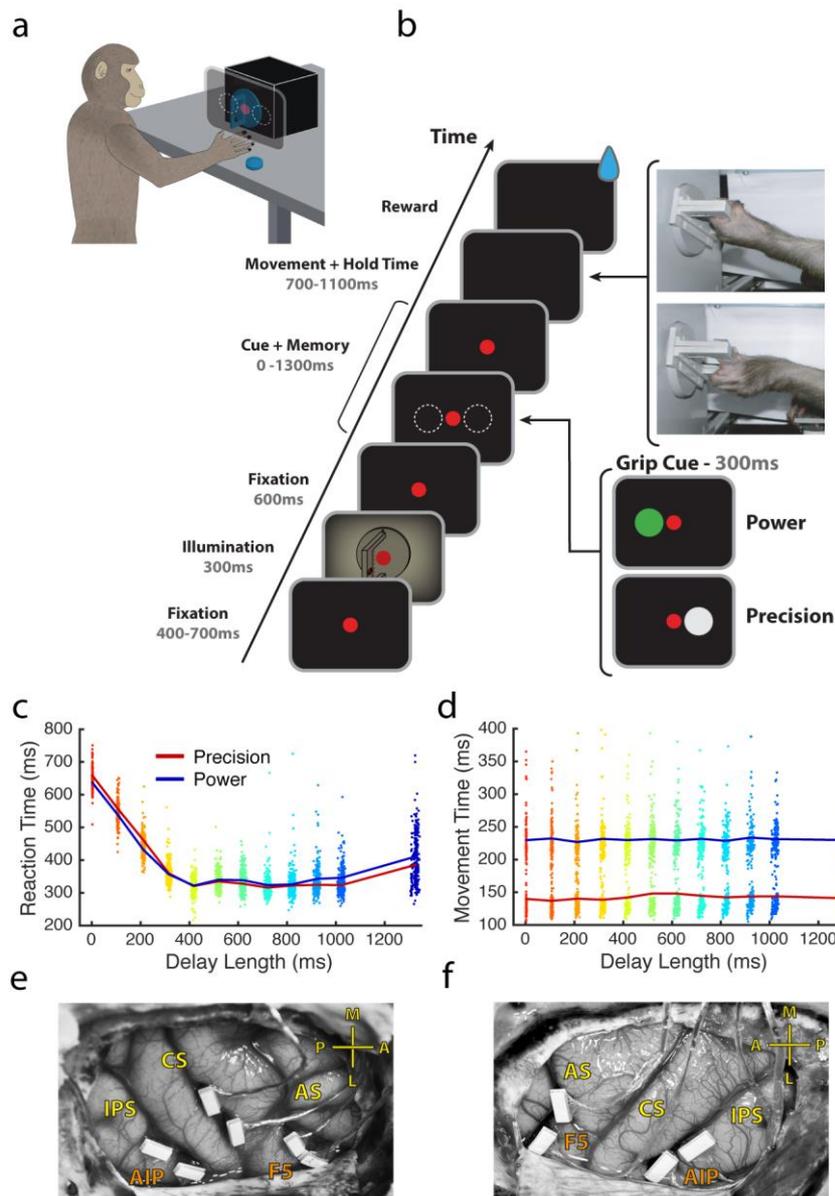
Figure 1. Task design, implantation, and behavior. (**a**) Illustration of a monkey in the experimental setup. The cues were presented on a masked monitor and reflected by a mirror such that cues appeared super-imposed on the grasping handle. (**b**) Delayed grasping task with two grip types (top: power grip, bottom: precision grip). Trials were presented in pseudorandom order in darkness and with the handle in the upright position. (**c and d**) Scatter plots of reaction time (**c**) and movement time (**d**) against delay length for monkey S. The solid line represents the mean for each delay bin. (**e and f**) Array locations for monkey B (**e**) and S (**f**). Two arrays were placed in F5 on the bank of the arcuate sulcus (AS) and two were placed in AIP toward the lateral end of the intraparietal sulcus (IPS). In monkey B two more arrays were placed on the bank of the Central sulcus (CS). The cross shows medial (M), lateral (L), anterior (A), and posterior (P) directions. Note that monkey S was implanted in the left hemisphere and monkey B the right hemisphere.

Neural responses

During the six recording sessions of each monkey, single- and multi-unit activity was recorded simultaneously in F5 and AIP using floating microelectrode arrays (Figure 1e,f; Materials and Methods). The number of units recorded in each session is summarized in Table 1. There were significantly more units recorded in area F5 of monkey B than in AIP (Paired *t*-test, p < 0.001), while there was no significant difference for monkey S (Paired *t*-test, p = 0.81). For all analyses we pooled single- and multi-units together (mean recorded per session: 75 single and 102 multi). We evaluated grip type tuning in both areas to ensure that the task successfully elicited task-related tuning. The average percentage of units tuned for grip type during the 200 ms following cue onset was 29% in F5 and 29% in AIP, 28% and 26% in the 200 ms before the go cue and 55% and 45% in the 200 ms following movement onset (*t*-test, *p* < 0.05), conservatively measured only for movements that could be fully planned (i.e. ≥500 ms delay). Amounts of grip tuning were very similar between monkeys and to previous studies of both F5 and AIP (Lehmann and Scherberger, 2013; Michaels et al., 2015; Schaffelhofer et al., 2015).
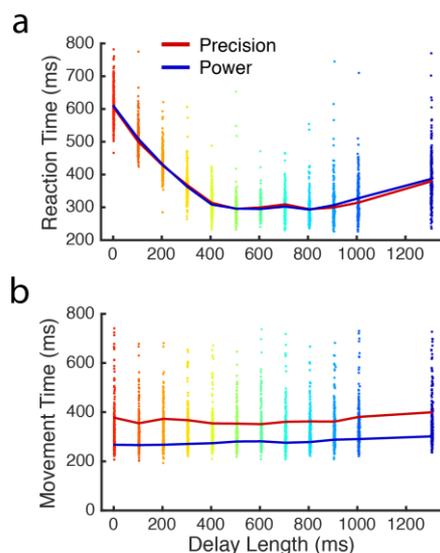
Figure 1 – Figure Supplement 1. Behavior of monkey B.

If the brain areas we investigated were specifically coding task-related visual features, we would expect similar responses to the grip cue regardless of whether grasps were cued immediately or not. Conversely, if single units were related to execution of the correct motor plan, we should observe similar neural responses during movement regardless of when go cues were presented. Interestingly, a wide

variety of mixed activity patterns were present in both areas, as shown by Figure 2, which directly compares the 1300 ms delay and the 0 ms delay for a number of example single-units. In many cases the initial cue response in the delayed condition was suppressed when the go cue appeared concurrently with the grip information (Figure 2a,e). In other cases, the initial cue response was present regardless of delay, but the activity quickly converged to an equivalent movement trajectory (Figure 2b,d,f,h). Other interesting responses were observed, such as a peak in activity during the memory period (Figure 2c), similar to hazard rate, i.e., the expectation of an event given that no event has occurred until now (Janssen and Shadlen, 2005), and activity during the movement period which differed between delayed and non-delayed grasps (Figure 2g), an interesting result returned to later. Figure 2d shows an example of strong grip type tuning after the cue and during movement, and Figure 2h shows an additional example where the pre-movement activity varies between delayed and non-delayed grasps. All of these diverse types of responses were present in both F5 and AIP, although movement period activity tended to be higher in F5. The broad variety of unit responses reveals a complex representation of non-delayed vs. delayed movement, making further analyses based on individual units insufficient for characterizing processing within these brain areas.
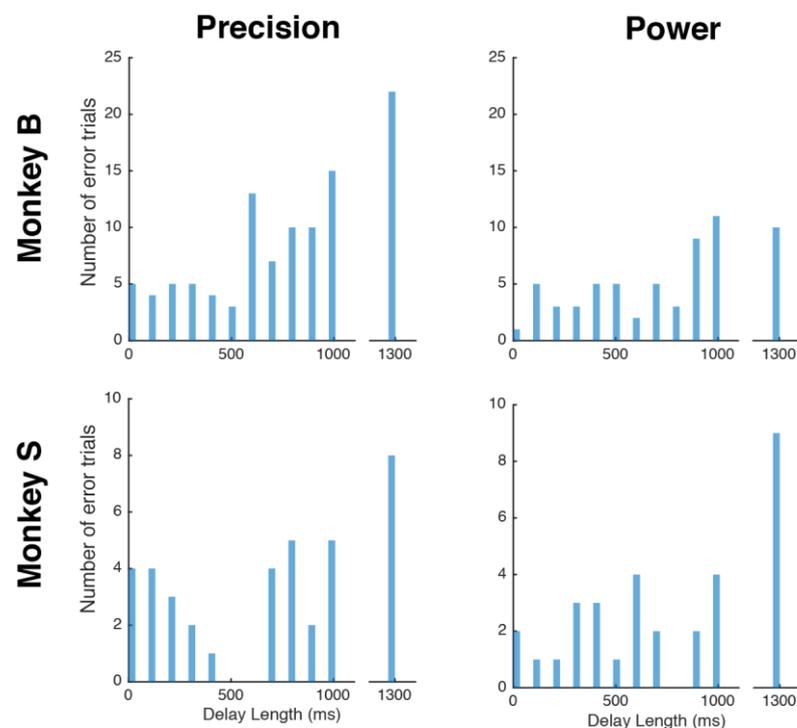
Figure 1 – Figure Supplement 2. Error trial distribution over all data sets.

Visualizing the population response

An alternative approach to single unit tuning is to use a state space framework, in which all units together are considered as a high-dimensional space in which the firing of each unit represents one dimension. In order to visualize the complex interactions between planning and movement, we normalized (soft-normalization, Materials and Methods) and then projected population activity of all units of each area, for all 12 delay lengths of one grip type, onto the first three principal components (PCs; Materials and Methods). Video 1 shows the neural trajectory of exemplar data sets in F5 (B5-Power) and AIP (S4-Power) from shortly before grip cue onset to shortly after movement onset. The first three PCs were able to explain between 62-75% of firing rate variance, suggesting that the first three components capture most of the dominant features in the entire population of units.
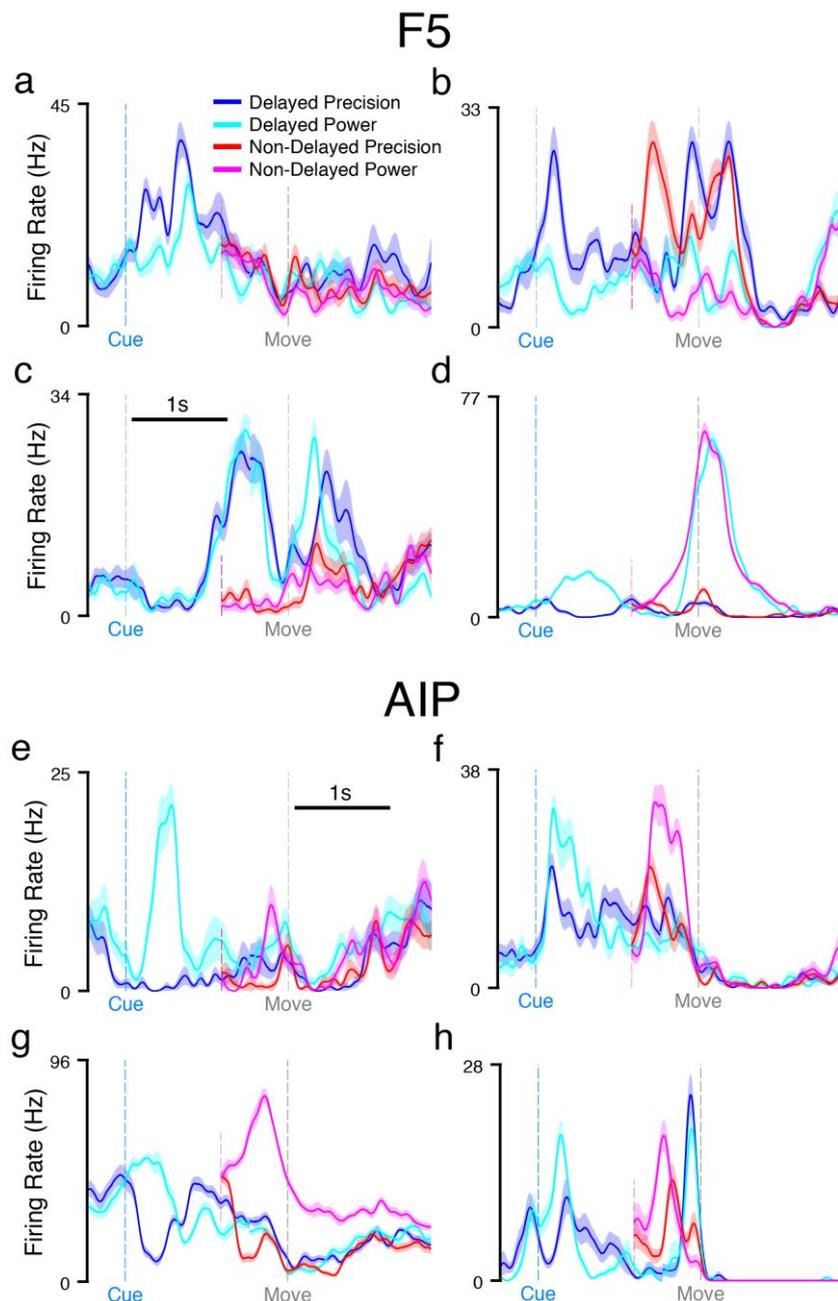
Figure 2. Example average firing rate curves of single-units for delayed (1300 ms) vs. non-delayed (0 ms) grasps. Data were aligned to two events, grip cue onset and movement onset. (**a-d**) Examples from area F5 showing (**a**) a suppressed cue response for the non-delayed condition in monkey B, (**b**) where the cue and movement responses are merged, (**c**) where the movement activity greatly differs between delayed and non-delayed grasps, (**d**) and an interesting additional example similar to the second example. All of these examples are from monkey B, but similar examples exist for monkey S. (**e-f**) Similar types of example single-units from AIP in monkeys B and S in the same order as the first section. Dotted red line represents approximate time of cue onset for non-delayed grasps. The cue was always presented for only 300 ms regardless of delay. Curves and shaded bands represent mean and standard error of the mean, respectively.

By following how the neural trajectories evolve in Video 1, we can generate a number of interesting hypotheses: (1) the states achieved by longer delays are completely bypassed for immediately cued grasps, in both areas, (2) the initial (first 300 ms) response to the cue is the same regardless of delay length or area, (3) the memory period activity in F5 changes continuously, never reaching a stable state, while the activity in AIP congregates in a stable state shortly after the cue, (4) activity at the go cue is very broad, but rapidly decreases in variance and converges towards a singular movement onset state, and (5) the activity during movement initiation is clustered into two groups, corresponding to delays shorter or longer than 500-600 ms. With this visualization as a guide, we now have a starting set of hypotheses to investigate.

The video can be viewed online at the following address:
http://www.jmichaels.me/docs/Video1.mp4

Video 1. Neural trajectory of all amounts of preparation in the first three principal components of F5 and AIP. (Left) Example data set from F5 of monkey B (B5-Power). (**Right**) Example data set from AIP of monkey S (S4-Power). Colors represent mean trajectory for each delay condition ranging from 0 ms (red) to 1300 ms delay (blue). To generate trajectories, data from each trial was aligned to two events, the grip cue presentation (magenta point) and movement onset (black point), averaged over all trials of each delay condition, and then interpolated to form a continuous trajectory. The grip cue always lasted for 300 ms regardless of delay. Trajectories begin 100 ms before the grip cue and end 50 ms after movement onset. The initial perspective is in the first two principal components and then rotates through the first three.
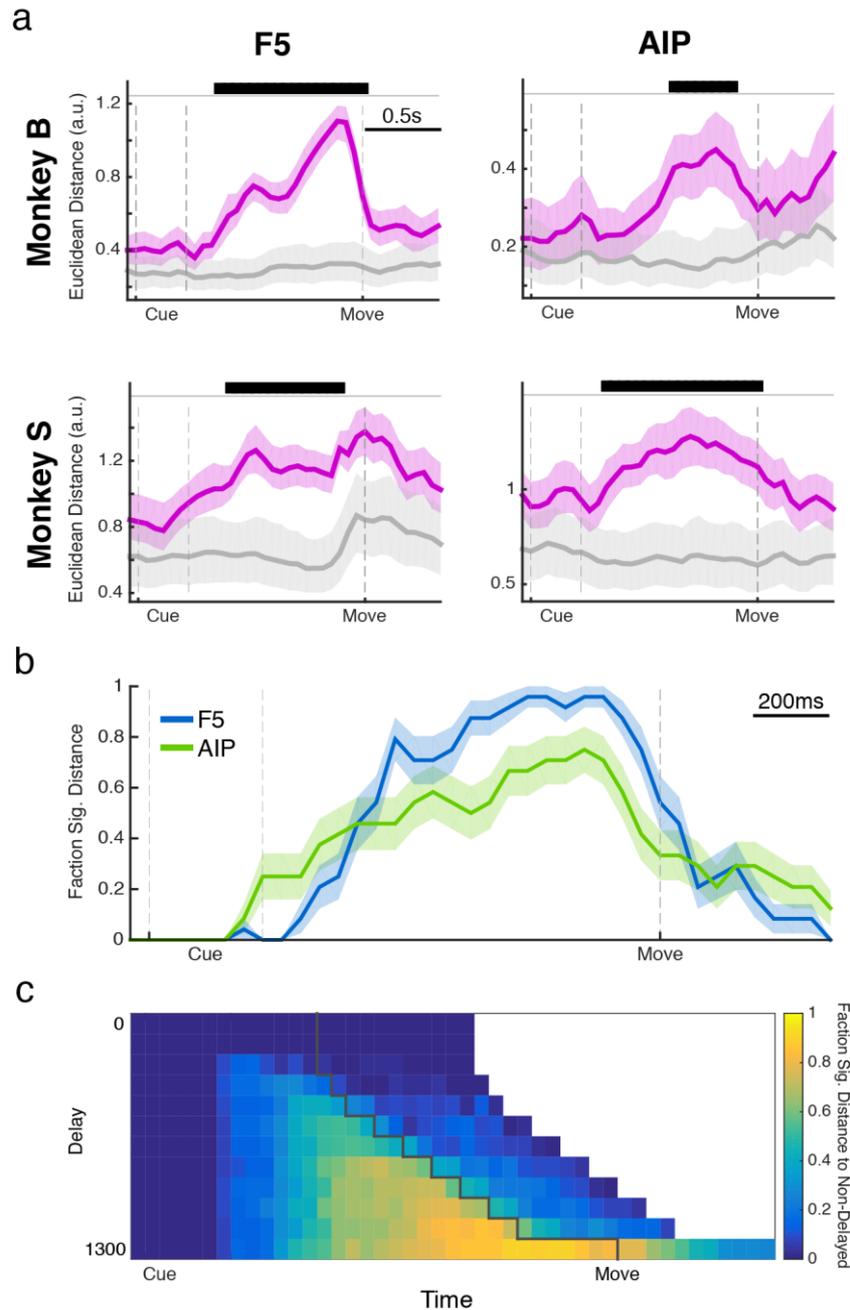
Figure 3. Point-to-curve distance between delayed and non-delayed trajectories.
(**a**) Minimum Euclidian distance between each time point on the delayed trajectory (1000 ms, in steps of 50 ms) and the entire non-delayed (0 ms) trajectory over time for 4 example data sets (F5: B5-Precision, S2-Power; AIP: B5-Power, S4-Precision) from both areas and monkeys. The magenta line represents the point-to-curve distance between the delayed and non-delayed trajectory, while the gray lines represent the same analysis when trials are drawn only from the delayed trajectory, acting as a chance level. Black bars along the top of plots denote times when the distance is significantly greater than chance level (Bootstrapping procedure with 1000 resamples, $p < 0.05$ Bonferroni-corrected for number of time points). Error bars represent the 5% and 95% confidence intervals generated by the bootstrapping procedure. (**b**) Fraction of significant distances (for the 1000 ms delay) over all data sets and conditions of both monkeys (12 data sets x 2 grip types). (**c**) Point-to-curve distance analysis

repeated for all time points on each delay (0-1300 ms) to the entire non-delayed (0 ms) trajectory reported as fraction significant distance over all data sets, grip types, and areas (12 data sets x 2 grip types x 2 area).

Required preparation and shortcuts to movement

When the monkey had no or only very little time to plan, the neural trajectory seemed to take a "shortcut" to movement initiation. To quantify if, and when, delayed and non-delayed trajectories diverged, we used a continuous distance analysis (Materials and Methods). We measured the minimum Euclidean distance (known as point-to-curve) between each time point on the delayed trajectory (1000 ms delay condition in steps of 50 ms) and the entire non-delayed (0 ms) trajectory in the full space of all units (i.e. without dimensionality reduction) to determine which points in the state space were traversed by both conditions and which were unique to longer delayed movements. This analysis was carried out separately for each grip type and data set.

Figure 3a displays the results separately for both monkeys and brain areas. After the cue, distance between the trajectories rose and stayed significantly above chance level (marked in gray) until around movement onset (Bootstrapping procedure with 1000 resamples, $p < 0.05$, Bonferroni-corrected for number of time points; Materials and Methods), at which point it decreased rapidly, suggesting that the trajectories during the movements were quite similar regardless of delay. Over all grip types and data sets the trend remains clear (Figure 3b), showing that maximal distance between the trajectories was maintained until near movement onset. These results statistically reinforce our first hypothesis that the population activity achieved during longer delays is unique to delayed grasps and is not achieved during non-delayed grasp and therefore not necessary to successfully complete the movement. The amount of divergence between the delayed and non-delayed trajectories was very similar in F5 and AIP, suggesting that when grasps are cued without a delay the neural population of both areas take shortcuts to movement, in agreement with our first hypothesis. Furthermore, our findings were not an artifact of our smoothing kernel, since a fixed spiking window of 100 ms yielded comparable results (data not shown). Due to differences in number of recorded units, the magnitude of the effect varied between monkeys (Figure 3 – Figure Supplement 1).
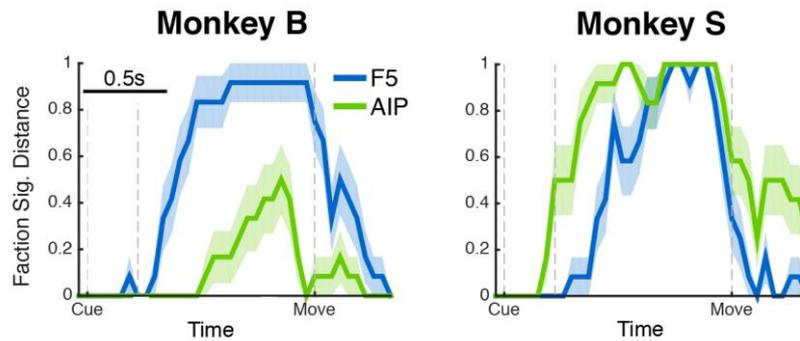
Figure 3 – Figure Supplement 1. Minimum point-to-curve distance separately by monkey.

Given that the delayed and non-delayed neural trajectories diverge, is there an area in the neural space that is traversed regardless of delay? In no data set or grip did the trajectories diverge before 300 ms after cue onset (Figure 3b), indicating that the initial response to the cue is identical whether or not the monkeys were given time to prepare. Therefore, it is likely that this time represents required processing, a point supported by the leveling off of the RT curve with preparation times exceeding around 300-400 ms (Figure 1c). Taken together, these results suggest that large portions of the state space that are traversed after the first 300 ms do not seem to be necessary for successfully executing grasping movements, in agreement with our second hypothesis. If we repeat the same distance analysis comparing all delayed trajectories with the non-delayed trajectory, we can gain a more detailed picture of how trajectories differ over varying amounts of preparation (Figure 3c). Interestingly, distance between delayed and non-delayed trajectories barely ever exceeded chance level for the shortest delays (100 and 200 ms) at any point during the trial, suggesting that neither the extra preparation time nor the difference in when the go cue was presented changed the trajectory at all.

Static and dynamic memory

Given that the trajectories of non-delayed grasps only overlap with longer delays for the first 300 ms of preparation, what is the function and dynamics of the memory period activity in longer delays? A striking feature of the visualization in Video 1 was that the F5 activity continually evolved throughout the course of the memory period, while activity in AIP seemed to congregate in an area of low variability. To analyze if the neuronal trajectory of the two areas stabilized during particular states, we

systematically compared the Euclidean distance between all pairs of time points along a trajectory for both the delayed (1300 ms) and no-movement trajectories (Figure 4, averaged over all data sets since results were very similar across grip types, monkeys, and sessions). For each time point, if the neural distance did not exceed a dynamic significance threshold, values were set to zero (Materials and Methods). Therefore, patches where the mean distance is zero (black) indicate that for no grip type or data set was there a significant difference between these two time points, while all other points contain at least one incidence of significance.

With the exception of the fixation epochs, all parts of the trial seemed to attain unique states that significantly differed from one another. The movement epoch was the most unique and dynamic epoch of the task, while the reward epoch and the time of cue offset stood out most in the no-movement condition. Most remarkably, the neuronal trajectory during the memory period in F5 is continuously and uniformly changing even though no event or movement occurs in this time. On the contrary, in AIP the neuronal trajectory stabilized about 300 ms after cue offset, representing a unique state, which is maintained throughout the memory period even after 2 seconds in the no-movement condition. These results suggest a considerably different code at the population level in AIP and F5.

In the current task, the expectation of having to perform a movement did not remain constant, since the probability of being in a no-movement trial increased with time spent in the memory period. Therefore, could it be that the dynamic nature of the memory period in F5 is due to this change in expectation? To rule out this possibility, we repeated the current analysis on data of a similar experiment in which movement expectation never changed during the memory period, due to the absence of a no-movement condition (Michaels et al., 2015). We found that even when the certainty of performing a movement remained constant, the same inter-area difference reported here was robustly present (Figure 4 – Figure Supplement 1), lending support to the observed dissociation between the two areas. These findings support our third hypothesis that population activity in F5 continually evolved throughout memory, while AIP activity became stable shortly after cue offset.
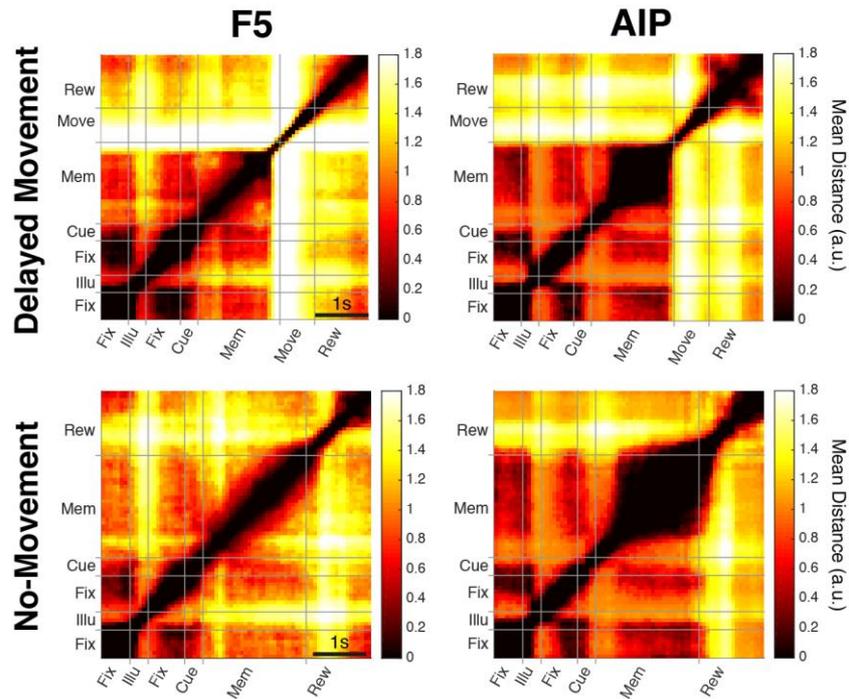
Figure 4. Uniqueness of neural state over the course of the trial. Mean Euclidean distance for delayed grasps (1300 ms) and no-movement trials between all pairs of time points over all data sets in F5 and AIP. For each pair of time points over all conditions (12 data sets x 2 grip types), distance results were tested for a significant difference using a bootstrapping procedure (6000 resamples in steps of 60 ms, $p = 0.01$, Bonferroni-corrected for number of time pairs). If a pair of time points did not significantly differ, their distance was set to zero. Therefore, all time points showing a distance of 0 did not significantly differ for any condition, while any value above zero showed significance for at least one condition. The abbreviations *Fix, Illu, Mem,* and *Rew,* correspond to the fixation, illumination, memory, and reward epoch, respectively. The median time reward onset was used for delayed movements.
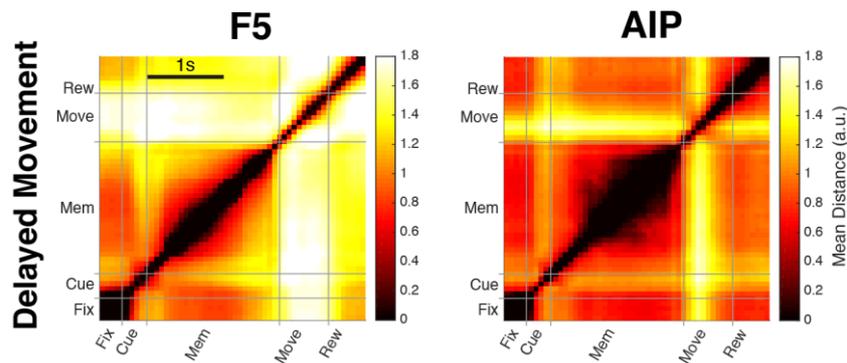
Figure 4 – Figure Supplement 1. Uniqueness of neural state over the course of the trial for an additional experiment. As described in Michaels et al. (2015), monkey Z performed a similar task to the current study (6 data sets x 2 grip types, Instructed condition). The same grip types were cued and the memory period was also variable; however, all trials resulted in movement, regardless of condition. Therefore, if the dynamic nature of the memory period observed in the present experiment were due only to the changing expectation of having to execute a movement over the course of the trial or the deterioration of a motor plan, we should observe stable activity. Yet, in this additional experiment the highly time dependent nature of the memory period activity is maintained, specifically in F5, suggesting that this variability is not due to the varying chance of subsequent movement, but represents an essential feature of the examined areas.

Signature of time

Given the constantly evolving population trajectory in F5, it should be possible to decode the exact time points of the task, while this should be more difficult in AIP, assuming the quality of recording is comparable between areas. To this end, we decoded the time course of the task, using a linear classifier (200 ms time bins, n-fold cross-validated; Matlab function: *fitcdiscr*). To ensure similar recording quality between areas, units were randomly discarded from each data set until there were an equivalent number of units in both areas (stratification; Materials and Methods). The average confusion matrix across data sets (decoded vs. real time bins) can be seen in Figure 5a for the no-movement condition and for a delayed grasp in Figure 5 – Figure Supplement 1a. as expected from Figure 4, the most decodable epochs were around movement and reward, while the least discernible were the first fixation period and the memory period. Figure 5b shows the average performance along the diagonal of the confusion matrices, revealing a significant advantage of F5 over AIP which was limited to the memory period (cluster-based permutation test across data sets, *p* <

0.05; Materials and Methods), confirming that there is less information regarding the passage of time in AIP. The fact that the reduction in decoding performance was limited to the memory period indicates that the observed difference between F5 and AIP represents a temporally localized difference in the dynamics at play while the monkey is remembering and withholding a grasping movement.

It is also notable that even when decoding was incorrect, most decoded epochs were assigned to those very close in time. Therefore, if we consider not just the diagonal of the confusion matrix when evaluating performance, but included also super- and sub-diagonal entries (i.e., entries just one step off the diagonal), performance increased by an average of 25 percentage points over both areas and delayed/no-movement trials (range: 23-26 percentage points). These results suggest that AIP holds a specific memory state during delays, while F5 actively tracks the temporal position within the delay, a novel finding in agreement with our third hypothesis.

Converging on movement

The population state at the time of go cue varied greatly between delays, especially in F5, but shortly before movement all activity converged towards a singular movement onset point. Taking a closer look at a few single units over all delay lengths (Figure 6a), we can see a large variety of paths to movement. However, all the trajectories converged at some point during the movement, with the exception of the first example in AIP, a point that will be returned to later.

To quantify how widespread variability around movement was, we tested for tuning to delay length independent of grip type (Figure 6b; sliding one-way ANOVA in steps of 50 ms, $p < 0.05$, Bonferroni-corrected for number of time points). Thirty percent of units showed delay length tuning and tuning in the population started to decrease about 200 ms before movement onset, reaching chance levels about 200-300 ms after movement onset in both areas. The progression of delay tuning was similar for both monkeys, although monkey B had more delay length tuning in F5 than monkey S. One possible explanation for the increased delay length tuning could be the slight, but significant, correlation between delay length and movement time only present in monkey B (Figure 1d and Figure 1 – Figure Supplement 1b).
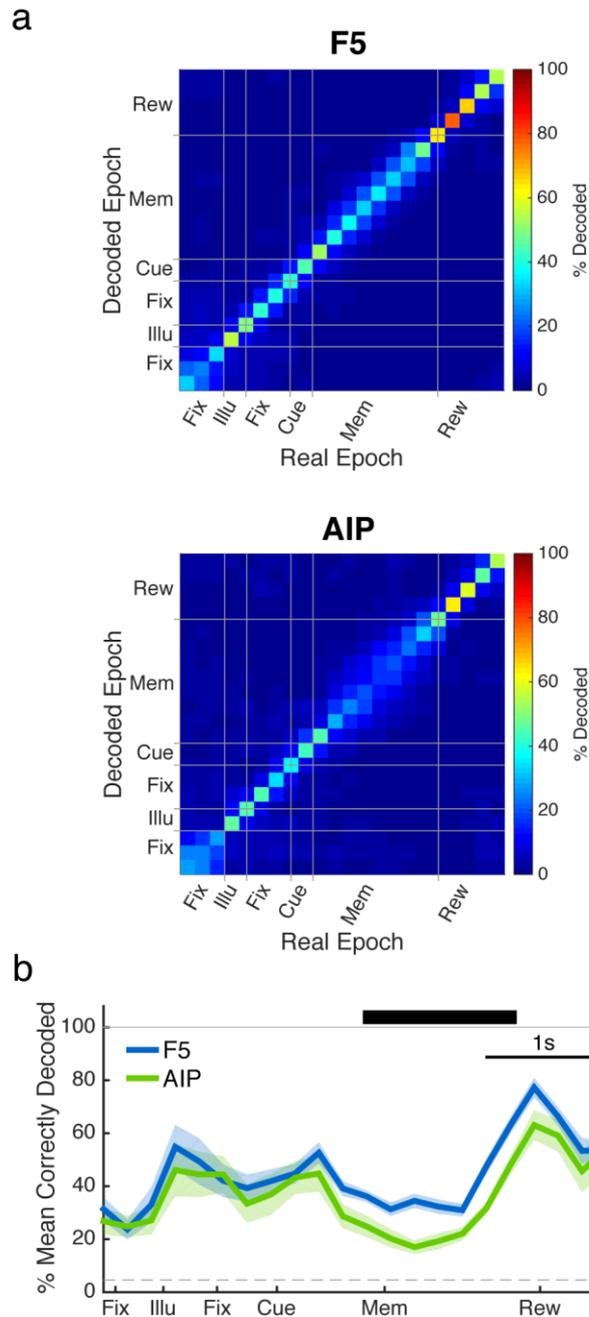
Figure 5. Time dependent decoding of task epochs of the no-movement condition using a linear classifier. (**a**) Confusion matrix of average decoding accuracy in identifying current epoch of each trial in non-overlapping 200 ms spike windows over all data sets in F5 (top) and AIP (bottom). (**b**) Mean correct decoding performance for each time epoch over all data sets in F5 and AIP. Black bar represents cluster of significant difference between F5 and AIP (cluster-based permutation test, $p < 0.05$) and the dashed line represents chance level. Units were discarded from each data set until an equivalent number of units were present for each area (stratification), then the decoding was repeated 100 times and the average performance taken. Abbreviations are as in Figure 4.
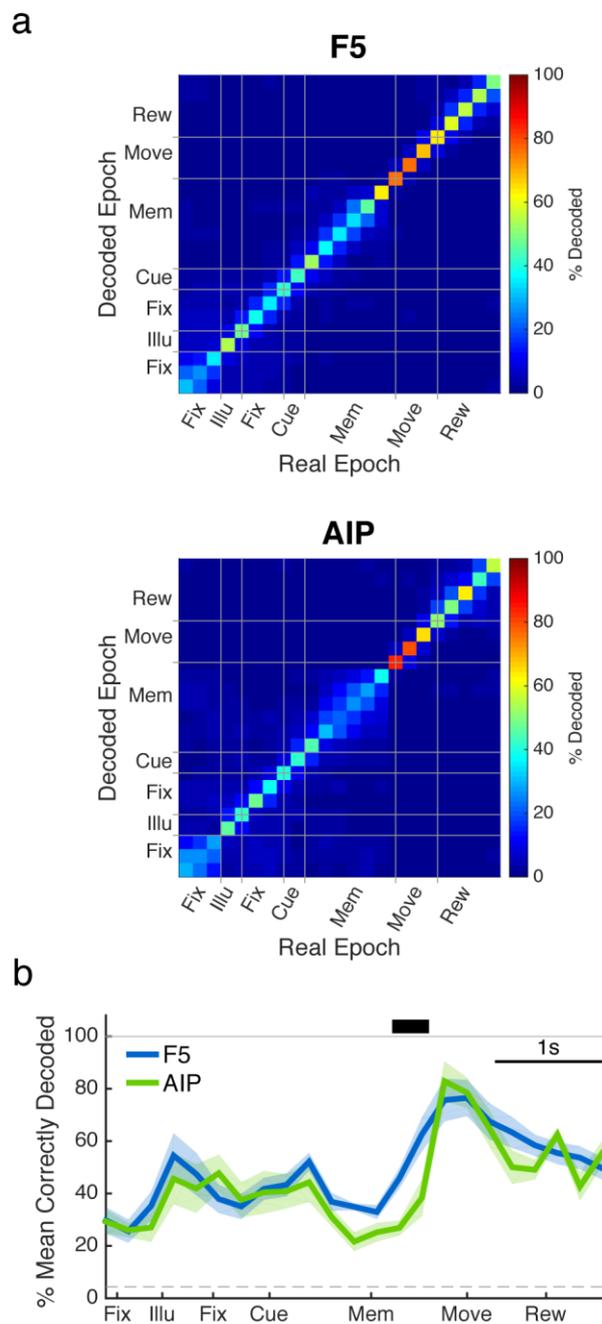
Figure 5 – Figure Supplement 1. Time dependent decoding of task epochs of delayed grasps (1300 ms) using a linear classifier.
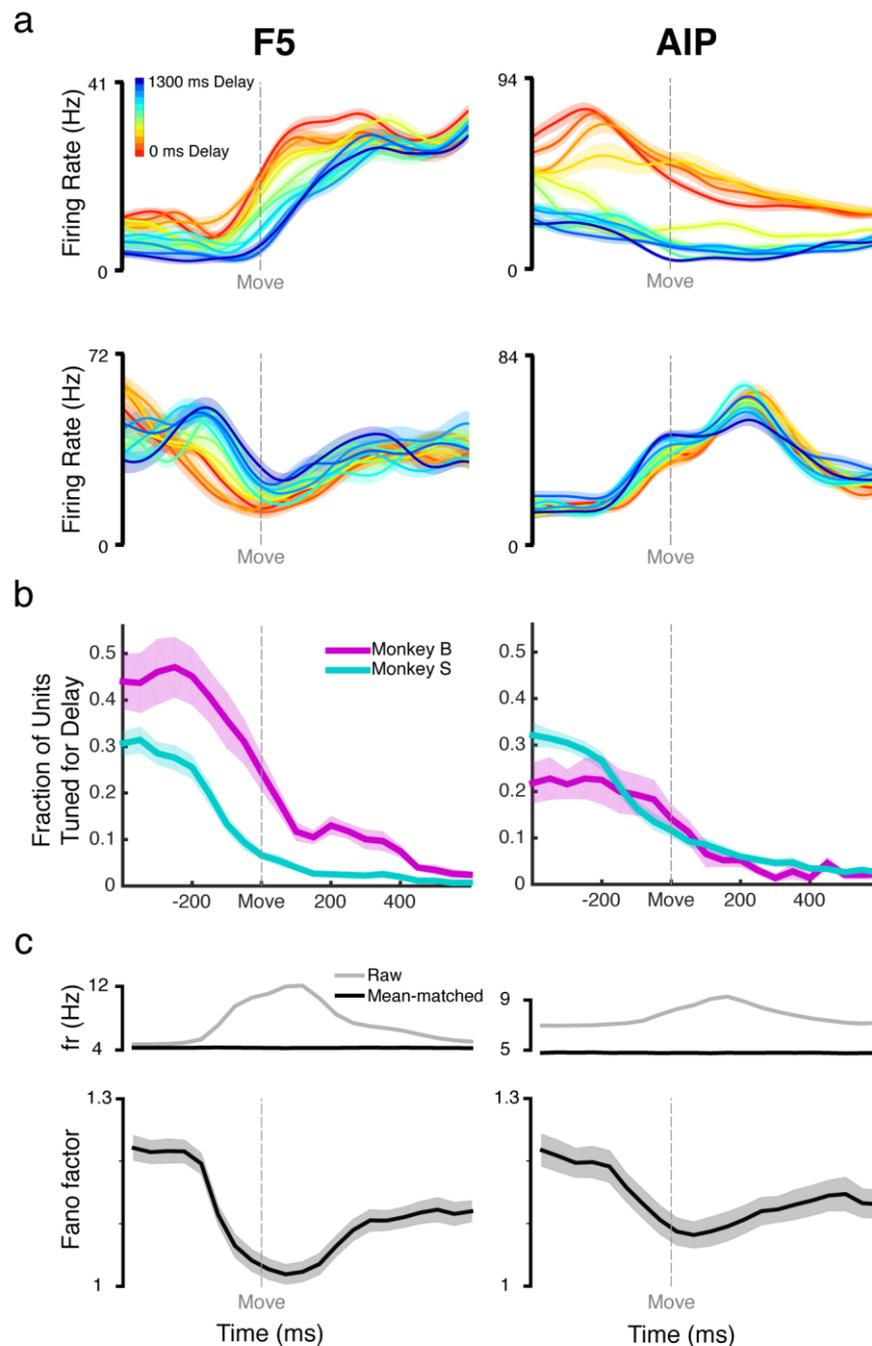
Figure 6. Rapid decrease in trial-to-trial variability during movement initiation. (**a**) Example average firing rates of single-units in F5 and AIP from both monkeys showing large firing rate differences between the various delay conditions of a single grip type. Only one grip type (precision or power) is plotted in each panel. (**b**) The fraction of units encoding delay length independent of grip type (one-way sliding ANOVA in steps of 50 ms, $p < 0.05$, Bonferroni-corrected for time points) in F5 and AIP. Error bars represent standard error of the mean over trials in (**a**), and over data sets in (**b**). (**c**) (Top) Mean firing rate before (gray) and after (black) mean-matching. (Bottom) Mean-matched Fano Factor over all units (pooled over data sets and conditions), showing a minimum shortly after movement onset. Error bars represent 95% confidence interval from regression.

At a specific time before movement onset, the high variability between delays rapidly begins to decrease. To quantify this effect further, we calculated the Fano factor over this same interval (Materials and Methods). The Fano factor provides a normalized measure of trial-to-trial spiking variability and has already been used to show that external stimuli decrease spiking variability in many cortical areas (Churchland et al., 2010a). Since the firing rate increases during movement (Figure 6c upper panel, gray line), which could possibly affect variability due to saturation of neurons at high firing rates, data were mean-matched (Figure 6c upper panel, black line) before calculating the Fano factor. Variability based on Fano factor is rapidly reduced 150-200 ms prior to movement onset (Figure 6c (bottom panel), reaching levels almost equivalent to the spontaneous spiking patterns of single neurons (neurons inherently do not spike in a completely predictable way, following a Poisson process). This strong effect is significant in both monkeys and areas ($p < 0.001$, confidence interval of regression; Materials and Methods) when comparing the Fano factor 300 ms before and 100 ms after movement onset, although the effect is stronger in F5. These results suggest that although the pre-movement activity for each grip is initially quite variable, it becomes mostly eliminated around movement onset, implicating an internal mechanism that brings all trajectories onto a similar path while the movement is being initiated. Although the path to movement onset may be broad, the state achieved when the monkey starts moving the hand seems to be in a very small area of the state space in order to successfully execute the grasp, confirming our fourth hypothesis and suggesting that activity must be very similar between trials in order to correctly execute the movement.
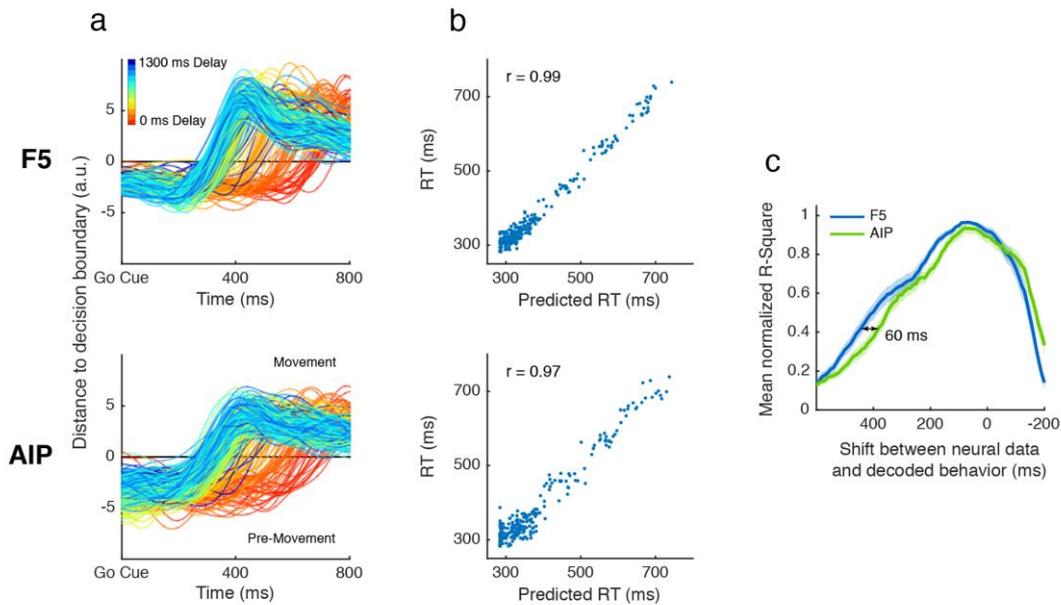
Figure 7. Movement initiation decoding in F5 and AIP. (**a**) Single trial distance to support-vector machine decision boundary between pre- and post-movement initiation generated from the population of units recorded simultaneously in F5 (top) and AIP (bottom) from an example data set (S2-Precision), aligned to the go cue. (**b**) Correlation of predicted time between go cue and movement onset (crossing of the decision boundary) and measured RT. Each point represents a single trial. (**c**) Difference in movement onset between F5 and AIP measured by systematically shifting the neural data relative to behavior and generating a new RT decoder for each time point (in steps of 10 ms), reported as mean normalized R-Square over all conditions (12 data sets x 2 grip types, normalized to max performance of each condition). F5 significantly led AIP by 60 ms (Wilcoxon sign-rank test, $p < 0.001$). Units were discarded from each data set until an equivalent number of units were present for each area (stratification), then the decoding was repeated 5 times and the average result taken.

Leading movement initiation

The decrease in variability preceding movement was not locked to the go cue itself, but to the internal generation of overt movement. This raised two questions: how is movement initiation represented in both F5 and AIP, and do both areas represent movement initiation at the same time? We can shed light on this process by testing the discriminability of movement initiation in each area and comparing the onset of this activity to see if one is leading the other, an analysis that is only possible when activity from many units in both areas is recorded in parallel. We implemented a support vector machine (SVM) to separate the times before and after movement onset by using activity in the 800 ms following the go cue (4-fold cross-validated; Matlab function: *fitcsvm*). Figure 7a shows the distance of each individual trial to the

discrimination boundary between pre- and post-movement onset of the SVM for an exemplar data set. Each trial crosses the decoder boundary at a specific time that can be used to make a prediction of RT. The time that each trial crossed the boundary (excluding unrealistic RTs: < 100 ms) was used to predict RTs and was correlated with the behavioral RT (Figure 7b), producing very high correlation coefficients. In order to make useful comparisons between the areas, we first discarded units from our recordings until the number of units in each brain area was the same (stratification), and repeated the current analysis 5 times. The average r-value over all data sets was 0.86 (0.95 and 0.77 for monkeys S and B, respectively; significant difference: Wilcoxon rank sum, $p < 0.001$) with no significant difference between areas or grip types (Wilcoxon rank sum, $p > 0.05$), suggesting that movement initiation is robustly encoded in both areas.

Since both areas were able to predict RT equally well, it is now possible to test if there was a difference in neuronal movement onset between the areas. We systematically shifted the neural data back in time relative to the behavior and tried to predict future RT using a new decoder for each time step (steps of 10 ms, 800 ms window; same SVM procedure as above). Figure 7c shows the mean performance of the decoders by time shift (normalized to maximum r-value for each session), showing that F5 significantly led AIP by approximately 60 ms in the representation of movement initiation (Wilcoxon sign-rank test, $p < 0.001$, 12 out of 12 data sets significant at the $p < 0.05$ level). Since our method relies on the decoding of RT on single trials, this small but robust difference in internal movement onset between the areas is only detectable when many unit are recorded in parallel, and suggests that the signal triggering movement to occur is more likely transmitted to F5 before AIP.

Clustering of immediate and withheld movements from memory
Based on the visualization of the many different delay trajectories between go cue and movement onset, the trajectories seemed to cluster into two groups dependent on the delay length. This interesting effect is especially visible in the example unit in the top right panel of Figure 6, where these two clusters remain present for over 600 ms after movement onset, signaling a long lasting effect. But is this clustering robustly present at the population level? To visualize the clustering for an example data set in F5, we plotted the activity of all linearly spaced delays (0-1000 ms) of a single grip type around movement onset in the largest principal component (Fig 8a, optimized to

extract delay variability; Materials and Methods). Looking specifically at the 500 and 600 ms delays just before movement onset, one can see that although both are following similar trajectories, the 500 ms delay is deflected downward and the 600 ms upward towards separate groups. This effect is even more striking in AIP, where delay lengths following the same trajectory diverge into two distinct groups (Figure 8 – Figure Supplement 1a).

To quantify this at the population level, we calculated the Euclidean distance between all pairs of delay lengths for each grip type in the full neural space (Figure 8b) and looked for clusters in the distance matrices without forcing clustering (Materials and Methods). Two clusters were identified for the example data set (Figure 8c), showing a split around the 500-600 ms delay point that lasts until under 50 ms before movement onset, after which no significant clusters are found (permutation test, $p < 0.05$, Bonferroni-corrected for number of time points; Materials and Methods). This pattern was very similar over all data sets and grips (Figure 8d) and was also present in AIP (Figure 8 – Figure Supplement 1), suggesting that the switch that occurs 500-600 ms after the cue spans both the frontal and parietal lobes.

Is it possible that the observed clustering could be due to factors other than a state switch? Clustering is not likely due to different movement kinematics, since the movement times were nearly identical for all delay lengths (Figure 1d), especially for monkey S. Additionally, since movement is initiated by the monkey, there could be a built-in bias in the amount of time that has elapsed since the grip cue was presented. However, differences in how long ago the grip cue was presented is unlikely to explain the two clusters, since repeating the same clustering analysis on the mean times between cue presentation and movement onset for all delays does not produce significant clustering for either grip type (permutation test, Precision grip: $p = 0.97$, Power grip: $p = 0.97$). These controls suggest that the separation of the neural trajectories into two distinct clusters reflects a robust effect of delay length in F5 and AIP, rather than some other trivial effect.
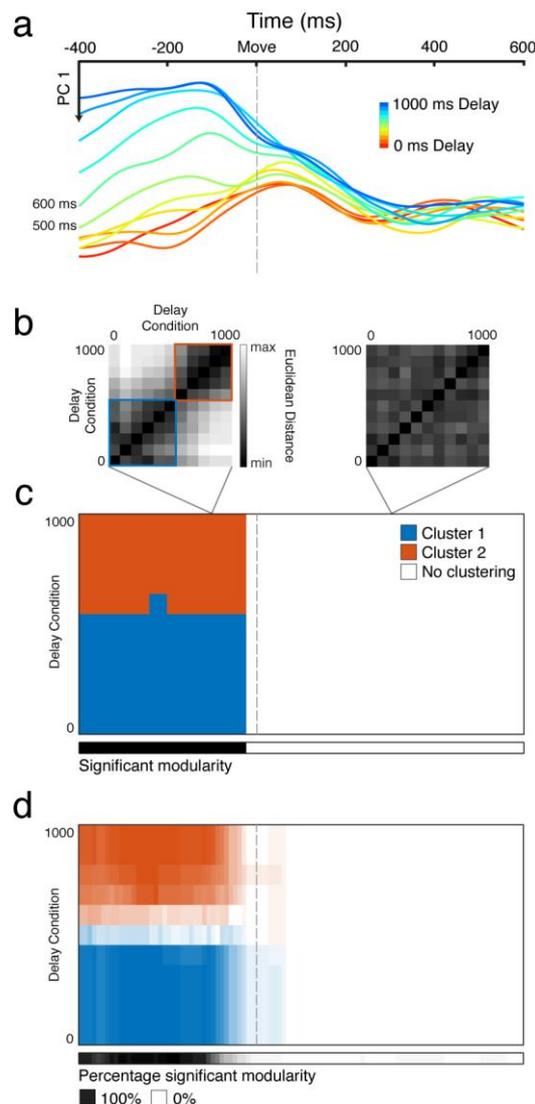
Figure 8. Clustering of movement initiation activity in F5. (**a**) PCA projection (optimized to distinguish delays) of population activity in F5 over all linearly spaced delays (0-1000 ms) for an example data set (B2-Precision), aligned to movement onset. (**b**) Euclidean distance between all pairs of delays in the full neural space for two example time points of the example data set including identified clustering using a clustering analysis that finds community structure. (**c**) Clusters identified in the distance matrices over time (in steps of 10 ms) for the example data set. Black significance bar shows time points where the modularity statistic exceeded chance level (permutation test, $p < 0.05$, Bonferroni-corrected for number of time points). (**d**) Same analysis as (**c**) averaged over all data sets and grip types of both monkeys (12 data sets x 2 grip types).
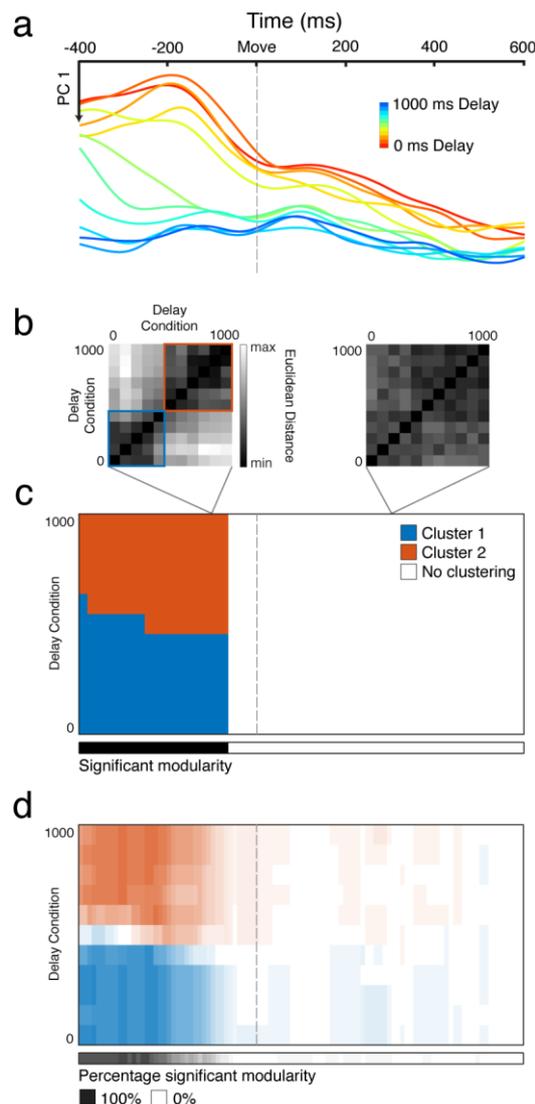
Figure 8 – Figure Supplement 1. Clustering of movement initiation activity in AIP. Same as Figure 8 for AIP (example data set in a-c: S3-Power).

Interestingly, as marked in the area of interest in Figure 8 – Figure Supplement 2, the delay at which the neural population transitions between the two clusters is not the same in F5 and AIP (Chi-Square test over all data sets, grips, and delay lengths, $p < 0.05$, Bonferroni-corrected for number of time points). For delays along the border between clusters in AIP, the 500 ms delay was more likely to be classified in the second cluster, suggesting that the state transition to the second cluster may occur earlier in AIP. This interesting finding suggests that although F5 leads AIP in movement initiation, AIP may lead F5 in the transition from immediate movements to withheld movements from memory.
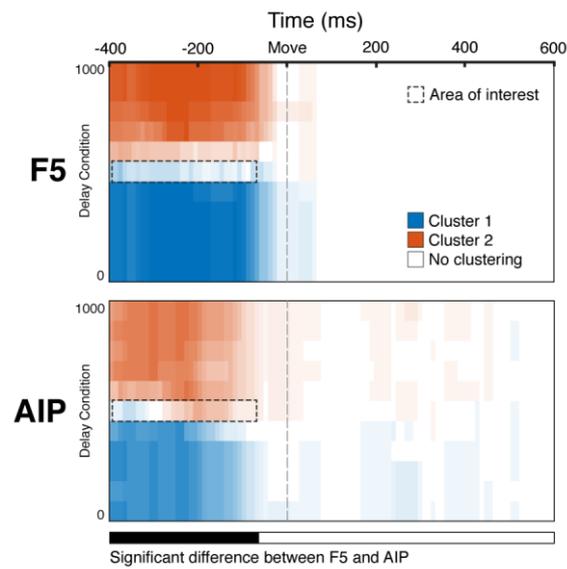
Figure 8 – Figure Supplement 2. Differences in cluster assignment between F5 and AIP. Presented data is identical to Figure 8d and Figure 8 – Figure Supplement 1d. Black bar marks time points that significantly differed in cluster assignment over all delays between F5 and AIP over all data sets and conditions (12 data sets x 2 grip types; Chi-Square test, $p < 0.05$, Bonferroni-corrected for number of time points). Note the difference in cluster assignment within the dashed area of interest.

Discussion

To systematically probe the interplay between planning and movement in the grasping network, we recorded large neural populations in premotor area F5 and parietal area AIP while two macaque monkeys performed a delayed grasping task with 12 distinct preparation times (0-1300 ms), from having to execute movements immediately to movements executed from memory after a long delay. Figure 9 illustrates the most significant findings using exemplar data sets. Firstly, the initial part of the neural space traversed was the same for all trajectories, regardless of delay length, and was grip specific, strongly suggesting that this activity was required to successfully prepare the correct movement. Secondly, once preparation was completed, population activity shifted into a separate state that is not achieved during short delays, revealing a unique memory state. Furthermore, while this area was static in AIP, it was highly dynamic in F5, tracking the evolution of time and suggesting a strong dissociation of roles between the two areas. Lastly, immediate movements followed a similar trajectory after the go cue, while withheld movements initiated from memory followed a separate trajectory to movement onset, forming two distinct clusters and suggesting a network-wide shift when movements are withheld.

Our results provide new insights for studies examining delayed reaching or grasping activity in premotor cortex (Cisek et al., 2003; Lucchetti et al., 2005; Fluet et al., 2010) or parietal cortex (Murata et al., 1996; Snyder et al., 2006; Baumann et al., 2009). The present study has some similarities with previous work on reaching movements (Ames et al., 2014). However, it provides several important advances. Firstly, we investigated the grasping network of the macaque, spanning two areas in multiple lobes of the brain, allowing comparison between the parietal and frontal aspects of grasp planning. Secondly, our paradigm is a memory-guided task, since our visual cue was phasic (300 ms) and was not presented either during the memory period or during the movement. This allowed us to dissociate strong visual inputs from internal planning and execution, unlike previous work that included a strong visual cue for the entire trial. Thirdly and crucially, previous work only compared a single delayed condition to a non-delayed condition, while we systematically varied the amount of preparation in 12 discrete steps, allowing fine resolution of the continuum between planning and movement.

Given that our task also involved a large reaching component, reach planning is likely a significant part of the observed activity. Still, the presence of robust grip

type tuning in our tasks, as well as previous research employing a grasp-only task (Hepp-Reymond et al., 1994), and a grasp-reach dissociation task (Lehmann and Scherberger, 2013) suggests that F5 encodes grasping quite independently of reaching, although both areas contain information about reach position (more so in AIP). Furthermore, reversibly inactivating F5 (Fogassi et al., 2001) or AIP (Gallese et al., 1994) selectively impairs hand-shaping and not reaching, suggesting that our results are an accurate representation of the grasping network.

The neural trajectory of short delays bypassed the states achieved by longer delays in both areas, taking shortcuts to movements. Yet, the first 300 ms of activity was identical regardless of delay, and grip specific (Figure 9), suggesting that this preparation is necessary for completing the movement. After this initial preparation, activity enters other states that are not necessary for preparation, but likely serve other purposes, such as holding a motor plan or withholding movement, a point discussed later. Furthermore, the trajectories of very short delays (0-200 ms) were almost completely identical for the entire trial (Figure 3c), suggesting that the planning and execution of movement does not differ within either F5 or AIP if the go cue comes before initial preparation is complete.

In F5 the memory period activity did not congregate in a specific region of the state space, a feature of the ventral premotor cortex never before observed to our knowledge. This finding is at odds with the conclusions of Ames et al. (2014), who postulated that delay period activity of the dorsal premotor cortex (PMd) may act as an attractor state into which all trials would congregate given enough preparation time. Nevertheless, since stability was not tested, it remains unclear if the attractor state described in their study is fixed or a temporally evolving state. An alternative explanation could be differences in task design. In Ames et al. (2014) the reaching target was present for the entire delay, a salient visual cue, while we utilized a memory period following a short presentation of the cue. Therefore, we predict that PMd activity would be much more dynamic if a memory period were used, a point supported by a study showing that PMd activity can encode prior knowledge of when events are likely to occur (Mauritz and Wise, 1986), a feature which requires continuously evolving population activity.
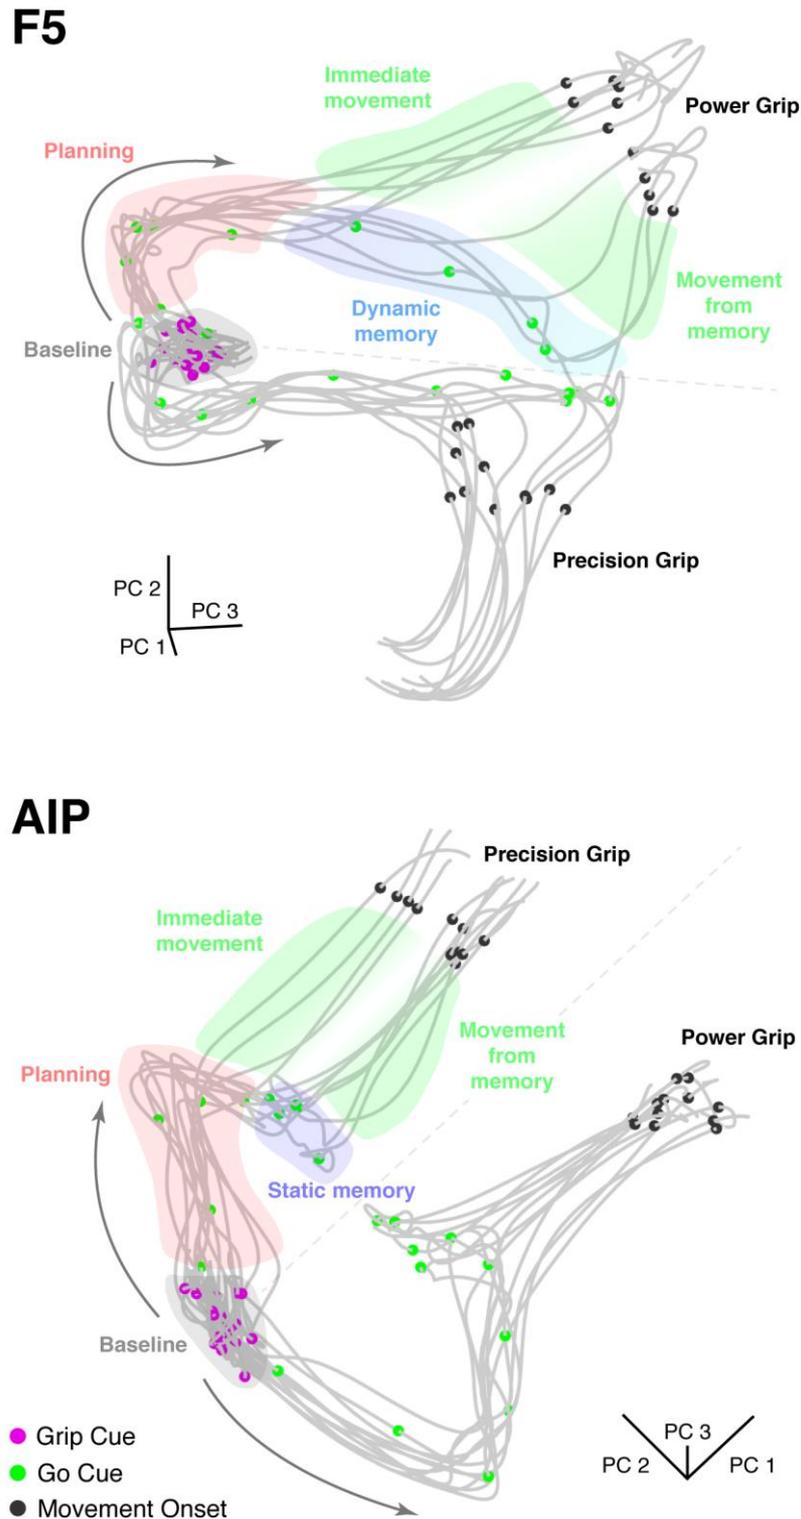
Figure 9. Visualization of fronto-parietal grasp preparation system. Visualization based on principal components of example data sets (F5: B4, AIP: B2). Each gray trace represents the mean activity for one delay length of one grip type. Following the grip cue activity in both areas follows a fixed preparation path, after which longer delays congregate in a memory state in AIP or continue to dynamically evolve in F5.

After the go cue, activity quickly branches off towards the appropriate grip onset position while adhering to two distinct clusters.

It could be that the temporal dynamics of F5 activity are a result of an internalized representation of the likelihood of task events occurring at specific times throughout the memory period, similar to hazard rate observed in the lateral intraparietal cortex (LIP) (Janssen and Shadlen, 2005) or the neural representation of time as observed in LIP (Leon and Shadlen, 2003). Nevertheless, in the current study we found significantly more time dependence in the memory period of F5 than AIP, even though LIP is connected to AIP (Borra et al., 2008). However, time dependence has also been identified in prefrontal areas (Genovesio et al., 2006), and in PMd (Lucchetti et al., 2005), suggesting that the temporal complexity we observe in F5 could be present in these nearby areas.

A mechanistic explanation for the dynamics observed during the memory period could be that recurrent networks of neurons in these areas generate temporal dynamics similar to a time code. Under this perspective, the complex temporal dynamics during memory would be a feature of a network optimized to produce the patterns of activity required to generate movement (Sussillo et al., 2015) and maintain the short-term memory of these motor plans in short time-scales. The observed temporal dynamics are in line with increasing literature suggesting that time keeping is an intrinsic property of all neural networks (for a review see Goel and Buonomano, 2014), as well as a feature of some sub-cortical areas (Gouvêa et al., 2015). Interestingly, a number of recent studies have shown that timing is a robust feature of chaotic recurrent networks (Buonomano and Laje, 2010; Laje et al., 2013; Goudar and Buonomano, 2014), and that neural networks can transition through internal states with or without the presence of external triggers (Ponce-Alvarez et al., 2012; Mazzucato et al., 2015), suggesting that F5 may be able to track the course of time internally and use this information to predict when an action is likely to be required.

One of the most striking features in both areas, but especially F5, was that the population activity of a single grip type was highly variable at the time of go cue, yet converged rapidly leading up to movement onset, raising the question of how the correct movement can be successfully initiated. Recently, alternative theories of movement generation have arisen, suggesting that preparatory activity in motor cortex may serve to set the initial conditions of a dynamical system (Churchland et al., 2012;

for a review see Shenoy et al., 2013; or Churchland and Cunningham, 2014). However, the large variability at go cue cannot be explained by a rotational dynamical system (Churchland et al., 2010b; 2012), since, under this model, all trials of a particular performed movement (e.g. power or precision grip) should have very similar preparatory activity and the movement activity should follow predictably from this state. This is especially interesting since F5 does indeed show strong rotational patterns during movement initiation (J.A. Michaels, S. Schaffelhofer, and H. Scherberger, unpublished data) that explain a similar amount of variance as in PMd/M1 (20-40%). We propose that the broadly tuned nature of activity at the go cue provides the motor system with a large flexibility in movement initiation. Similar to the dynamics observed during the memory period, it could be that once movement is triggered, recurrent networks of neurons within these areas rapidly reduce variability within particular regions of the neural space in order to ensure correct muscle activation during initiation. Under this framework, selecting between multiple movement plans would only require the neural population to be within a general region of activity. Once movement is initiated, activity would fall onto a common trajectory unique to each action plan and rotational dynamics could proceed as proposed previously (Churchland et al., 2012). Yet, the movement itself must be triggered in order to engage this process. Interestingly, the internal initiation of movement seems to appear first in F5, suggesting that the source of this trigger is either frontal or subcortical, and only transferred to AIP after a delay (Figure 7). Future work must tackle the question of to what degree local or extrinsic processes can account for the rapid decrease in trial-to-trial variability taking place before movement execution, bringing all trajectories onto a similar path.

Although variability is strongly suppressed leading up to movement onset, the trajectories did not follow individual paths to movement onset. Instead, trajectories seemed to cluster into two distinct groups, splitting around the 500-600 delay mark. Given that the preparation itself likely takes ~300 ms, in the case of short delays the arrival of the go cue appears before preparation has completed, thereby triggering a movement 'as fast as possible'. Conversely, in longer delays the monkey must first wait for the go signal. We propose that shifting between immediate movements and movements from memory causes a state shift in the fronto-parietal network that produces the two clusters during movement initiation (Figure 8-9). Once the state has been shifted, the trajectories continue to cluster for the entirety of movement

initiation, and occasionally remain visible into the movement itself in AIP (Figure 8 – Figure Supplement 1).

The clustering observed in both areas during movement initiation could be a function of the transition from reactive to proactive control, i.e., the increased ability to properly anticipate a go cue after long preparation times. This kind of neural state transition is inherent in highly trained tasks, as is the case in supplementary motor area (SMA)(Chen et al., 2010). Execution of timed behavior is reduced in humans with SMA lesions (Halsband et al., 1993) and supports our findings, since F5 is especially connected to the pre-SMA (Luppino et al., 1993). Yet, this explanation does not explain the strength of the clustering in the activity of AIP, or the interesting finding that AIP seemed to transition to the second state earlier than F5 (Figure 8 – Figure Supplement 2).

More and more evidence proposes that beta-band activity, normally measured in the local field potential (LFP) acts as a signal of maintenance and re-emerges during delayed tasks where movements must be withheld (for reviews see Engel and Fries, 2010; or Khanna and Carmena, 2015). A strong beta rebound occurs when movements must suddenly be withheld (Zhang et al., 2008). Therefore, it is possible that once movement planning is completed, beta could provide a network-wide signal of maintenance during proactive control that spans both F5 and AIP, explaining the presence of movement initiation clustering in both areas. Future studies should probe the boundary between immediate and withheld movements by assessing whether beta activity appears during some, but not all, amounts of preparation time, a finding that would implicate a state shift once movements are fully planned and await a go signal.

In the current study we probed the continuum of grasp preparation by systematically varying the amount of time monkeys had to prepare grasping movements. We showed robust differences in population dynamics between areas F5 and AIP, revealing that F5 may play a larger role in tracking time or expectation during memory. Interestingly, population activity in both areas converged on movement initiation along two separate trajectories, comprised of movements 'as fast as possible' and withheld movements executed from memory. The separation of trajectories into two distinct clusters suggests that a widespread and long-lasting state shift occurs during the transition from immediate to memory-guided movement.

Materials and Methods

Basic procedures

Neural activity was recorded simultaneously from area F5 and area AIP in one male and one female rhesus macaque monkey (*Macaca mulatta,* monkeys B and S; body weight 11.2 and 9.7 kg, respectively). Animal care and experimental procedures were conducted in accordance with German and European law and were in agreement with the *Guidelines for the Care and Use of Mammals in Neuroscience and Behavioral Research* (National Research Council, 2003).

Basic experimental methods have been described previously (Michaels et al., 2015). We trained monkeys to perform a delayed grasping task. They were seated in a primate chair and trained to grasp a handle with the left (monkey B) or the right hand (monkey S) (Fig. 1a). A handle was placed in front of the monkey at chest level at a distance of ~26 cm and could be grasped either with a power grip (opposition of fingers and palm) or precision grip (opposition of index finger and thumb; Fig. 1b insets). Two clearly visible recessions on either side of the handle contained touch sensors that detected thumb and forefinger contact during precision grips, whereas power grips were detected using an infrared light barrier inside the handle aperture. The monkey was instructed which grip type to make by means of two colored LED-like light dots projected from a TFT screen (CTF846-A; Screen size: 8" digital; Resolution 800x600; Refresh rate: 75Hz) onto the center of the handle via a half mirror positioned between the monkey's eyes and the target. A mask preventing a direct view of the image was placed in front of the TFT screen and two spotlights placed on either side could illuminate the handle. Apart from these light sources, the experimental room was completely dark. In addition, one or two capacitive touch sensors (Model EC3016NPAPL; Carlo Gavazzi) were placed at the level of the monkey's mid-torso and functioned as handrest buttons. The non-acting arm of monkey B was placed in a long tube, preventing it from interacting with the handle. Monkey S was trained to keep her non-acting hand on an additional handrest button.

Eye movements were measured using an infrared optical eye tracker (model AA-ETL-200; ISCAN) via a heat mirror directly in front of the monkey's head. To adjust the gain and offset, red calibration dots were shown at different locations at the beginning of each session for 25 trials that the monkey fixated for at least 2 seconds.

Eye tracking and the behavioral task were controlled by custom-written software implemented in LabView Realtime (National Instruments) with a time

resolution of 1 ms. An infrared camera was used to monitor behavior continuously throughout the entire experiment.

Task Design

The trial course of the delayed grasping task is shown in Figure 1b. Trials started after the monkey placed the acting hand on the resting position and fixated a red dot (fixation period). The monkey was required to keep the acting hand, or both hands (monkey S), completely still on the resting position until after the go cue. After a variable period of 400 to 700 ms two flashlights illuminated the handle for 300 ms, followed by 600 ms of additional fixation. In the cue period a second light dot was then shown next to the red one to instruct the monkey about the grip type for this trial (grip cue). Either a green or white dot appeared for 300 ms, indicating a power or a precision grip, respectively. After that, the monkey had to either react immediately or memorize the instruction for a variable memory period (also referred to as delay length). This memory period lasted for 0 to 1300 ms (i.e. the go cue could appear simultaneously with the grip cue, which was always presented for 300 ms regardless of the length of the delay), in discrete memory period bins of 0, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000, or 1300 ms. Switching off the fixation light then cued the monkey to reach and grasp the target (movement period) in order to receive a liquid reward. Monkeys were required to hold the appropriate grip for 300 ms. Additionally, no-movement trials were randomly interleaved (8% of trials), in which a go cue was never shown and the monkey only received a reward if it maintained fixation and the hands on the hand rests for 2000 ms following the grip cue. All trials were randomly interleaved and in total darkness.

Surgical procedures and imaging

Upon completion of behavioral training, each monkey received an MRI scan to locate anatomical landmarks, for subsequent chronic implantation of microelectrode arrays. Each monkey was sedated (e.g., 10 mg/kg ketamine and 0.5 mg/kg xylazine, i.m.) and placed in the scanner (GE Healthcare 1.5T or Siemens Trio 3T) in a prone position. T1-weighted volumetric images of the brain and skull were obtained as described previously (Baumann et al., 2009). We measured the stereotaxic location and depth

orientation of the arcuate and intra-parietal sulci to guide placement of the electrode arrays.

An initial surgery was performed to implant a head post (titanium cylinder; diameter, 18 mm). After recovery from this procedure and subsequent training of the task in the head-fixed condition, each monkey was implanted with floating microelectrode arrays (FMAs; MicroProbe for Life Science) in a separate procedure. Monkey B was implanted with six electrode arrays in the right hemisphere, each with 32 electrodes (Fig. 1e). Two such arrays were implanted in area F5, two in area AIP, and two in area M1. Monkey S was implanted with four FMAs in the left hemisphere and received two arrays in each area (Fig. 1f). The arcuate sulcus of monkey S did not present a spur, but in the MRI a small indentation was visible in the posterior bank of the arcuate sulcus, about 2 mm medial to the knee, which we treated as the spur. We placed both anterior FMAs lateral to that mark. FMAs consisted of non-moveable monopolar platinum-iridium electrodes with initial impedances ranging between 300 and 600 kΩ at 1 kHz measured before implantation and verified in vivo. Lengths of electrodes were between 1.5 and 7.1 mm.

All surgical procedures were performed under sterile conditions and general anesthesia (e.g., induction with 10 mg/kg ketamine, i.m., and 0.05 mg/kg atropine, s.c., followed by intubation, 1–2% isofluorane, and analgesia with 0.01 mg/kg buprenorphene). Heart and respiration rate, electrocardiogram, oxygen saturation, and body temperature were monitored continuously and systemic antibiotics and analgesics were administered for several days after each surgery. To prevent brain swelling while the dura was open, the monkey was mildly hyperventilated (end-tidal CO_2, ~30 mmHg) and mannitol was kept at hand. Monkeys were allowed to recover fully (~2 weeks) before behavioral training or recording experiments commenced.

Neural recordings and spike sorting

Signals from the implanted arrays were amplified and digitally stored using a 128 channel recording system (Cerebus, Blackrock Microsystems; sampling rate 30 kS/s; 0.6-7500Hz hardware filter). Data were first filtered using a median filter (window-length: 3ms) and the result subtracted from the raw signal, corresponding to a nonlinear high-pass filter. Afterwards, the signal was low-pass filtered with a non-causal Butterworth filter (5000 Hz; 4th order). To eliminate movement noise (i.e., common component induced by reference and ground), PCA artifact cancellation was

applied for all electrodes of each array (as described in Musial et al., 2002). In order to ensure that no individual channels were eliminated, PCA dimensions with any coefficient greater than 0.36 (with respect to normalized data) were retained. Spike waveforms were extracted and semi-automatically sorted using a modified version of the offline spike sorter *Wave_clus* (Quiroga et al., 2004; Kraskov et al., 2009).

Units were classified as single- or non-single unit, based on five criteria: (1) the absence of short (1–2 ms) intervals in the inter-spike interval histogram for single units, (2) the homogeneity and SD of the detected spike waveforms, (3) the separation of waveform clusters in the projection of the first 17 features (a combination for optimal discriminability of principal components, single values of the wavelet decomposition, and samples of spike waveforms) detected by *Wave_clus*, (4) the presence of well known waveform shapes characteristics for single units, and (5) the shape of the inter-spike interval distribution.

After the semiautomatic sorting process, redetection of the average waveforms (templates) was done in order to detect overlaid waveforms (Gozani and Miller, 1994). Filtered signals were convolved with the templates starting with the biggest waveform. Independently for each template, redetection and resorting was run automatically using a linear classifier function (Matlab function: *classify*). After the identification of the target template, the shift-corrected template (achieved by up and down sampling) was subtracted from the filtered signal of the corresponding channel to reduce artifacts for detection of the next template. This procedure allowed a detection of templates up to an overlap of 0.2 ms. As a control, unit isolation was evaluated again as described before to determine the final classification of all units into single- or multi-units. Units were only classified as single if they unambiguously met the five criteria. Stationarity of firing rate was checked for all units and in case the firing rate was not stable over the entire recording period (more than 30% change in firing rate between the first 10 min and the last 10 min of recording), the unit was excluded from further analyses (<3% of all units).

Data preprocessing
Although units were classified as single- or multi-units, all recorded units were used in our main analyses. A detailed list of data set information can be found in Table 1. After spike sorting, spike events were binned in non-overlapping 1 ms windows and smoothed with a Gaussian window ($\sigma = 50$ ms) to produce a continuous firing rate

signal (1 kHz). Data were aligned to two events, the presentation of the grip cue and movement onset, i.e. the time when the monkey's hand left the handrest. No go cue alignment was necessary, since the length of the delay period was identical for all trials of the same delay. For most analyses these two alignments were combined to produce a continuous signal with as small an overlap as possible. In this case the two signals were simply concatenated in time. Average firing rates were then calculated by averaging over all trials of the same condition. All units received soft-normalization before further population analysis, which is computed by dividing the activity of each unit over time, $f(t)$, by it's overall firing rate range, f_r, plus a softening term, s, giving $f'(t) = f(t)/(f_r + s)$. For all analyses we used a softening term of 10 Hz, meaning a unit that fires between 0 and 10 Hz would correspond to a normalized range of 0 to 0.5. In contrast, a softening term of 0 would correspond to normalization by range. Normalization was utilized in order scale the impact of extremely high firing units while not over-emphasizing weak units.

Dimensionality reduction

Dimensionality reduction was carried out for the purpose of visualization only. All statistical analyses relied on the full space of all units. Neural trajectories were generated by performing standard principal component analysis (Matlab function: *princomp*) on normalized firing rate curves. In general, data was prepared by creating a matrix of size ct x n, where c is the number of conditions analyzed (e.g., delayed and non-delayed grasps / precision and power grips), t is the number of time points, and n is the number of units recorded, then finding principal components as a linear combination of the n units. The only exception to this procedure was the PCA analysis in Figure 8a, where data were first averaged over all time points, t, in order to find components that ideally separated the various delays, then the entire trajectory was transformed into this space.

Distance analysis

In order to find the neural distance between two conditions over time, we calculated the minimum Euclidean distance (point-to-curve distance) between the two trajectories in the full space of all units. Two versions of this analysis were performed. For the distance in Figure 3, we selected time points on specific delayed trajectories

(in steps of 50 ms) and calculated the Euclidean point-to-curve distance of the delayed (100-1300 ms) trajectory to the non-delayed (0 ms) trajectory, calculated as the minimum of all distances between a specific point on the delayed trajectory to all points on the non-delayed trajectory. Minimum distance, as a conservative measure, was used in order to overcome the different time courses of the conditions being compared. Small distances indicate that the two trajectories achieve the same point in neural space at some point in time, while large distances indicate that the two trajectories do not pass through the same point in the high dimensional space.

For the distance analysis in Figure 4, the Euclidean distance was calculated between all pairs of time points on the same trajectory and used in conjunction with the bootstrapping procedure (next section) to determine if two points significantly differed.

Bootstrap procedure

In order to gain an estimate of underlying trial-to-trial variability, we performed a bootstrap analysis. This procedure varied slightly for the different distance analyses presented above, but the general principle was constant. We resampled trials from each condition randomly, with replacement, of the same size as the number of recorded trials in that condition. We then constructed normalized average firing rates and carried out the appropriate distance analysis as described above (e.g., minimum distance between delayed and non-delayed trajectory). This resampling was done 1000 times, producing a distribution of distances.

To obtain an estimate of how much distance is expected between trajectories by chance, we carried out another resampling in which a trajectory was resampled from itself to determine its underlying variability. Trajectories were resampled once with the number of trials observed in that condition, and once using the number of trials recorded in the other trajectory in the comparison, then the Euclidean distance was calculated as described in the previous section. We could then calculate the percentage of observations where the distance between delayed and non-delayed trajectories was greater than the distance from a trajectory to itself (1000x1000 comparisons) and use this as a *p*-value. This analysis allowed us to determine when an observed distance was significantly greater than the distance expected if two trajectories were generated from the same underlying distribution.

For chance analyses in Figure 4, resampling of trials was carried out 6000 times, with replacement, for each condition and data set, due to the large number of comparisons being made. For each of the 6000 resampling steps the same trajectory was resampled twice, termed $\boldsymbol{p}$ and $\boldsymbol{p}'$. Then, for every pair of time points (t_1 and t_2), the resampled distance along the first trajectory $d = d\big(\boldsymbol{p}(t_1), \boldsymbol{p}(t_2)\big)$ was compared to the two inter-trajectory distances at time t_1 and t_2: $d_1 = d\big(\boldsymbol{p}(t_1), \boldsymbol{p}'(t_1)\big)$ and $d_2 = d\big(\boldsymbol{p}(t_2), \boldsymbol{p}'(t_2)\big)$. We determined the fraction of resamples (across all 6000) for which the along-trajectory distance d exceeded both inter-trajectory distances: $d > max(d_1, d_2)$. This fraction determined a specific p-value for each time pair (t_1, t_2). The resampled distance, d, was then considered significant if $p < 0.01/c$, given the Bonferroni correction $c = T^2$, where T is the number of time points on the trajectory. In this way, the underlying threshold for significance was dependent on which time points were compared along the trajectory, establishing a conservative estimate of the underlying trial-to-trial variability.

Linear classifier

To evaluate the uniqueness of different time points in the trial, a linear classifier (linear discriminant analysis) was used to discriminate non-overlapping 200 ms bins of spikes distributed throughout the trial-course (Matlab function: *fitcdiscr*). In order to make a meaningful comparison between F5 and AIP, units were first randomly discarded from each recording session until both areas had an equivalent number of units (stratification). For classification, decoding was n-fold cross-validated, where n is the total number of trials. All trials of the delay length being investigated (either no-movement or 1300 ms delay) were decoded together regardless of grip type.

To determine at which times during the trial F5 and AIP differed, decoding performance was compared using a cluster-based permutation test (Maris and Oostenveld, 2007). Briefly, this test evaluates the t-statistic (independent samples) between two conditions over all time points and extracts clusters (consecutive time segments) of activity whose t-statistic exceeds a predefined threshold (α = 0.05), then the absolute t-statistics within each cluster were summed to produce cluster-level statistics. To generate a chance-level distribution from which the appropriate threshold could be determined, trials were randomly partitioned between the two conditions and the t-test and clustering redone (1000 partitions). From every partition

the largest cluster-level statistic was used to generate a largest chance cluster distribution. By comparing the real cluster-level statistic against the largest chance cluster distribution, significant clusters could be determined if the observed cluster value exceeded a set percentage of largest chance cluster values ($p = 0.05$). In this way, sensitivity to short time-scale differences is greatly reduced, but the overall false-alarm rate remains below the designated p-value.

Fano factor

In order to obtain a measure of how spike rate variability changes over time, we employed the frequently used measure of Fano factor. The current analysis was performed using a freely available toolbox (http://churchlandlab.neuroscience.columbia.edu/code/) that was originally introduced by Churchland et al. (2010a). Briefly, Fano factor is based on the ratio of spiking variance (across trials) to spiking mean rate. The total data set consisted of all units (pooled over recording sessions), and both grip types. Data for each unit and grip type were initially treated separately. Spike counts were computed in a 100 ms sliding window in steps of 50 ms from 400 ms before movement onset to 600 ms after.

For each time point, the variance across trials was plotted against the mean spike count (one point per unit x grip type). The weighted regression was calculated through these points. For the regression, values were weighted by the estimated sampling error of the variance, which is the square of the mean divided by the number of trials, and the resulting slope of the regression represented the raw Fano factor. A value of one indicates purely Poisson spiking.

In order to control for increases in firing rate over time, which could bias spike timing, data were first mean-matched. The mean-matching procedure consisted of calculating the histogram of mean rates over all units and grip types for each time point, then finding the largest common distribution over all time points, i.e., the height of each bin in the common distribution was equal to the smallest height of that bin over all time points. Afterwards, data points were randomly discarded from each bin until the distribution at each time point matched the common distribution. This procedure was carried out 50 times and the resulting Fano factors averaged to produce the mean-matched Fano factor. During mean-matching, 21% of data points were discarded in F5 and 15% in AIP. This procedure ensures that the overall mean does

not increase over time, thereby eliminating any reduction in Fano factor that is purely a result of an increase in the mean.

To evaluate if the reduction in Fano factor was significant, the sampling distributions estimated from the 95% confidence intervals provided by the regression were compared between 300 ms before movement onset and 100 ms after movement onset to produce a *p*-value.

Clustering analysis

To evaluate whether or not delay trajectories leading up to movement onset clustered in a distinct way, we calculated the Euclidean distance between all pairs of linearly spaced delays (0-1000 ms, in steps of 10 ms) in the full neural space and looked for community structure (i.e. distinct clusters of similar value) in the resulting distance matrix. We employed a well-known modularity analysis that iteratively finds non-overlapping groups of conditions that minimizes the within-group distance between conditions and maximizes the between-group distance (Newman, 2004; Reichardt and Bornholdt, 2006) with a gamma sensitivity of 0.8. Using this analysis, the number of clusters obtained is purely data-driven and not specified by the experimenter. To ensure that the found structure was not due to chance, we randomly permuted the distance matrix (1000 permutations, while conserving matrix symmetry) and compared the modularity statistic, Q, between the empirical and permuted data. The percentage of instances where the empirical value exceeded the permuted distribution was used to generate a *p*-value.

2.3 – Neural population dynamics during reaching are better explained by a dynamical system than representational tuning

Authors:

Jonathan A. Michaels
Benjamin Dann
Hansjörg Scherberger

"Even something as mundane as watching a movie on a TV screen is an example of the importance of emergent properties: one cannot comprehend the scene by looking at individual pixels but instead needs to simultaneously view many pixels to decipher the image. Although the neuron doctrine and single neuronal techniques have focused on the exhaustive analysis of the individual 'pixels' of the brain, it is possible that the function of neural circuits may not be apparent unless one can visualize many, or most, 'pixels' in the screen."
(Yuste, 2015)

PLOS | COMPUTATIONAL BIOLOGY

RESEARCH ARTICLE

Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning

Jonathan A. Michaels[1], Benjamin Dann[1], Hansjörg Scherberger[1,2]*

1 German Primate Center, Göttingen, Germany, 2 Faculty of Biology, Georg-August-Universität Göttingen, Göttingen, Germany

* hscherberger@dpz.eu

Abstract

Recent models of movement generation in motor cortex have sought to explain neural activity not as a function of movement parameters, known as representational models, but as a dynamical system acting at the level of the population. Despite evidence supporting this framework, the evaluation of representational models and their integration with dynamical systems is incomplete in the literature. Using a representational velocity-tuning based simulation of center-out reaching, we show that incorporating variable latency offsets between neural activity and kinematics is sufficient to generate rotational dynamics at the level of neural populations, a phenomenon observed in motor cortex. However, we developed a covariance-matched permutation test (CMPT) that reassigns neural data between task conditions independently for each neuron while maintaining overall neuron-to-neuron relationships, revealing that rotations based on the representational model did not uniquely depend on the underlying condition structure. In contrast, rotations based on either a dynamical model or motor cortex data depend on this relationship, providing evidence that the dynamical model more readily explains motor cortex activity. Importantly, implementing a recurrent neural network we demonstrate that both representational tuning properties and rotational dynamics emerge, providing evidence that a dynamical system can reproduce previous findings of representational tuning. Finally, using motor cortex data in combination with the CMPT, we show that results based on small numbers of neurons or conditions should be interpreted cautiously, potentially informing future experimental design. Together, our findings reinforce the view that representational models lack the explanatory power to describe complex aspects of single neuron and population level activity.

100

Author Summary

The question of how the brain generates movement has been extensively studied, yet multiple competing models exist. Representational approaches relate the activity of single neurons to movement parameters such as velocity and position, approaches useful for the decoding of movement intentions, while the dynamical systems approach predicts that neural activity should evolve in a predictable way based on population activity. Existing representational models cannot reproduce the recent finding in monkeys that predictable rotational patterns underlie motor cortex activity during reach initiation, a finding predicted by a dynamical model in which muscle activity is a direct combination of neural population rotations. However, previous simulations did not consider an essential aspect of representational models: variable time offsets between neurons and kinematics. Whereas these offsets reveal rotational patterns in the model, these rotations are statistically different from those observed in the brain and predicted by a dynamical model. Importantly, a simple recurrent neural network model also showed rotational patterns statistically similar to those observed in the brain, supporting the idea that dynamical systems-based approaches may provide a powerful explanation of motor cortex function.

Introduction

Throughout the history of neuroscience research, the question of how motor cortex generates movements has been investigated deeply [1]. Yet, substantial and conflicting models have been proposed [2–7]. According to the representational view, motor cortex neurons encode abstract or high-level aspects of movements, such as kinematic parameters [8]. In contrast, in the dynamical systems view the firing of each neuron is a function of a population optimized to control muscles directly [9]. It remains a point of considerable debate which model better explains existing neural data and provides a mechanistic explanation of how movements can be generated.

The representational view of neuron tuning, or 'neuron doctrine', is strongly rooted in the history of neuroscience [10] and detailed models of single neuron tuning have been indispensable tools for a basic understanding of the brain's computations [11–13]. However, recent advances in electrophysiological recording technology [14,15] have made it possible to examine network level hypotheses of movement generation that require large populations of neurons to study [16–19].

Recently, it was suggested that motor cortex, operating as a dynamical system, might be sufficient for generating required muscle activity [20–22]. Using simultaneous recordings in the dorsal premotor cortex (PMd) and primary motor cortex (M1) of non-human primates, Churchland et al. [22] proposed that preparatory activity may act to prepare a dynamical system, which, like a spring box, could be released to act as an 'engine of movement' and produce muscle activity from a basis set of oscillators, which they termed the generator model or dynamical model [9,23]. They supported their theory by developing a dimensionality reduction method (jPCA), which revealed that predictable rotational dynamics underlie a large portion of the variance observed in PMd/M1 during reach initiation, a direct prediction of the dynamical model. Importantly, they showed that representational models of movement activity, including those based on velocity tuning in single neurons [24] and complex kinematic models [25], did not contain the robust rotational patterns they observed empirically, and therefore are weak descriptive models [23].

However, it has been shown that when fitting neural activity to kinematic variables, decoding of movement intention can be improved by including variable time lags between single neuron activity and kinematics (neuron-kinematic latency, [24,26–28]) and these offsets are highly variable (SD: 70 ms; re-digitized data, Moran & Schwartz [24], their Fig 13A). Yet, these offsets were not included in the comparison to representational models made by Churchland et al. [22]. Furthermore, given that representational models of single neuron tuning have been widely implemented in both an experimental and clinical setting, such as in the development of neural prosthetics, it is not clear how those results can be interpreted under the dynamical systems framework.

To clarify this, we first investigated whether or not jPCA would reveal rotational dynamics in a velocity-based model for center-out reaching in which neuron-kinematic latencies were built into single neuron activity. We found that jPCA alone revealed rotational structure in both the representational model and the dynamical model, but that implementing a novel covariance-matched permutation test (CMPT) readily distinguished between these two, showing that variable neuron-kinematic latency did not uniquely produce rotational structure due to the condition structure. Secondly, we show that movement intention could be decoded from a recurrent neural network (RNN) trained to complete the same task using representational methods, such as the population vector, even though the preferred directions of single neurons were highly unstable, suggesting that high levels of decoding performance using representational models do not necessarily inform the mechanistic operation of the underlying circuit. Importantly, both simulated RNN data and real data collected in PMd/M1 of macaque monkeys show similar and significant rotations under the CMPT, providing further support for the dynamical systems view. Furthermore, repeating the CMPT on subsets of the PMd/M1 data showed clear minima in number of neurons and conditions required to draw statistical conclusions, cautioning the use of such analysis methods on low numbers of conditions or neurons, and thus informing the design of future experiments.

Results / Discussion

Incorporating variable neuron-kinematic latencies into the representational model

Velocity-based models without variable neuron-kinematic latencies were shown to exhibit little to no rotational structure [22]. To investigate how variable neuron-kinematic latencies may affect rotational structure, we simulated 200 cosine-tuned motor cortex neurons in a standard 13-direction center-out reaching task with variable neuron-kinematic latencies (Fig 1A; Methods) [13]. The simulation was based on the assumption of bell-shaped velocity profiles (Fig 1B). For activity with a movement duration of 300–400 ms and a latency distribution with a standard deviation (SD) of 72 ms, we found that the first principal component (PC) of our population of simulated neurons resembled a condition-independent representation of the individual neuron profile, while the second PC resembled a condition-dependent representation (Fig 1C). Interestingly, all higher order PCs resembled a sequence of harmonic Fourier bases. In general, it is well known that time-shifted versions of identical signals preferentially produce PCs very similar to a Fourier series (S1A and S1B Fig) as a result of sinusoidal eigenvectors of increasing frequency. This feature introduces a potential confound, since the higher-order PCs show patterns of activity that are not present in any individual neurons. Furthermore, these PCs produce rotational 'horseshoe' patterns when plotted in a plane (S1C Fig) [29], revealing how rotations can emerge from signals that are not present in any individual neuron (for an example of false interpretations made from application of PCA, see this well-known example from genetics research [30,31]).
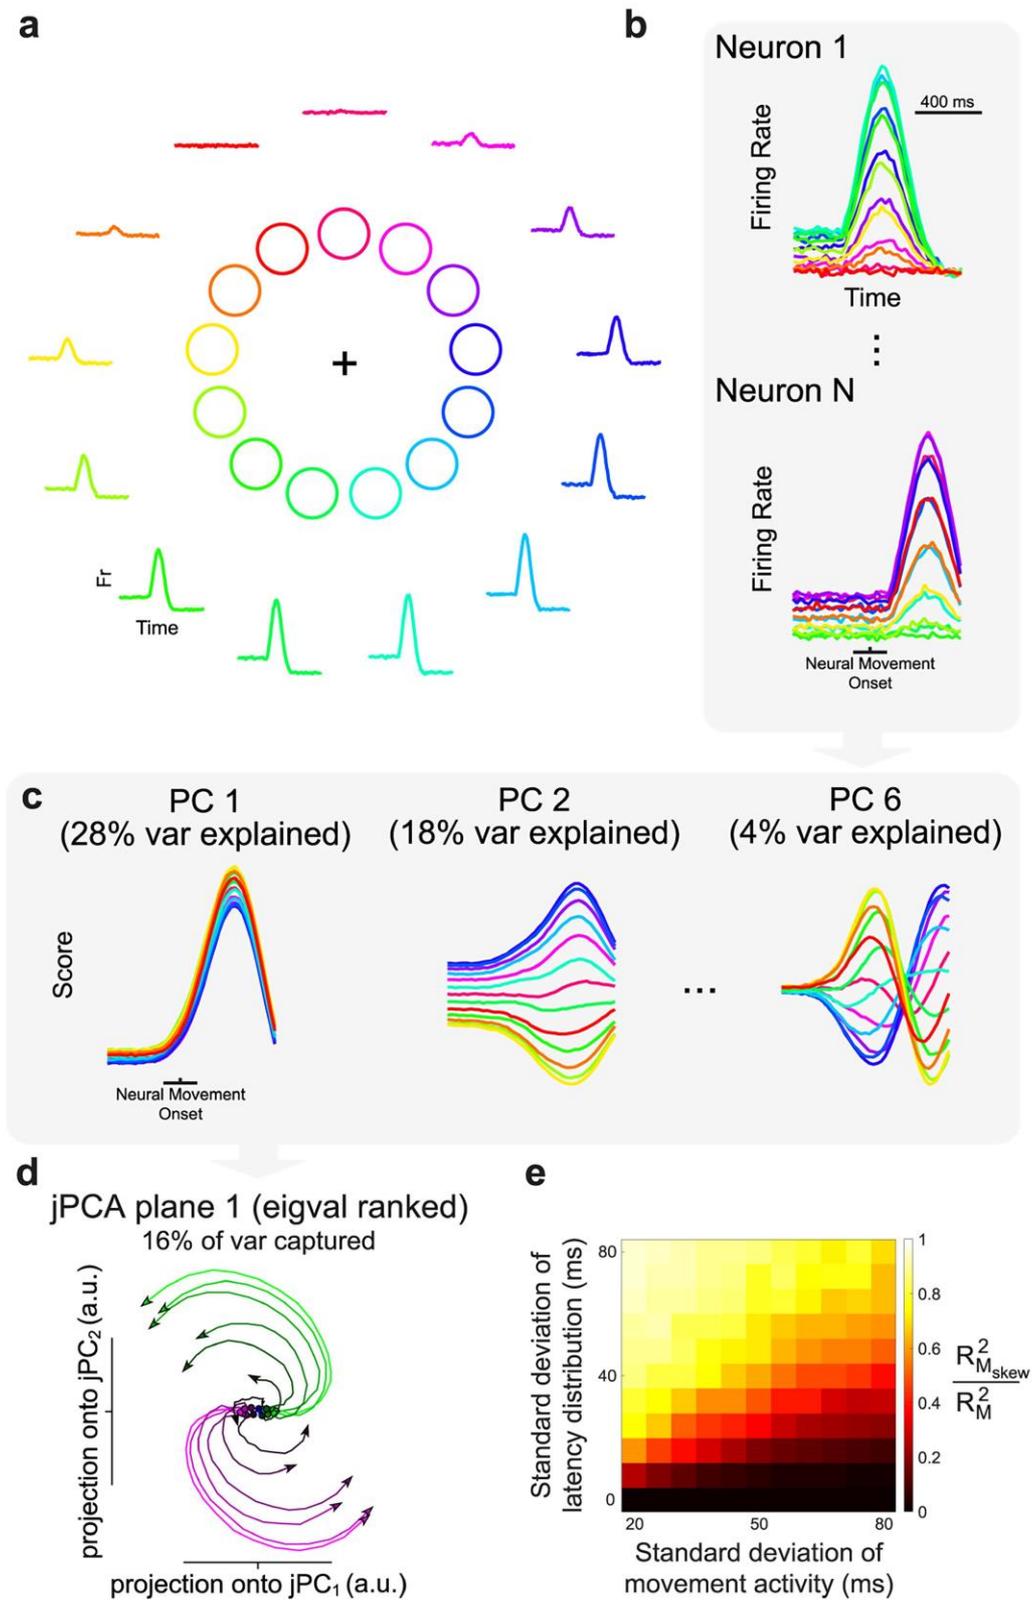
Fig 1. Simulation of a velocity-tuning based model with variable neuron-kinematic latencies. (a) Task design of a 13-direction center-out reaching task. The firing of a simulated neuron is plotted around the reach directions. **(b)** Two example neurons with differing latencies. **(c)** Principal components (PCs) for a simulated population of 200 neurons (latency SD: 72 ms, movement SD: 56 ms). **(d)** Exemplar jPCA plane for the first 6 PCs of the simulated population from 0 ms before to 200 ms after neural movement onset (analysis was computed on entire movement). Individual conditions are colored based on their activity at neural movement onset in the first jPC. **(e)** Proportion of change in neural trajectory explained by rotational dynamics (in all jPCA planes) for various latency offsets and movement durations. A value of 1 indicates that rotational dynamics completely explain the transformation between each time point and its temporal derivative.

In order to test the presence of rotational structure at the population level, we implemented the same analysis developed by Churchland et al. [22], termed jPCA (Methods). jPCA is a method for finding linear combinations of principal components that capture rotational structure in a population of neurons. In essence, jPCA finds low-dimensional planes in which neural activity follows a predictable rotational trajectory from time point to time point (analogous to a circular flow-field). We found that the introduction of the above-mentioned variable neuron-kinematic latencies were sufficient to produce rotational dynamics (Fig 1D) when explored with jPCA, unlike the representational model results of Churchland et al. [22], who found only weak rotations. The level of rotational dynamics observed here is similar to empirically recorded PMd/M1 data in terms of visualization of the jPCA planes, amount of variance explained per plane (30% in the first two planes, 16% in the first plane), rotational goodness-of-fit ratio (RGR) between $R^2_{M_{skew}}$ and R^2_M (0.79 in the first three planes; Methods), which provides a measure of how much variance can be explained by purely rotational dynamics, and how circular the rotation (0.72, where 1 is purely circular, computed as the average dot product of angle between x and $\dot{x}$, and $\pi/2$; Methods).

To characterize more generally how rotational structure arises with the addition of variable lags, we varied the duration of movement period activity (expressed as the SD of normally distributed movement activity; Methods) and the SD of the latency distribution systematically in repeated simulations (Fig 1E). Interestingly, when the SD of the latency distribution exceeded the SD of the movement activity, the level of underlying rotational structure increased rapidly. Therefore, our results show that the application of jPCA alone on a population where neuron-kinematic latency is more variable than the duration of movement leads to rotational dynamics.

Disrupting the underlying condition structure–covariance-matched permutation test

Based on the above results, it is clear that jPCA alone is not sufficient to distinguish between a representational model with lags and the dynamical model proposed by Churchland et al. [22]. While Churchland et al. [22] performed extensive shuffling controls to test the possibility that rotations emerge purely as a consequence of high-dimensional data, their controls do not differentiate between the above cases. Therefore, we developed a covariance-matched permutation test (CMPT) to differentiate these models. The objective of our test was to determine if the underlying condition structure, i.e., whether or not shuffling the neural data between different task conditions independently for each neuron, uniquely determined the rotational structure as is predicted by the dynamical model.

To provide intuition about the rationale of the test, consider the dynamical model proposed by Churchland et al. [22]. They observed that muscle activity during reaching could be fit extremely well (correlation coefficients ≥ 0.97) by a summation of two sinusoidal oscillators, each with fixed frequency, but whose phase, amplitude, and constant offset varied from

condition to condition (Methods). They proposed that these oscillators underlie the neural population activity during movement, providing a basis set from which the muscle activity can be generated, while the preparatory activity sets the phase and amplitude of these rotations. Since the phase and amplitude of these rotations are unique to each condition and defined jointly across the entire neural population, disrupting the condition structure should eliminate rotational structure. In Fig 2A we show one of two example oscillators (2.8 Hz), which consisted of a pair of leading and lagging sinusoids. To simulate neurons in the model, we randomly combined the oscillatory signals and offset, where each condition had a different phase, amplitude, and offset (Fig 2B; Methods; see Churchland & Cunningham [23], their Fig 2, for another illustration).

After applying jPCA, Fig 2C shows that strong rotations exist at the population level for both the representational model (same as Fig 1D) and the dynamical model (28% variance explained in first two planes, 14% in the first plane, 0.97 RGR, 0.98 circularity). In order to test if the underlying condition structure was uniquely responsible for the observed rotations, the CMPT consisted of reassigning task conditions within individual neurons while maintaining the overall covariance matrix between all neurons to a reasonable threshold (95% similarity; Methods). This method disrupts the underlying relationship between neurons and conditions, but not other measures, such as average rate per neuron, relationship between neurons in the population (covariance), and each neuron's contribution to each PC, since the results of PCA are dependent on covariance. If rotations are disrupted as a result of our control, the underlying relationship between neurons and conditions is uniquely essential to the emergence of rotations. On the other hand, intact rotations indicate that many possible condition assignments produce similar rotational patterns, at odds with the findings of Churchland et al. [22] in PMd/ M1 data.

Initially randomly permuting conditions without covariance matching destroyed rotational structure in both the representational and dynamical models (Fig 2D). However, after repeating the CMPT procedure (1000 repetitions) and comparing the RGRs between the observed and permuted data sets to generate a p-value (Methods), we found that the rotational structure found in Fig 2C was restored after covariance matching in the representational model (Fig 2E, $p = 0.71$), but not for the dynamical model ($p < 0.001$, Fig 2E). As a further measure of statistical power, the effect size of rotations in the dynamical model was quite high (effect size: 3.2; Eq 4 in Methods).

In the representational model, permuting disrupts the condition structure, but not the lag relationships, since no data is exchanged between neurons. Once the overall neuron-to-neuron relationship is restored after covariance matching, the rotations are restored as well, even though the condition structure is still disrupted, showing that rotational structure in the representational model does not emerge because of a unique condition structure, as it does in the dynamical model. Repeating the same analysis on additional simulations where neurons were permitted to achieve both positive and negative firing rates ($b_{n,c} = \cos[\theta_c - \theta_n]$ in Eq 1), or when the magnitude of kinematic tuning per neuron varied randomly, did not alter this result ($p = 0.92$ and $p = 0.22$, respectively). Furthermore, the CMPT did not simply 'unshuffle' the data, as there was no significant correlation between the RGR of a given permutation and how similar the condition assignment in that permutation was to the original condition assignment in the observed data (Methods; representational model: r = 0.03, $p = 0.30$; dynamical model: r = 0.03, $p = 0.49$).

It remains an open question whether or not the CMPT can also distinguish rotations arising in a dynamical model from those generated by a complex-kinematic model with varying neuron-kinematic latencies, in which neurons are not only sensitive to velocity, but also to position, acceleration, and occasionally jerk [25]. Therefore, we simulated a population of neurons

106
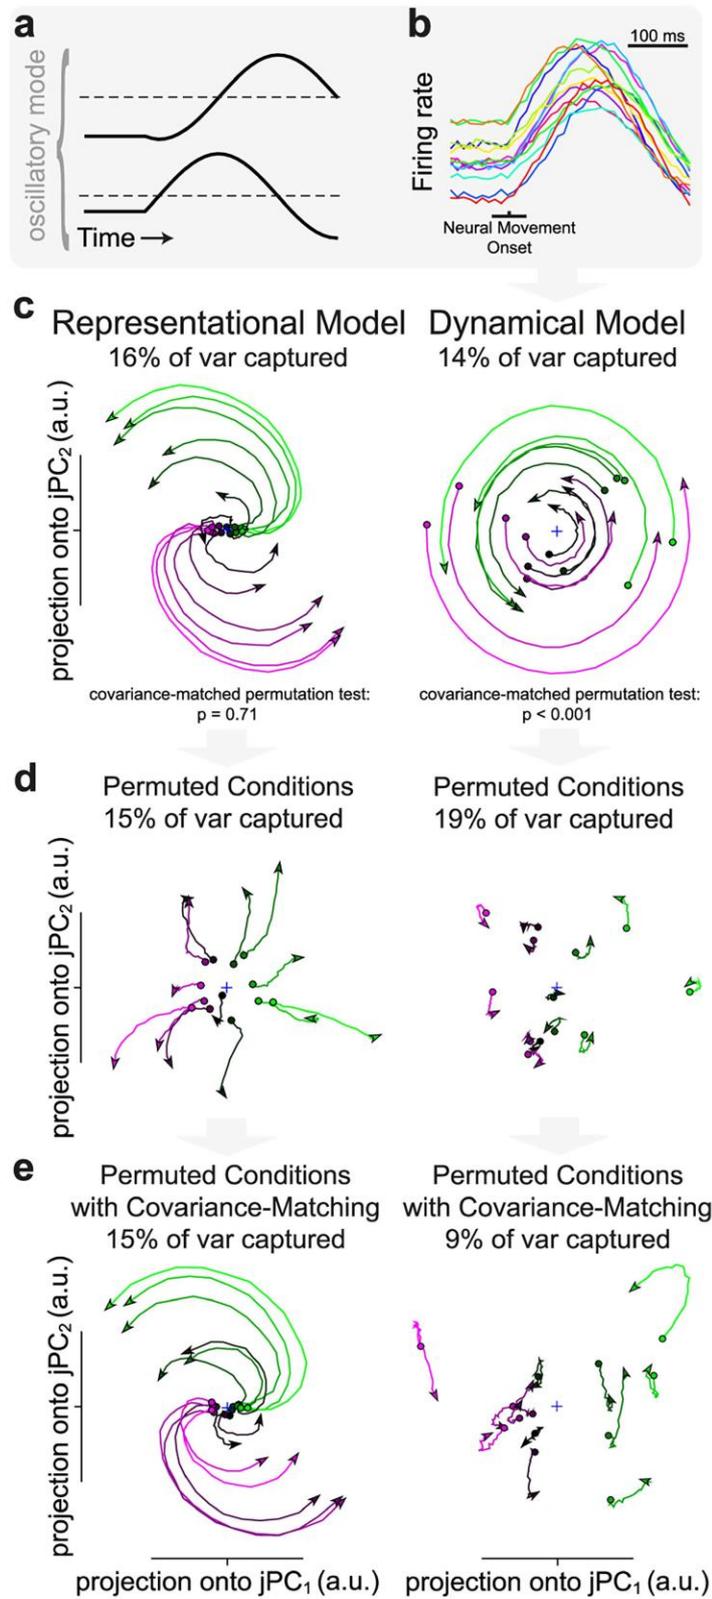
Fig 2. Comparing rotational structure between the representational and the dynamical models. (a) One of the two oscillatory modes (2.8 Hz) used to generate the simulated muscle activity of all conditions (2.8 Hz and 0.3 Hz). **(b)** Firing rate of an example neuron of the dynamical model for all 13 conditions. Each neuron is generated from a random combination of the two underlying oscillatory modes and offset for each condition. **(c)** Rotational dynamics in the first jPCA plane for the observed data. *p*-value shows results of CMPT for the representational and dynamical models evaluated by the rotational goodness-of-fit ratio (RGR: $R^2_{M_{skew}}/R^2_M$). **(d)** Same as **c**, but for permuted data without covariance matching. **(e)** Same as **c**, but for covariance-matched data. Data is plotted for 200 ms regardless of time period used to generate statistics. Colors are based on the preparatory activity in the first jPC.

doi:10.1371/journal.pcbi.1005175.g002

identically to the representational model (Methods; Fig 1D), but further implemented sensitivity to these additional kinematic parameters with the same weights as Churchland et al. [22] (S2A Fig; assuming a reach radius of 20 cm). While complex-kinematic model simulations with no varying neuron-kinematic latencies only produce weak rotations (see Churchland et al. [22], their Fig 4), the inclusion of lags generated rotational structure (S2B Fig; RGR: 0.89, circularity: 0.82). However, similar to the representational model, these rotations were not significant under the CMPT ($p = 0.09$), further emphasizing the power of the CMPT in identifying rotations that are uniquely dependent on the underlying condition structure.

Repeating the CMPT on the representational model for all parameter combinations in Fig 1E revealed that these data generally had no significant rotational structure (*p*-values above 0.05, 100 permutations). Occasionally, *p*-values below 0.05 occurred, but the magnitude of these effects were extremely small and completely disappeared for stricter implementations of the CMPT (similarity 99%), a modification that had no impact on the dynamical model. Taken together, these findings suggest that a broad variety of simulated populations of classically cosine-tuned neurons can exhibit reasonably strong rotational dynamics when explored using jPCA, but that proper controls disrupting the underlying relationship between conditions while conserving other features can distinguish these rotations from those proposed by the dynamical model.

Hallmarks of representational tuning and rotational structure in a recurrent neural network model

Given that representational tuning models have been used extensively to characterize motor cortex activity, how can findings of robust single neuron tuning be reconciled with a dynamical model of movement generation? To address this question, we implemented a simple recurrent neural network (RNN), operating as a dynamical system, from which the velocity profiles required to complete the previously described center-out reaching task can be read out (Fig 3A; Methods). Recent studies have augmented the original findings of Churchland et al. (2012) by generating biologically plausible RNNs that seek to produce complex activity patterns [20,32,33] and using cortical circuit models to explain population activity [34].

In accordance with recent work [20,32], we constructed two time-varying inputs representing the location of the target in 2-D space, and one input representing a hold signal that is released at the go cue. As in the representational model, we generated a network with 200 internal neurons (Methods). The outputs of the network were the x- and y-velocity profiles of the reach. After training, the RNN was able to withhold movement for the entire delay period and execute accurate velocity profiles with a normalized error of less than 0.1% (Fig 3B). Integrating the decoded velocity over time produced the desired kinematics for each reach direction (Fig 3C). A benefit of such a framework is that preparatory activity cancels out at the level of the output signal (null-space), as output must be suppressed during planning to avoid premature movement, a quality observed empirically between PMd/M1 and muscles [35].

107
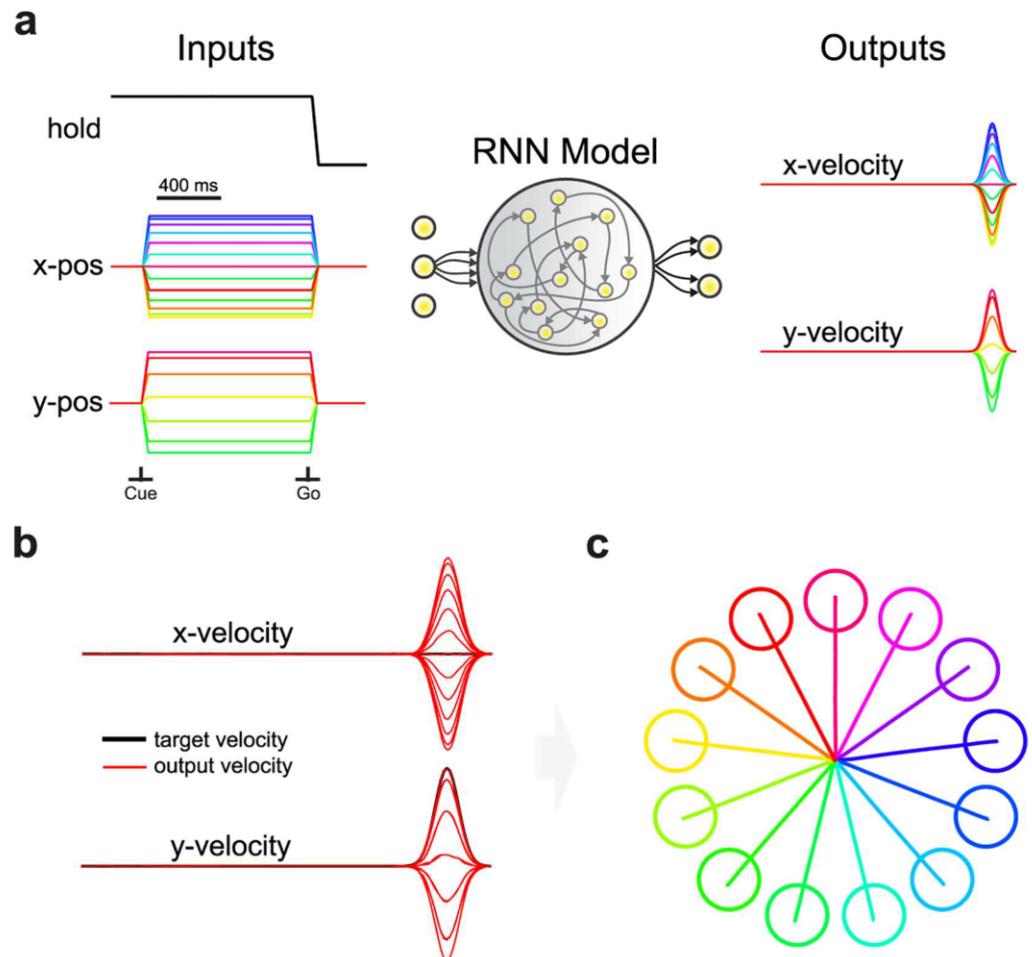
Fig 3. Schematic of recurrent neural network performing center-out reaching. (a) Schematic of RNN, with input layer, hidden layer, and output layer. The three inputs were a condition-independent hold signal that was released at the go cue and two inputs representing the target angle. The two outputs were a linear combination of the internal neurons and read out velocity in the x and y direction. All weights were modified during training. The network received no feedback from the output layer. **(b)** Output velocity profiles produced by the RNN compared with target velocity used in training. The normalized error was less than 0.1%. **(c)** Simulated kinematics produced by integrating the velocity profiles over time, with corresponding targets for illustration.

Fig 4A shows the responses of three example neurons that showed very similar tuning patterns during the delay and movement. Fig 4B shows examples in which the delay tuning was unrelated to movement tuning, and Fig 4C shows examples where the tuning preference flipped at various times during the movement. The overall diversity of tuning is similar to motor cortex neurons presented in Churchland et al. [22] and Sussillo et al. [20].

Fig 4D shows the preferred reach direction (highest firing) of all 200 simulated neurons over time. Preferred directions remained relatively stable during the late delay period, but shortly after the go cue the preferred directions changed rapidly [36]. In this framework the neurons themselves are not explicitly tuned for any given reach direction and are expected to vacillate when the network is released, a property observed previously in a feed-forward network with state feedback (Lillicrap & Scott [7], their Fig 2F).
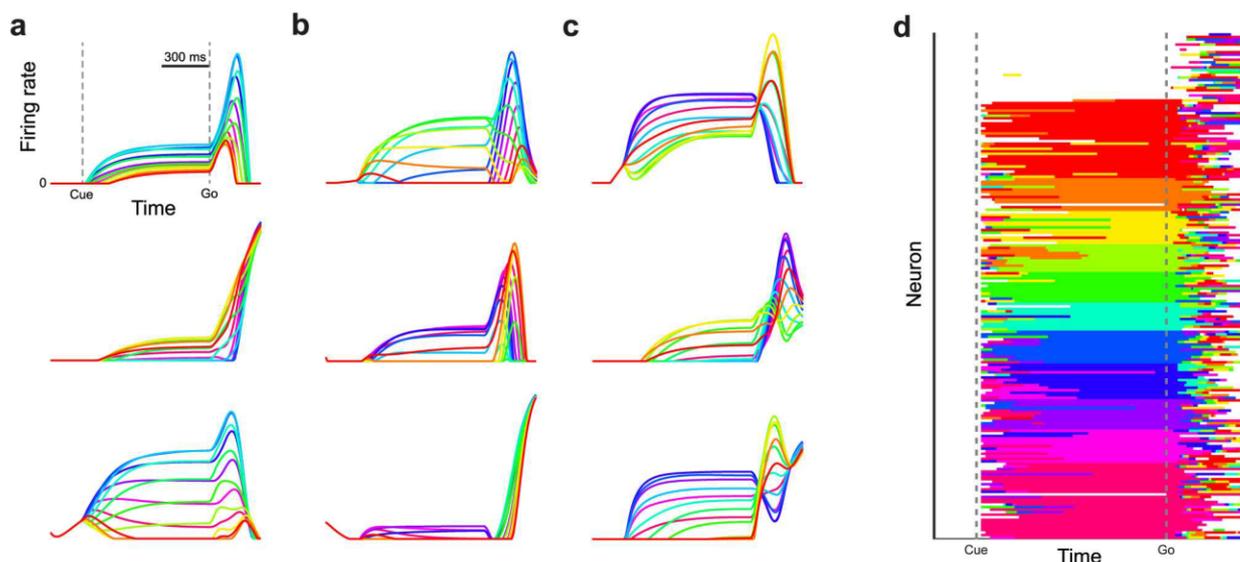
Fig 4. Tuning properties of RNN neurons. (a) Three example units for which the pattern of directional tuning remained highly correlated between the delay period and movement. **(b)** Same as **a**, but for example units that have delay tuning that is not correlated with movement activity. **(c)** Same as **a**, but for example units that invert their tuning between delay and movement. **(d)** Preferred reach direction (highest firing) of all 200 units, sorted by preferred direction at go cue. If there was no firing rate difference (< 1e-4) between the preferred direction and non-preferred direction, units were deemed un-tuned and are marked in white. Firing rates are displayed from 0 to the maximum firing rate of each neuron.

One of the most iconic movement prediction techniques is the population vector, which has been used extensively to decode intended movement direction and instantaneous velocity using knowledge about the preferred direction of all neurons in a population [37,38]. Fig 5A shows the preferred directions of our model neurons (Methods), which were distributed throughout the Cartesian space. Fig 5B shows contribution vectors of all individual neurons over the entire movement of each condition, revealing a remarkably good prediction of movement direction (mirroring results of Georgopoulos et al. [38], their Fig 1). Lastly, Fig 5C shows the result of integrating all population vectors over the course of movement, producing predicted trajectories that well match the desired trajectories (mirroring results of Georgopoulos et al. [38], their Fig 5). In addition, tuning curves of individual RNN neurons visually resembled those observed empirically (S3 Fig). Together, these results reveal that readouts based on the assumption of "preferred direction" can accurately reproduce intended trajectories even when consistent individual neuron tuning was neither included nor observed in the model, a feature of the population vector that has been mathematically outlined by Sanger [39].

As we saw in Fig 4D, preferred direction seemed to fluctuate throughout movement. By correlating the average firing of each neuron for each condition between neural movement onset and later time points during the movement, we can track the stability of tuning over time. The more time has elapsed since neural movement onset, the lower the correlation between delay tuning and movement tuning (Fig 5D; mirroring results of Churchland et al. [40], their Fig 4), both in the model and in example data from PMd/M1 (data from [22], Monkey N). Furthermore, the distribution of correlation coefficients across the population is not bimodal, a finding that would be expected if one subpopulation of neurons was positively correlated over time and one subpopulation inversely tuned during movement.
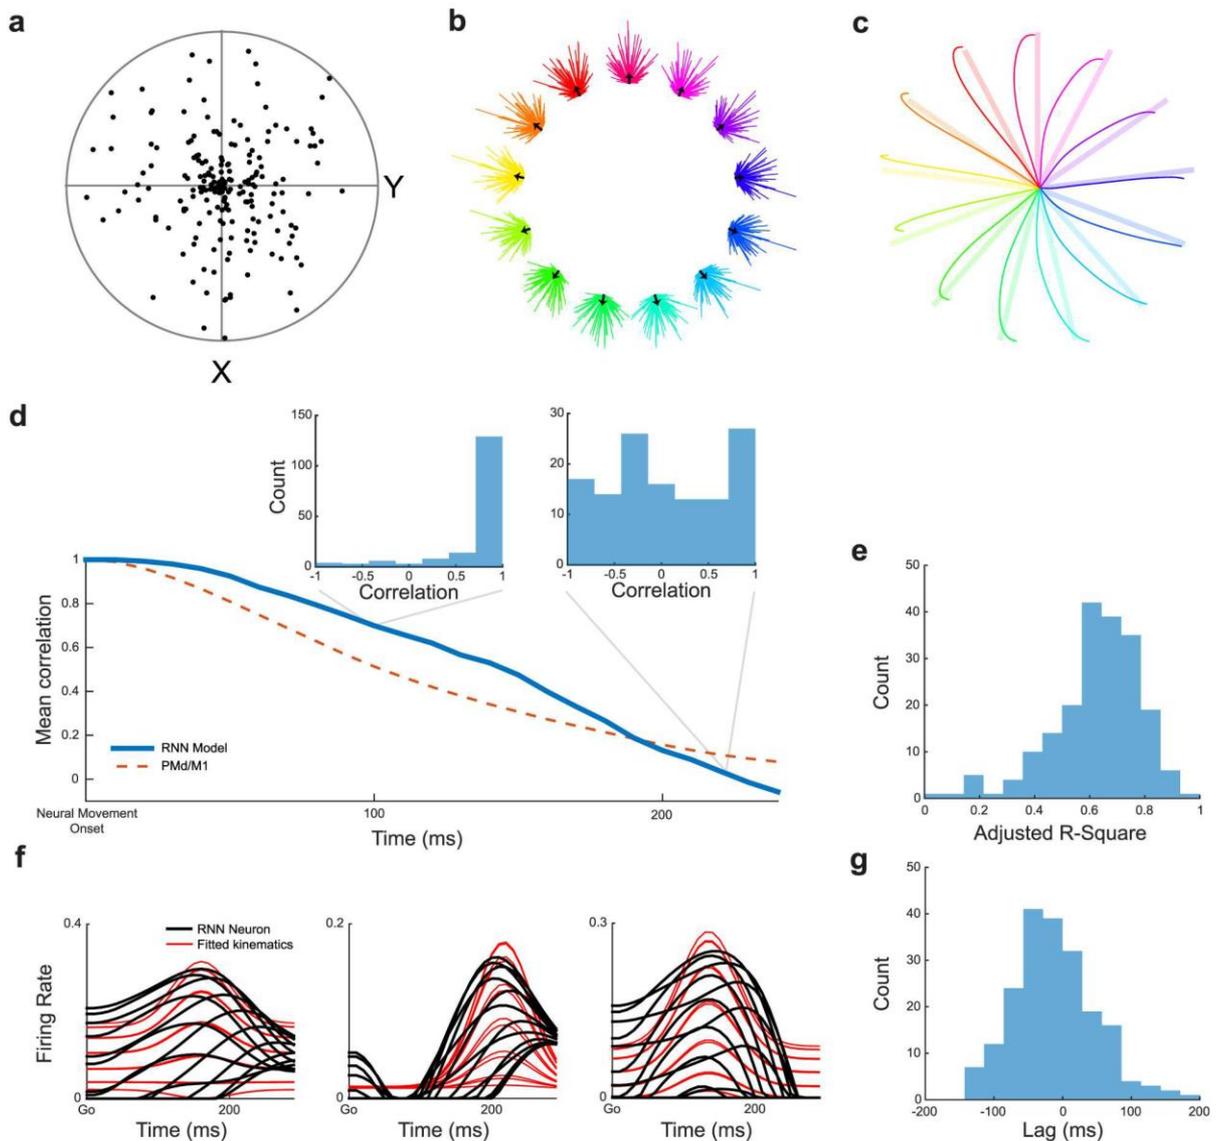
Fig 5. Representational tuning in an RNN for center-out reaching. (**a**) Preferred movement direction in Cartesian space of all units, corresponding to the magnitude of $b_{i,2}$ and $b_{i,3}$ in Eq 9. (**b**) Summary of contribution vectors of all individual neurons (one vector each) over the entire movement, with black population vector showing the overall predicted movement direction. (**c**) Integrating the population vectors in panel **b** over time traces out a predicted trajectory (solid) that largely matches the actual trajectory (translucent). (**d**) Mean correlation between condition tuning order at neural movement onset compared to later time points during movement (in steps of 10 ms) for the RNN model and an example PMd/M1 data set presented in Churchland et al. [22]. Insets show full correlation histograms for two time points. (**e**) Adjusted R-Square obtained by regressing the activity of each neuron (from the go cue to the end of movement 300 ms after go) on a representational cosine model of velocity tuning (Methods). (**f**) Movement activity of three example neurons and the corresponding velocity based regression fits. The overall fit performance to these units is high (Adjusted R-Square above 0.8), but the regression fails to capture the multiphasic and varied nature of the underlying signal. (**g**) Time lag between neural activity and velocity, per neuron, obtained from the velocity tuning regression in panel **e**, showing a large range of values.

Based on the above finding that preferred directions are highly variable during movement, how can it be that representational tuning models explain large amounts of variance in firing rate in empirical studies [24]? Interestingly, regressing the movement activity of each neuron on a full model of velocity tuning (Methods) produces fits very similar to empirical data (Fig

5E, mean Adjusted R-Square: 0.63, mirroring results in Moran and Schwartz [24], their Fig 12A and 12B). However, the actual model fits do not well capture the dynamic properties of the individual units (Fig 5F), such as the changes in preferred direction that occur over the course of the movement or non-linear changes such as when neurons cease firing (0 Hz). Importantly, the optimal neuron-kinematic offsets obtained in the regression cover a range of values, very similar to those observed previously (Fig 5G, mirroring results in Moran and Schwartz [24], their Fig 13A and 13B), providing a potential explanation of how variable neuron-kinematic latencies can improve the performance of representational tuning models even when fixed offsets between neurons is not a property of the underlying circuit.

Yet, it remains unclear if significant rotational structure underlies the activity of our RNN. Therefore, we repeated the jPCA analysis and CMPT with both example data from PMd/M1 and our RNN model. As seen in Fig 6, the PMd/M1 data contained robust rotational structure explaining 56% of the variance in the first two planes (40% in the first plane), an RGR of 0.77 over all jPCA planes, a circularity of 0.63, and the rotational structure was highly significant ($p < 0.001$, CMPT with 1000 repetitions). Importantly, the RNN model also produced robust rotations, explaining 54% of the variance in the first two planes (26% in the first plane), an RGR of 0.74 over all jPCA planes, a circularity of 0.73, and the rotational structure was highly significant ($p < 0.001$, CMPT with 1000 repetitions). In both cases the effect size was also very large, 4.1 and 3.7 for the PMd/M1 data and RNN model, respectively. In addition, similarly to the representational and dynamical models, the CMPT did not simply 'unshuffle' the condition assignment, as the correlations between the RGR of each permutation and the similarity in condition assignment to the observed data was not significant for the PMd/M1 data (r = 0.06, $p = 0.06$) or the RNN model (r = -0.002, $p = 0.94$).

Although significant rotational structure was found in the PMd/M1 data, it is unclear how many recorded neurons and conditions are necessary for jPCA to reveal this result. Therefore, we repeated the CMPT on many subsets of the PMd/M1 data by randomly sampling conditions and neurons to determine how many neurons or conditions might be required to produce statistically significant rotations (Fig 7). This analysis revealed that our test was able to identify clear minima in number of neurons and conditions that are necessary to achieve significance, in general more than 30 neurons and more than 8 conditions, a finding that may guide the design of future experiments and encourages skepticism of experiments with small numbers of neurons or conditions.

It is important to note that the CMPT may not necessarily distinguish between all possible models, as there exist cases of the dynamical model for which our test would find no significant rotational structure. For example, if the required oscillator phases required to fit muscle activity were identical between all conditions, while rotational structure would be found using jPCA, our test would find these rotations to be non-significant. Therefore, we do not propose the CMPT as a singular test of rotational structure to accompany jPCA, but rather as an additional control.

We posit that future studies should seek to explain single neuron characteristics as a function of population or circuit activity rather than imbue single neurons with complex tuning characteristics [9,10]. Furthermore, RNNs provide an ideal medium for more detailed study, as the ground truth of synaptic connectivity, plasticity, noise, trial-to-trial variability, and responses to unexpected perturbations are known and can be manipulated directly. However, "exploring an artificial model universe comes with its own risk" [41] and proper models must resist the temptation of explaining purely idiosyncratic properties, but rather those that are able to explain large amounts of variance in electrophysiological data. Our results also emphasize that explaining a large amount of variance in neural data in and of itself does not necessary lead to mechanistic insight [42], as the observation of rotational structure arose under multiple
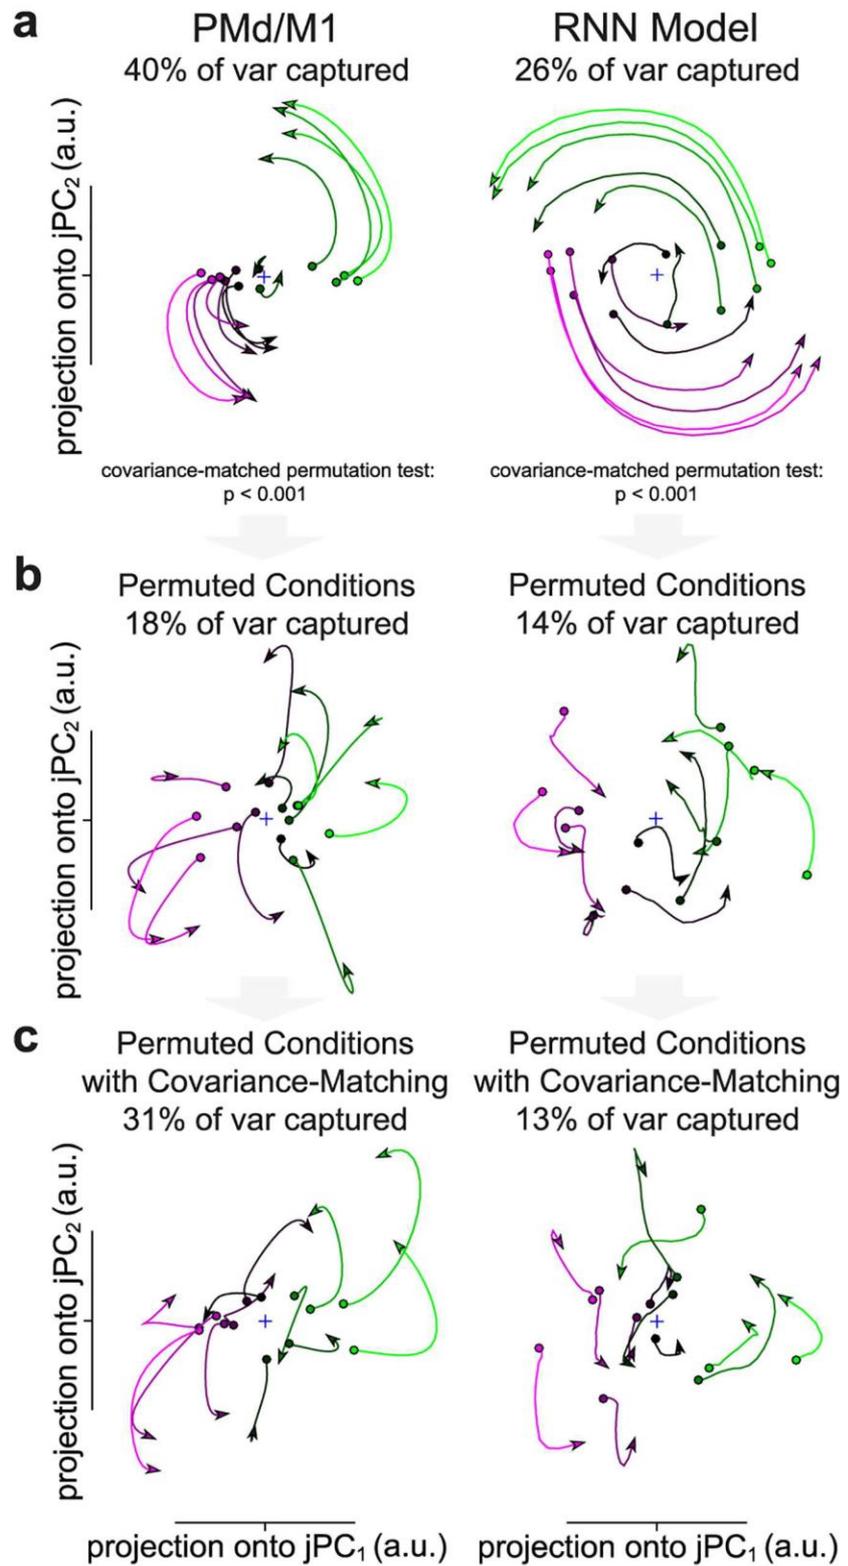
Fig 6. Significant rotational structure in PMd/M1 data and RNN model. Comparison of rotational dynamics for (**a**) observed, (**b**) permuted without covariance matching, and (**c**) covariance-matched data in the first jPCA plane. *p*-values in **a** are from the CMPT for the rotational goodness-of-fit ratio (RGR: $R^2_{M_{skew}}/R^2_M$) in all jPCA planes. Conditions and neurons were randomly down-sampled in the PMd/M1 data to match the RNN model. Data is plotted for 200 ms regardless of time period used to generate statistics. Colors are based on the preparatory activity in the first jPC.

doi:10.1371/journal.pcbi.1005175.g006

models, and future work is needed to determine the biological circuit mechanism underlying population level rotational structure.

Fundamentally, as representational [43,44] and dynamical [20,32,45] models become more complex in their implementations, their ability to explain empirical data becomes more striking and convincing. Ultimately, what will signify the usefulness of either framework will be their utility in generating testable hypotheses of how the brain executes complex behavior in basic research contexts, and in developing new solutions in applied research contexts. In terms of application, the representational view has been indispensable in developing neural prosthetics for paralyzed patients [46–48], but this trend may be changing as prosthetic algorithms are augmented by the inclusion of dynamical systems into their underlying framework [49,50].

Methods

Representational model

Preparatory and movement activity were simulated for a population of 200 neurons in a 13-direction center-out reaching task. Neurons were cosine-tuned for velocity during both preparation and movement with respect to their randomly assigned (uniform) preferred direction. The average firing rate, $f_{n,c}$, of a given simulated neuron, n, for a particular reach condition, c, at time t is given by,

$$f_{n,c}(t, \tau_n, \sigma) = \begin{cases} b_{n,c}e^{-\frac{(t - \tau_n - \mu_0)^2}{2\sigma^2}} + \varepsilon, & t \geq \tau_n \\ \varphi b_{n,c} + \varepsilon, & t < \tau_n \end{cases} \quad, \quad b_{n,c} = \frac{1 + cos[\theta_c - \theta_n]}{2} \quad (1)$$
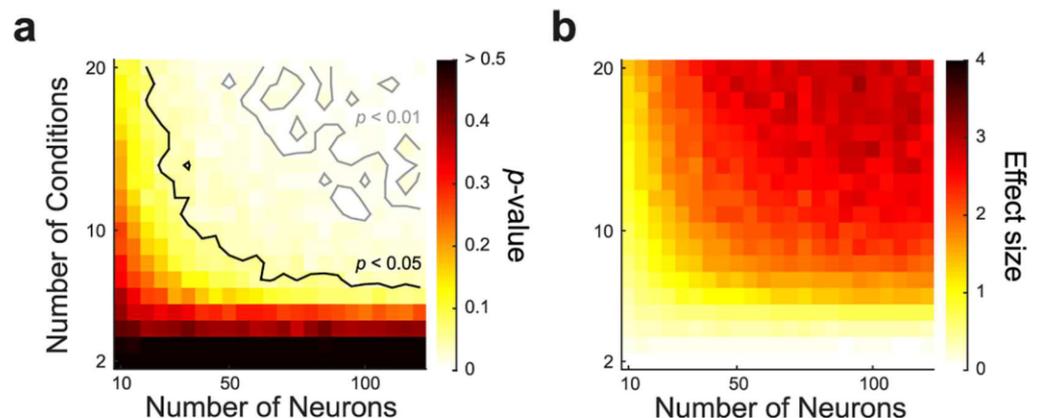

Fig 7. Number of neurons and conditions required for statistically significant rotations. The CMPT was carried out (500 repetitions) for many subsets of example PMd/M1 data including from 10–120 neurons and 2–20 conditions. (**a**) Map of *p*-values for the rotational goodness-of-fit ratio (RGR: $R^2_{M_{skew}}/R^2_M$). (**b**) Map of effect size (difference between observed RGR and mean of permuted distribution, divided by the SD of the permuted distribution). For every permutation, random neurons and conditions were drawn from the example set. Contours show the 0.05 and 0.01 significance levels.

doi:10.1371/journal.pcbi.1005175.g007

where τ_n is the neural response latency (normally distributed) of each neuron, σ is the duration parameter of the movement activity, which never differed between neurons of the same simulation, $b_{n,c}$ is the gain factor for each neuron and condition, θ_c is the angle of the reach target in condition c, θ_n is the preferred reach angle of neuron n, φ a constant which determines the magnitude of preparatory activity, μ_0 a constant and given by $\mu_0 = \sigma\sqrt{-2ln\varphi}$, and ε is random noise drawn from a normal distribution.

For all analyses, φ was fixed at 0.2, i.e., preparatory activity was always one fifth of the maximum movement activity for that condition; however, our results do not depend on this factor. The distribution of latency factors, τ_n, and the movement duration parameter, σ, were varied systematically to produce the results in Fig 1E. For visualization purposes we chose the noise distribution, ε, to have an SD of 0.01 for all analyses. However, this value did not greatly affect the outcome. We found it necessary to increase the noise more than 300 times to eliminate all structure.

Rotational dynamics

jPCA is a method for finding linear combinations of principal components that capture rotational structure in data (Churchland et al., 2012). The method is based on finding a transformation between a neural system at each time point and its temporal derivative, using the following steps. First, the average firing rate of many neurons is extracted and aligned to the execution of a movement, starting whenever the neural activity begins rapidly changing preceding movement onset, termed neural movement onset (typically 100–200 ms before overt movement). Next, each neuron is normalized and reduced, using standard principal component analysis, to a set of principal components, X_{red}, of size $d \times ct$, in which the d largest components are retained, and c is the number of conditions and t is the number of time points selected. Then, via linear regression, the unconstrained matrix M and the skew-symmetric matrix M_{skew} (where $M_{skew} = -M_{skew}{}^T$) can be found to satisfy $\dot{X}_{red} = MX_{red}$ and $\dot{X}_{red} = M_{skew}X_{red}$, where $\dot{X}_{red}$ is the difference in adjacent time points of X_{red} (temporal derivative). The jPCA planes are then constructed from the eigenvectors of M_{skew}, with the added constraint that the net rotation in each plane is anticlockwise.

In order to avoid finding spurious rotations, only the first 6 PCA dimensions explaining the most variance were fed into the jPCA algorithm (sampled in steps of 10 ms). For the representational model, the data fed into jPCA began at neural movement onset, which was defined as the time when the average signal exceeded 10% of the difference between preparatory activity and maximum activity, and ended when the average activity fell below this level. Given the variable lags between neurons, it was necessary to define the above procedure for determining neural movement onset, which is similar to the one performed by Churchland et al. [22]. For the dynamical model and RNN model, neural movement onset was simply defined as the time of the go cue, and the entire movement (300 ms) was used. On the other hand, for the example PMd/M1 data (presented as Monkey N in Churchland et al. [22]) the analysis was replicated as done in Churchland et al. [22], from -50 to 150 ms relative to neural movement onset, a time-course specifically chosen to avoid sensory feedback not present in our simulation. jPCA was performed using a freely available toolbox (http://churchlandlab.neuroscience.columbia.edu/links.html).

Dynamical model

The dynamical model is based on the finding that muscle activity during reaching can be well explained by a summation of the lagging components of two oscillatory modes, each with a fixed frequency, but with varying phase, amplitude, and offset for each movement condition

[22,23]. Following Churchland et al. [22], we simulated for each condition $c = 1,\ldots,13$ an offset o_c, and the two complex oscillations ($k = 1,2$)

$$F_{c,k} = a_{c,k}\, e^{i(2\pi f_k\ t - \theta_{c,k})} \tag{2}$$

for the two underlying frequency modes $f_1 = 2.8\ Hz$ and $f_2 = 0.3\ Hz$ (however, the specific frequencies used did not alter the results). Phases, $\theta_{c,k}$, amplitudes, $a_{c,k}$, and offset, o_c, were randomly drawn for each condition to match both the variance explained per plane and the similarity between conditions in the representational model (phase drawn from uniform distribution, range: 0 to $\pi/2$; amplitude drawn from uniform distribution, range: -1.5 to -2.5; offset drawn from uniform distribution, range: -4.5 to -5.5). For simplicity, we did not implement the windowed gamma functions used in Churchland et al. [22], as these only increase the realism of neural responses and do not contribute to the main result.

To generate simulated neurons in the dynamical model (N = 200), the activity $r_{n,c}(t)$ of each neuron $n \in \{1,\ldots,N\}$ and condition c was generated as a neuron-specific combination of the condition-specific oscillations and offset ($F_{c,1}, F_{c,2}, o_c$)

$$r_{n,c}(t) = Re(\boldsymbol{w}_{n,1}\, F_{c,1}(t) + \boldsymbol{w}_{n,2}\, F_{c,2}(t)) + s_n o_c + \varepsilon_{n,c}(t) \tag{3}$$

with the real and imaginary components of the complex coefficients, $\boldsymbol{w}_{n,1}$ and $\boldsymbol{w}_{n,2}$, and the offset coefficient s_n drawn from a standard normal distribution (zero mean, unit variance). As described above, each neuron had a unique set of 5 weights that were used for all conditions, and the small amount of normally distributed noise $\varepsilon_{n,c}(t)$ that was matched to the representational model. Preparatory activity ($r_{n,c}(t)$ for t<0) was generated by simply extending the first data point, i.e., $r_{n,c}(0)$ including noise, for 100 ms back in time.

Covariance-matched permutation test for rotational dynamics

In order to test if the rotational dynamics found in neural population data depended on the underlying condition structure, we developed a covariance-matched permutation test to disrupt the condition-wise relationships while sparing other features of the data. In this iterative procedure the entire time-course of each condition, for each neuron separately, was first randomly reassigned to another condition. Then, individual pairs of conditions were randomly exchanged (within, but not between random neurons) and the similarity of the covariance of all neurons was compared to the observed data for the time period of interest (i.e. the time period analyzed using jPCA). Covariance between neurons was calculated from the matrix $n \times ct$, where n is the number of neurons, c the number of conditions, and t the time period of interest. Covariance similarity was calculated as the sum of the squared difference between the observed covariance matrix and the covariance matrix of the permuted data, divided by the variance of the observed covariance matrix. If the similarity was increased by a given permutation, it was accepted, otherwise it was rejected and the process continued. When the covariance similarity between the observed data and the permuted data exceeded 95%, the process was complete (this process generally lasted many thousands of permutations). In this way, no data values were altered. Correspondingly, the average firing rates, total firing rates of all neurons, and the approximate covariance relationship between all neurons were conserved.

To test significance, the permutation procedure was repeated many times (100–1000 repetitions) and the covariance-matched data was fed into jPCA in the same fashion as the observed data. The rotational goodness-of-fit ratio (RGR: $R^2_{M_{skew}}/R^2_M$) over all three jPCA planes (spanning 6 principal components), which provides a measure of how much variance can be explained by rotational dynamics, was evaluated for all permutations. Subsequently, the

PLOS COMPUTATIONAL BIOLOGY

fraction of repetitions that the above statistic computed from the permuted data exceeded the observed data determined the p-value, as is standard procedure for permutation tests.

To measure statistical power, effect size was computed as

$$effect\ size = \frac{RGR_{observed} - \overline{RGR}_{permuted}}{\sigma_{RGR_{permuted}}} \tag{4}$$

similar to Cohen's d.

To test if the results of the CMPT procedure were simply due to 'unshuffling' the conditions and restoring the original condition assignment, we checked what percentage of the assignment matrix ($c \times n$) retained its original condition assignment at the end of the CMPT and correlated this measure with the RGR of the corresponding permutation repetition. Importantly, since even 'unshuffled' data would not be guaranteed to be in the same order as in the original data, before correlating we first sorted the rows of the above mentioned assignment matrix by the most common condition in each row (condition 1 most common assigned to row 1, condition 2 most common assigned to row 2, etc.), to achieve the most conservative comparison possible.

Recurrent neural network

In order to examine a system in which velocity profiles of a 13-direction center-out reaching task could be read out over time, we implemented the dynamical system, $\dot{x} = F(x, u)$, using a standard continuous RNN equation of the form

$$\tau \dot{x}_i(t) = -x_i + \sum_{k=1}^{N} J_{ik} r_k(t) + \sum_{k=1}^{I} B_{ik} u_k(t) \tag{5}$$

where the network has N units and I inputs, x are the activations and r the firing rates in the network, which were related to the activations by the rectified hyperbolic tangent function, such that $r = \begin{cases} 0, & x < 0 \\ \tanh(x), & x \geq 0 \end{cases}$. The units in the network interact using the synaptic weight matrix, J. The inputs are described by u and enter the system by input weights, B. The time integration constant of the network is τ.

For all simulations N was fixed at 200. The three inputs were a condition-independent hold signal that was released at the go cue, and two inputs representing the target position that corresponded to $sin(\theta)$ and $cos(\theta)$, where θ is the angle of the target around the circle. The elements of B were initialized to have zero mean (normally distributed values with $SD = 1/\sqrt{N}$). The elements of J were initialized to have zero mean (normally distributed values with $SD = g/\sqrt{N}$), where the synaptic scaling factor, g, was set at 1.5 [51]. We used a fixed time constant of 50 ms for τ, with Euler integration every 10 ms.

To produce the two desired velocity profile outputs, which were the x-velocity and the y-velocity of the 13-direction center-out reaching task described previously (bell-shaped velocity profiles lasting 300 ms), we defined a linear readout of the internal network

$$z_i(c, t) = \sum_{k=1}^{N} W_{ik} r_k(c, t) \tag{6}$$

where z represents the two velocity readouts ($i = 1,2$) and is a linear combination of the internal firing rates using weight matrix W, which was initialized to all zero values.

The input weights, B, internal connectivity, J, and output weights, W, were trained using Hessian Free Optimization [52] using freely available code (https://github.com/sussillo/hfopt-matlab)

116

also utilized in Sussillo et al. [20]. The error function used to optimize the network considered the difference between the output of the linear readout and the desired velocity profiles, v,

$$E_i(c,t) = z_i(c,t) - v_i(c,t) \tag{7}$$

at each time point, t, each output dimensions, i, and each movement direction, c. We report normalized error, which is the sum of the squared error from Eq 7 over all times, dimensions, and conditions, divided by the total variance of the target signal. In addition to the above error signal, we also implemented three regularizations designed to encourage the network to produce biologically-plausible activity (implemented as in Sussillo et al. [20]). The three penalties were a cost on the mean firing rate, the squared-sum of the input and output weights, and a penalty encouraging the network to avoid complex state trajectories (similar to local space contraction [53]). The hyper-parameters used for these regularization were 1e-2, 2e-5, and 5e-5, respectively.

In order to discourage internally-timed responses, the network was trained to produce movements after three varying delays of 600, 800, and 1000 ms. All results used came from the 800 ms delay set and the reaction time (time between go cue and movement onset) of the network was fixed at 100 ms. We opted not to model any feedback, since the goal of the study was to illustrate the main points parsimoniously and without relying on confronting the issue of what kind of feedback is most biologically plausible in such a network.

Population vector

The population vector decoding technique was performed as described in Georgopoulos et al. [38] and Schwartz et al. [37]. Specifically, the preferred direction of each neuron was determined via linear regression

$$R_{i,c} = b_{i,1} + b_{i,2} \sin\theta_c + b_{i,3} \cos\theta_c \tag{8}$$

where R is the average firing rate of neuron i over time from the go cue to the end of movement (300 ms after go) for condition c, b are constants, and θ_c is the angle of the current target. The preferred direction of each neuron was then defined as

$$C_i = \left[\frac{b_{i,2}}{k_i}, \frac{b_{i,3}}{k_i} \right] \tag{9}$$

where

$$k_i = \sqrt{b_{i,2}^2 + b_{i,3}^2} \tag{10}$$

To make predictions about direction and magnitude of movement [38], the population vector at time t during movement was computed using the instantaneous firing rate of all neurons (R) and each neuron's previously determined preferred direction (C), such that

$$P(t) = \sum_{i=1}^{N} (R_i(t) - b_{i,1}) C_i \tag{11}$$

where N is the number of neurons. The sum of P over all time points during the movement of a given trial then determined the overall predicted direction of movement. Alternatively, P could be integrated over time points to trace out a predicted trajectory, as in Fig 5C. Fitting procedure was performed using the Matlab *fit* function using the least-squares method.

Velocity regression

In order to investigate the presence of representational tuning in the RNN, we regressed the movement period activity of each neuron (starting at the go cue until the end of movement 300 ms after go) on the following model of directional and speed tuning,

$$R(t - \tau) = a_1 + \|\overrightarrow{V}(t)\|(a_2 + a_3 sin[\theta] + a_4 cos[\theta]) + a_5(a_3 sin[\theta] + a_4 cos[\theta]) \tag{12}$$

where R is instantaneous neural activity, τ is the time lag between neural activity and its expression as movement, a are constants, θ is the direction of the current target, which stays constant during center-out reaches, and $\overrightarrow{V}$ is the velocity profile. Fitting procedure and resulting goodness-of-fit statistics were obtained using the Matlab *fit* function using the least-squares method. The final term of the equation was appended in addition to the factors presented in Moran & Schwartz (1999) in order to account for differences in preparatory activity between reach directions, an aspect not utilized in the original experiment when no delay period was present. Tuning during the preparatory period was the same as during movement, scaled by a factor, a_5, which also allowed for inverted tuning during movement.

Supporting Information

S1 Fig. Latency offsets produce derivative-like principal components. (a) Firing rates of six simulated neurons (normal distributions with identical SD) over time with random time offsets (drawn from normal distribution). (b) The first three principal components of the simulated units. (c) The plane formed by the first two principal components, showing a 'horseshoe' pattern.
(EPS)

S2 Fig. Simulation of a complex-kinematic tuning based model with variable neuron-kinematic latencies. (a) Four example neurons with differing latencies. (b-d) Comparison of rotational dynamics for (b) observed, (c) permuted without covariance matching, and (d) covariance-matched data in the first jPCA plane. *p*-value in **b** are from the CMPT for the rotational goodness-of-fit ratio (RGR: $R^2_{M_{skew}}/R^2_M$) in all jPCA planes. Data is plotted for 200 ms regardless of time period used to generate statistics. Colors are based on the preparatory activity in the first jPC.
(EPS)

S3 Fig. Tuning curves of RNN neurons during movement. Mean firing rate during the movement epoch of all movement directions for 16 randomly selected RNN neurons.
(EPS)

Author Contributions

Conceptualization: JAM BD HS.

Funding acquisition: HS.

Investigation: JAM BD.

Writing – original draft: JAM.

Writing – review & editing: JAM BD HS.

References

1. Scott SH. Inconvenient truths about neural processing in primary motor cortex. J Physiol (Lond). 2008; 586: 1217–1224.

2. Scott SH, Gribble PL, Graham KM, Cabel DW. Dissociation between hand motion and population vectors from neural activity in motor cortex. Nature. 2001; 413: 161–165. doi: 10.1038/35093102 PMID: 11557980

3. Kurtzer I, Herter TM. Contrasting interpretations of the nonuniform distribution of preferred directions within primary motor cortex. J Neurophysiol. 2007; 97: 4390–author reply 4391–2.

4. Kurtzer I, Herter TM, Scott SH. Nonuniform distribution of reach-related and torque-related activity in upper arm muscles and neurons of primary motor cortex. J Neurophysiol. 2006; 96: 3220–3230. doi: 10.1152/jn.00110.2006 PMID: 17005623

5. Naselaris T, Merchant H, Amirikian B, Georgopoulos AP. Large-scale organization of preferred directions in the motor cortex. I. Motor cortical hyperacuity for forward reaching. J Neurophysiol. 2006; 96: 3231–3236. doi: 10.1152/jn.00487.2006 PMID: 16971681

6. Georgopoulos AP, Naselaris T, Merchant H, Amirikian B. Reply to Kurtzer and Herter. J Neurophysiol. 2007; 97: 4391–4392.

7. Lillicrap TP, Scott SH. Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. Neuron. 2013; 77: 168–179. doi: 10.1016/j.neuron.2012.10.041 PMID: 23312524

8. Georgopoulos AP, Carpenter AF. Coding of movements in the motor cortex. Current Opinion in Neurobiology. 2015; 33: 34–39. doi: 10.1016/j.conb.2015.01.012 PMID: 25646932

9. Shenoy KV, Sahani M, Churchland MM. Cortical Control of Arm Movements: A Dynamical Systems Perspective. Annu Rev Neurosci. 2013; 36: 337–359. doi: 10.1146/annurev-neuro-062111-150509 PMID: 23725001

10. Yuste R. From the neuron doctrine to neural networks. Nat Rev Neurosci. 2015; 16: 487–497. doi: 10.1038/nrn3962 PMID: 26152865

11. Tanji J, Evarts EV. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. J Neurophysiol. 1976; 39: 1062–1068. PMID: 824409

12. Evarts EV. Pyramidal tract activity associated with a conditioned hand movement in the monkey. J Neurophysiol. 1966; 29: 1011–1027. PMID: 4961643

13. Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci. Society for Neuroscience; 1982; 2: 1527–1537. PMID: 7143039

14. Gao P, Ganguli S. On simplicity and complexity in the brave new world of large-scale neuroscience. Current Opinion in Neurobiology. 2015; 32: 148–155. doi: 10.1016/j.conb.2015.04.003 PMID: 25932978

15. Churchland MM, Yu BM, Sahani M, Shenoy KV. Techniques for extracting single-trial activity patterns from large-scale neural recordings. Current Opinion in Neurobiology. 2007; 17: 609–618. doi: 10.1016/j.conb.2007.11.001 PMID: 18093826

16. Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV. Single-trial neural correlates of arm movement preparation. Neuron. 2011; 71: 555–564. doi: 10.1016/j.neuron.2011.05.047 PMID: 21835350

17. Michaels JA, Dann B, Intveld RW, Scherberger H. Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. J Neurosci. 2015; 35: 11415–11432. doi: 10.1523/JNEUROSCI.1714-15.2015 PMID: 26269647

18. Churchland MM, Shenoy KV. Delay of movement caused by disruption of cortical preparatory activity. J Neurophysiol. 2007; 97: 348–359. doi: 10.1152/jn.00808.2006 PMID: 17005608

19. Ames KC, Ryu SI, Shenoy KV. Neural Dynamics of Reaching following Incorrect or Absent Motor Preparation. Neuron. 2014; 81: 438–451. doi: 10.1016/j.neuron.2013.11.003 PMID: 24462104

20. Sussillo D, Churchland MM, Kaufman MT, Shenoy KV. A neural network that finds a naturalistic solution for the production of muscle activity. Nat Neurosci. 2015; 18: 1025–1033. doi: 10.1038/nn.4042 PMID: 26075643

21. Pandarinath C, Gilja V, Blabe CH, Nuyujukian P, Sarma AA, Sorice BL, et al. Neural population dynamics in human motor cortex during movements in people with ALS. eLife. 2015; 4: e07436. doi: 10.7554/eLife.07436 PMID: 26099302

119

22. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, et al. Neural population dynamics during reaching. Nature. 2012; 487: 51–56. doi: 10.1038/nature11129 PMID: 22722855

23. Churchland MM, Cunningham JP. A Dynamical Basis Set for Generating Reaches. Cold Spring Harb Symp Quant Biol. 2014; 79: 67–80. doi: 10.1101/sqb.2014.79.024703 PMID: 25851506

24. Moran DW, Schwartz AB. Motor cortical representation of speed and direction during reaching. J Neurophysiol. 1999; 82: 2676–2692. PMID: 10561437

25. Todorov E. Direct cortical control of muscle activation in voluntary arm movements: a model. Nat Neurosci. 2000; 3: 391–398. doi: 10.1038/73964 PMID: 10725930

26. Ashe J, Georgopoulos AP. Movement parameters and neural activity in motor cortex and area 5. Cerebral Cortex. 1994; 4: 590–600. PMID: 7703686

27. Schwartz AB. Motor cortical activity during drawing movements: population representation during sinusoid tracing. J Neurophysiol. 1993; 70: 28–36. PMID: 8360717

28. Humphrey DR, Schmidt EM, Thompson WD. Predicting measures of motor performance from multiple cortical spike trains. Science. 1970; 170: 758–762. PMID: 4991377

29. Podani J, Miklos I. Resemblance coefficients and the horseshoe effect in principal coordinates analysis. Ecology. 2002; 83: 3331–3343.

30. Novembre J, Stephens M. Interpreting principal component analyses of spatial population genetic variation. Nat Genet. 2008; 40: 646–649. doi: 10.1038/ng.139 PMID: 18425127

31. Menozzi P, Piazza A, Cavalli-Sforza L. Synthetic maps of human gene frequencies in Europeans. Science. 1978; 201: 786–792. PMID: 356262

32. Hennequin G, Vogels TP, Gerstner W. Optimal control of transient dynamics in balanced networks supports generation of complex movements. Neuron. 2014; 82: 1394–1406. doi: 10.1016/j.neuron.2014.04.045 PMID: 24945778

33. Renart A. Bringing the dynamics of movement under control. Neuron. 2014; 82: 1193–1195. doi: 10.1016/j.neuron.2014.06.002 PMID: 24945762

34. Murakami M, Mainen ZF. Preparing and selecting actions with neural populations: toward cortical circuit mechanisms. Current Opinion in Neurobiology. 2015; 33C: 40–46.

35. Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. Cortical activity in the null space: permitting preparation without movement. Nat Neurosci. 2014; 17: 440–448. doi: 10.1038/nn.3643 PMID: 24487233

36. Churchland MM, Shenoy KV. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. J Neurophysiol. 2007; 97: 4235–4257. doi: 10.1152/jn.00095.2007 PMID: 17376854

37. Schwartz AB, Kettner RE, Georgopoulos AP. Primate motor cortex and free arm movements to visual targets in three- dimensional space. I. Relations between single cell discharge and direction of movement. Society for Neuroscience; 1988; 8: 2913–2927.

38. Georgopoulos AP, Kettner RE, Schwartz AB. Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. 1988; 8: 2928–2937.

39. Sanger TD. Theoretical considerations for the analysis of population coding in motor cortex. Neural Comput. 1994; 6: 29–37.

40. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? Neuron. 2010; 68: 387–400. doi: 10.1016/j.neuron.2010.09.015 PMID: 21040842

41. O'Leary T, Sutton AC, Marder E. Computational models in the age of large datasets. Current Opinion in Neurobiology. 2015; 32: 87–94. doi: 10.1016/j.conb.2015.01.006 PMID: 25637959

42. Jonas E, Kording K. Could a neuroscientist understand a microprocessor? bioRxiv. 2016.

43. Agarwal R, Thakor NV, Sarma SV, Massaquoi SG. PMv Neuronal Firing May Be Driven by a Movement Command Trajectory within Multidimensional Gaussian Fields. J Neurosci. 2015; 35: 9508–9525. doi: 10.1523/JNEUROSCI.2643-14.2015 PMID: 26109672

44. Castellanos L, Vu VQ, Perel S, Schwartz AB. A multivariate gaussian process factor model for hand shape during reach-to-grasp movements. Statistica Sinica. 2015; 25.

45. Zibner S, Tekülve J, Schöner G. The neural dynamics of goal-directed arm movements: a developmental perspective. 5th International Conference on Development and Learning and Epigenetic Robotics. Providence, RI; 2015.

46. Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, et al. High-performance neuroprosthetic control by an individual with tetraplegia. The Lancet. Elsevier; 2012.

120

47. Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. Nature. 2012; 485: 372–375. doi: 10.1038/nature11076 PMID: 22596161

48. Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. Cortical control of a prosthetic arm for self-feeding. Nature. 2008; 453: 1098–1101. doi: 10.1038/nature06996 PMID: 18509337

49. Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, Shenoy KV. Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. Nat Commun. 2015; 6: 7759. doi: 10.1038/ncomms8759 PMID: 26220660

50. Sussillo D, Nuyujukian P, Fan JM, Kao JC, Stavisky SD, Ryu S, et al. A recurrent neural network for closed-loop intracortical brain-machine interface decoders. J Neural Eng. 2012; 9: 026027. doi: 10.1088/1741-2560/9/2/026027 PMID: 22427488

51. Sussillo D, Abbott LF. Generating coherent patterns of activity from chaotic neural networks. Neuron. 2009; 63: 544–557. doi: 10.1016/j.neuron.2009.07.018 PMID: 19709635

52. Martens J, Sutskever I. Learning recurrent neural networks with hessian-free optimization. Proceedings of the 28th International Conference on Machine Learning. Bellevue, WA, USA; 2011.

53. Rifai S, Muller X, Glorot X, Mesnil G, Bengio Y. Learning invariant features through local space contraction. arXiv. 2011;1104.4153.
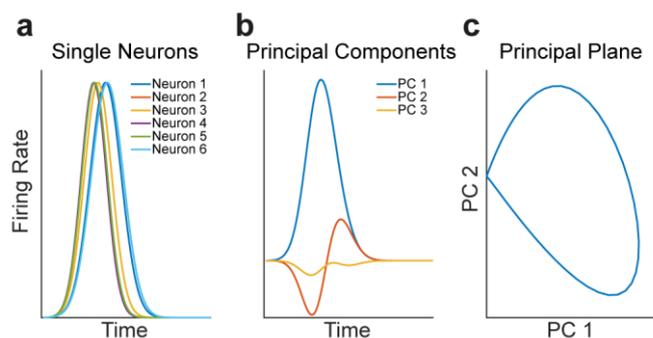
121

Supporting Information

S1 Fig. Latency offsets produce derivative-like principal components. (**a**) Firing rates of six simulated neurons (normal distributions with identical SD) over time with random time offsets (drawn from normal distribution). (**b**) The first three principal components of the simulated units. (**c**) The plane formed by the first two principal components, showing a 'horseshoe' pattern.

S2 Fig. Simulation of a complex-kinematic tuning based model with variable neuron-kinematic latencies. (**a**) Four example neurons with differing latencies. (**b-d**) Comparison of rotational dynamics for (**b**) observed, (**c**) permuted without covariance matching, and (**d**) covariance-matched data in the first jPCA plane. *p*-value in **b** are from the CMPT for the rotational goodness-of-fit ratio (RGR: $R^2_{M_{skew}}/R^2_M$) in all jPCA planes. Data is plotted for 200 ms regardless of time period used to generate statistics. Colors are based on the preparatory activity in the first jPC.

S3 Fig. Tuning curves of RNN neurons during movement. Mean firing rate during the movement epoch of all movement directions for 16 randomly selected RNN neurons.

2.4 – Hand-invariant to hand-specific representation of grasping movements in the macaque fronto-parietal network

Authors:

Jonathan A. Michaels
Hansjörg Scherberger

"Performance of reaching movements appears to require control at multiple levels of abstraction. For example, the neural mechanisms involved in deciding on the target for a reach need not necessarily take into account all the details of muscular contraction which must ultimately be controlled to accomplish the selected movement. Conversely, mechanisms involved in overt muscular control need not be sensitive to the criteria by which a particular action was selected. One therefore expects that different neural populations represent a given movement in different ways, emphasizing some cognitive, temporal, or spatial aspects while ignoring others."

(Cisek et al., 2003)

Abstract

Preparing and executing grasping movements demands the coordination of sensory information across many scales. Factors such as the position of an object, the shape of the hand required to grasp, and which of our hands to use must all be coordinated in parallel. The network formed by macaque anterior intraparietal area (AIP) and hand area (F5) of the ventral premotor cortex is essential in the generation of grasping movements. However, the role of these areas in hand selection for grasping is unclear. We recorded from 1342 single and multi units in areas AIP and F5 of macaque monkeys (*Macaca mulatta*) during a delayed grasping task, in which monkeys were instructed by a context cue (LED) to perform power or precision grips on a handle that was presented in five different orientations, with either the left or right hand, as instructed by an auditory tone. We found that orientation tuning was prevalent in AIP throughout the trial, even in darkness. Units in AIP maintained identical orientation preferences between cue and movement, and were not modulated by the hand used before movement, suggesting that AIP encodes task information in a largely visuo-spatial frame during preparation. In contrast, preparatory activity in F5 increasingly represented the intended hand to be used, but orientation tuning was only present for contralateral movements, revealing a selective transfer of orientation information depending on hand use. Interestingly, preferred grip type was shared between hands, suggesting that the underlying representation of grip information is identical regardless of hand use. Together, our results indicate a more direct form of sensorimotor integration in F5 than previously thought, and a surprising lack of hand dependence in AIP, suggesting that the circuit formed by AIP and F5 is an essential step in the visual to motor specific transformation.

Introduction

Our everyday reaching and grasping movements demand the coordination of information across multiple scales. While grasping a cup requires determination of the physical position and orientation of the cup, one must also resolve the appropriate shaping of the hand, which hand to use, and the muscle forces required. Given this, and given the flexibility with which we switch between hands, it is expected that both hand independent and muscle specific representations should be found at various levels of abstraction throughout cortex.

Indeed, a number of studies have probed how neural circuits represent laterality of reaching movements in macaque monkeys. Integration of arm specific and arm independent information has been found in the parietal reach region (PRR) (Chang et al., 2008), dorsal premotor cortex (PMd) (Hoshi and Tanji, 2000; Cisek et al., 2003; Hoshi and Tanji, 2006; Kurata, 2010), and primary motor cortex (M1) (Muakkassa and Strick, 1979; Tanji et al., 1988; Kermadi et al., 1998), although the outputs from M1 have been identified as mostly contralateral (Penfield and Rasmussen, 1950; Evarts, 1966).

Yet, little is known about the laterality of grasping movements. It has been shown that when all inter-hemispheric connections of macaques has been severed, the ipsilateral hemisphere can generate reaching movements towards food, but cannot properly pre-shape the fingers of the hand (Brinkman and Kuypers, 1973), suggesting that grasping is a highly lateralized process. The hand grasping circuit (Luppino et al., 1999) consisting of the hand area (F5) of the ventral premotor cortex (PMv) and the anterior intraparietal area (AIP) is an essential anatomical and function circuit in grasp preparation and execution. Neural activity in these areas is strongly modulated by visual object properties (Murata et al., 1997; 2000), extrinsic goals (Kakei et al., 2001), performed grip types (Baumann et al., 2009; Fluet et al., 2010), and preparatory activity in these areas can be used to decode the visual properties of objects and complex hand shapes required to grasp a diverse range of objects (Carpaneto et al., 2011; Townsend et al., 2011; Schaffelhofer et al., 2015), as well as predict reaction times (Michaels et al., 2015). Although laterality has been studied in PMv, these studied either employed no delay period (Rizzolatti et al., 1988), simple movements (Tanji et al., 1988), or required only reaching movements (Hoshi and Tanji, 2006; Kurata, 2007; 2010). Additionally, to our knowledge, no studies of laterality have been undertaken in AIP.

In the current study, laterality of grasping movements were investigated using a delayed grasping task (for review see Janssen and Scherberger, 2015) while neural activity was recorded in AIP and F5. Two monkeys visually fixated a central fixation point throughout the trial. During a cue phase, monkeys received a visual cue indicating which of two grip types to perform in one of five possible grasping handle orientations as well as an auditory tone indicating the hand to use on that trial. Following a memory period, a go cue instructed monkeys to grasp the handle in the dark.

We found that activity in AIP and F5 during the movement itself robustly reflected which hand was used, but preparatory activity representing the intended hand was only found in F5, suggesting that AIP represents task information independent of hand during preparation. Furthermore, the amount of grip tuning and preferred grip type of each unit did not depend on hand used, indicating a shared framework for grasp planning. However, although orientation tuning was abundant in AIP, orientation tuning was only present in F5 for contra- and not ipsi-lateral movements, revealing a functional differentiation between hemispheres. Finally, the functional representation of task conditions in the neural population was significantly more correlated between AIP and F5 during contra- rather than ipsi-lateral movement preparation, but equal during execution, suggesting that preparatory coordination between the areas may be limited to contralateral grasping movements.

Results

Behavior

To investigate the laterality of grasp movement coding in premotor and parietal cortex, two monkeys performed a delayed grasping task in which the hand the monkey had to use, as well as the appropriate grip type and hand orientation, was cued on each trial. Concurrently with behavior, single- and multi-unit activity was recorded from premotor area F5 and parietal area AIP simultaneously (Figure 1). Both monkeys successfully performed the task. After initiating trials to the point of obtaining specific trial information, monkeys S and P successfully completed 85% and 84% of trials, respectively. In detail, monkeys S and P correctly selected the correct hand on 89% and 93% of trials, respectively, while grip type selection was correct 99% and 98% of the time. As the motion of the hand during the memory period was tracked with an infrared camera (Materials and Methods), monkeys were required to keep the hands completely still on the hand rests. Trials were completed successfully without premature movement 99% and 94% of the time, for monkeys S and P, respectively. Median reaction time, i.e. the time between the go cue and the hand leaving the handrest, was 230 and 265 ms for monkeys S and P, respectively, while median movement time, i.e. the time between the hand leaving the handrest and executing the appropriate grip on the handle was 305 and 325 ms.

Neural recordings

The analyzed data sets included a collection of 178 individual recording sessions, 91 from monkey S and 87 from monkey P. In monkey S, 861 single- and multi-units were successfully recorded (single: 459, multi: 402), of which 581 were task-related (AIP: 189, F5: 392) and used in further analysis (Materials and Methods). In monkey P, 481 units were recorded (single: 263, multi: 218), of which 390 were task-related (AIP: 207, F5: 183). Units were classified as task-related if they were tuned for any of the three task factors (hand, grip, or orientation) at any point during the course of the trial as determined by a cluster-based permutation test (CBPT; Materials and Methods), which finds contiguous segments of time tuned for one of the three task factors, while keeping the overall false-positive rate below 1.7%. Only units found to be task-related were used in further analysis.
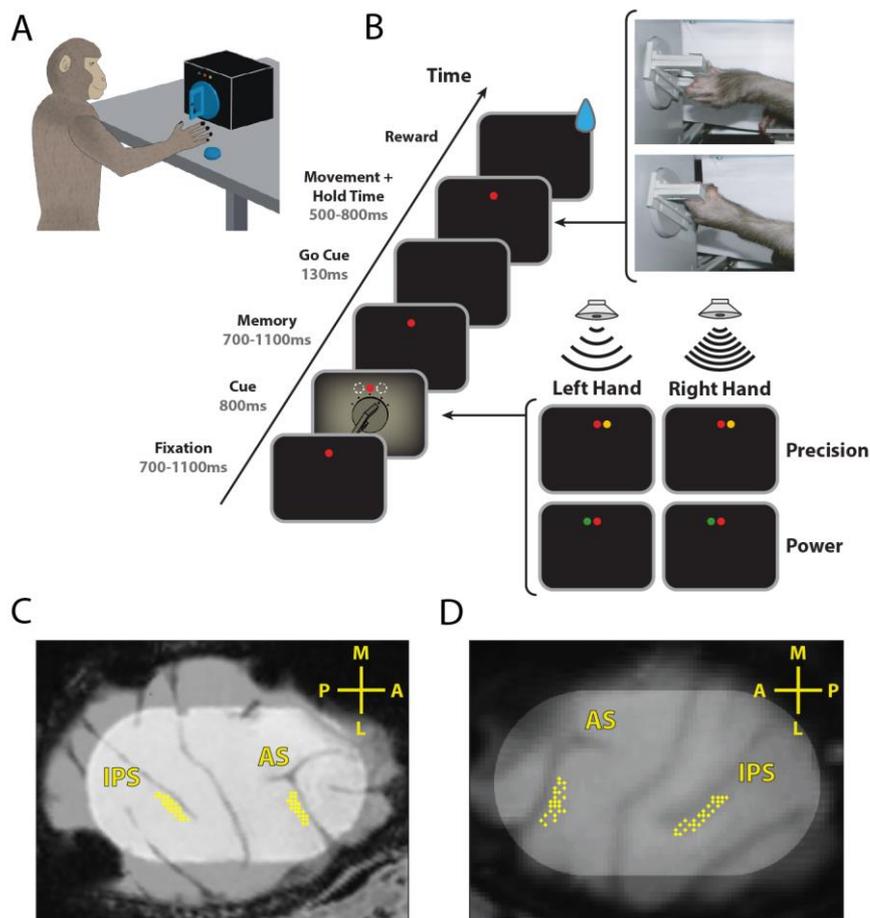
Figure 1. Task design and recordings. *A*, Illustration of a monkey in the experimental setup. The cues were presented via LEDs above the handle. *B*, Delayed grasping task with two grip types (top: power grip, bottom: precision grip), five orientations of the grasping handle, and grasped with either the left or right hand. Grips and orientation were cued using LEDs and handle illumination, while hand was cued by two auditory tones. Trials were presented in pseudorandom order in darkness. *C-D*, Recording locations for monkey P (*C*) and S (*D*) overlaid on a structural MRI. The illuminated oval marks the outline of the recording chamber. Recordings were made in F5 on the bank of the arcuate sulcus (AS) and in AIP toward the lateral end of the intraparietal sulcus (IPS). The cross shows medial (M), lateral (L), anterior (A), and posterior (P) directions. Note that monkey S was implanted in the left hemisphere and monkey P in the right hemisphere.

To get an overview of what kind of task-related responses were present, we averaged over all trials of each condition to produce average firing rate curves and combined them with the significance testing described above (CBPT). Figure 2 shows a number of example single units recorded from both areas and monkeys. One of the most common responses in AIP was tuning for a specific orientation of the handle that was sustained from cue onset to the end of movement, even though the handle was

only illuminated during the cue (Fig. 2 - Top Left). Another common response in AIP was units that did not respond to the cue at all, but showed strong grip and hand tuning specifically during the movement (Fig. 2 – Middle Left). Interestingly, many units in F5 were tuned for the hand used not only during the movement, but also from the end of the cue period onwards, showing a preference for both ipsi- and contra-lateral movements (Fig. 2 – Top Right). Additionally, units showing sustained tuning for grip were widely present (Fig. 2 – Middle Right), and occasionally units that were tuned for all three factors (Fig. 2 – Bottom Right), although these were relatively rare.
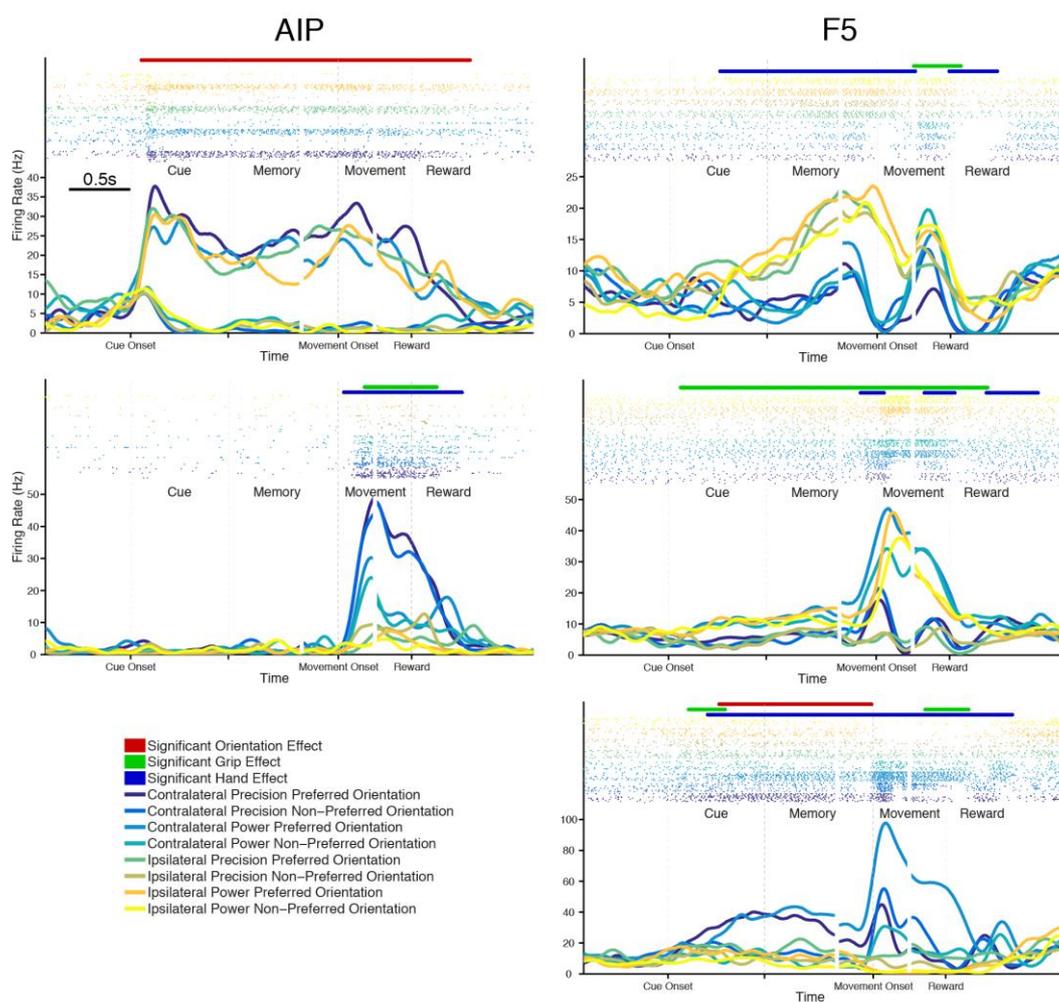
Figure 2. Example average firing rate curves of single-units in AIP and F5. (Top Left) A unit tuned to a single orientation of the handle throughout the trial, even in darkness. (Middle Left) A unit tuned only during movement both for the grip performed and the hand used. (Top Right) A unit tuned for hand used through the trial, showing a preference for ipsilateral movements. (Middle Right) A unit tuned for performed grip throughout the trial, along with phasic hand tuning during movement. (Bottom Right) A unit tuned for all task factors at different points in the trial. Data were aligned to three events, cue onset, movement onset, and reward. Raster plots

above curves show single spikes over all trials of each condition. Significance bars represent tuning for each of the three factors, as determined by cluster-based permutation test ($p < 0.017$). Examples were taken from both monkeys.

Population tuning

While examining the average responses of individual units is an essential step, characterization of population function requires examination of a large set of units. Before looking at the laterality of population responses, we examined population tuning for grip type and orientation of the handle. Figure 3A shows the times of significant grip and orientation tuning for all recorded units from both monkeys and areas, aligned to time of first tuning. Data from both monkeys were pooled together since the amounts of tuning were very similar for both grip and orientation. Tuning increased for both factors typically during the cue epoch and around movement onset. Interestingly, the amount of grip tuning in AIP was quite low, especially during the memory period, while orientation tuning on AIP was extremely prevalent and often sustained for the whole trial, as seen in the example unit in Figure 2. In contrast, grip tuning in F5 ramped steadily during the cue and was often present during the memory period, while orientation tuning was very rare. Also notable is that the preferred grip type of each unit was frequently unstable over the course of the trial, showing a switch in significant tuning at various time points, an interesting effect that will be returned to later.

To see the difference in prevalence in tuning between contra- and ipsi-lateral grasps we repeated the previous tuning analysis for trials of each hand separately and summed over the population, producing plots of the total amount of tuning in Figure 3B. Interestingly, the amount of grip tuning did not differ between contra- and ipsi-lateral trials, suggesting that the intended grip is equally prevalent in both hemispheres during planning and execution. Intriguingly, while there was some orientation tuning in F5 during contralateral grasps, there was virtually none during ipsilateral grasps, implicating a differing function of the area depending on the hand being used for grasping.
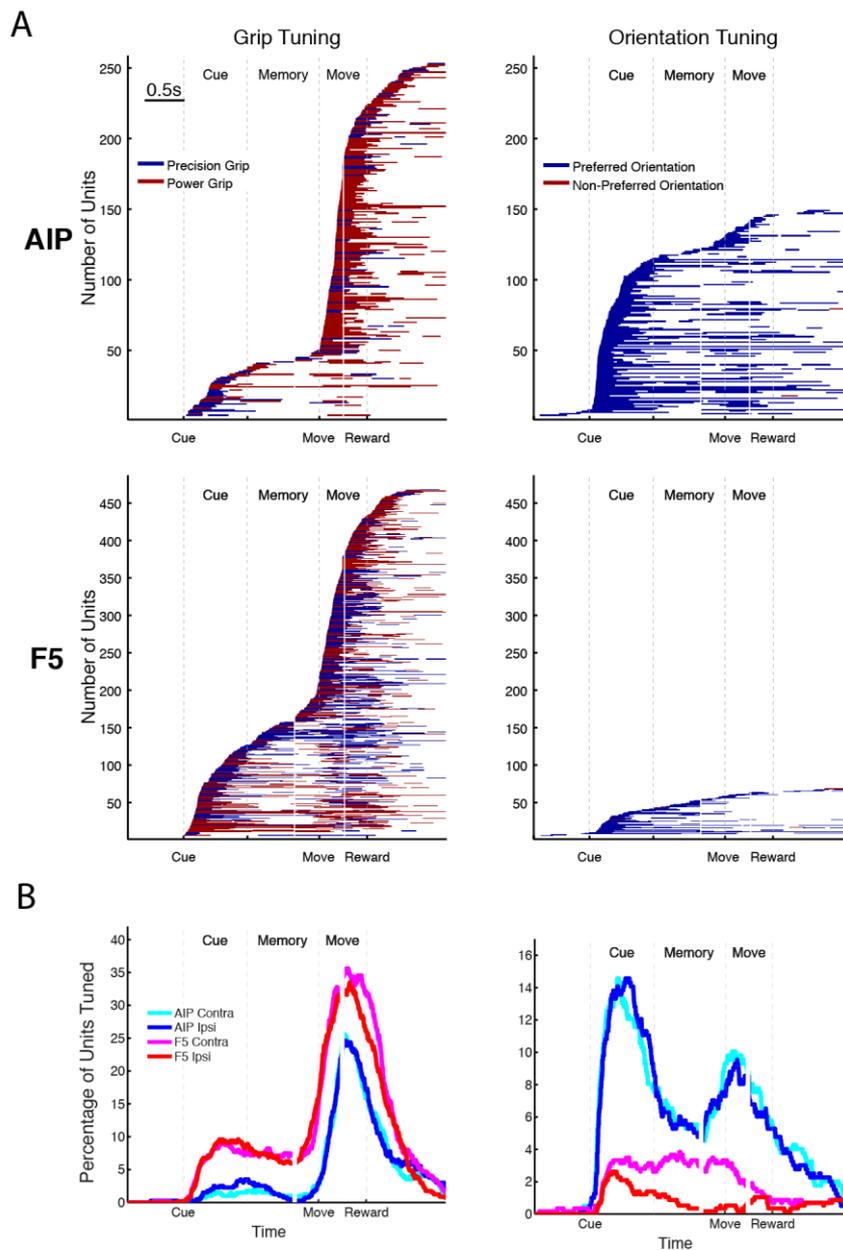
Figure 3. **Grip type and orientation tuning over all recordings.** *A*, Times of significant tuning for all units from both areas and monkeys, aligned by onset of first tuning, as determined by cluster-based permutation test ($p < 0.017$). Labeled condition represents the one that elicited the higher firing rate. *B*, Percentage of units tuned for each effect over time. Same analysis as *A* repeated for contra- and ipsi-lateral trials separately and collapsed for each time point. Results were very similar between monkeys and were therefore pooled. Data were aligned to three events, cue onset, movement onset, and reward.

Another interesting question is whether or not the preferred grip type was shared between contra- and ipsi-lateral grasps. To test this, we compared the preferred grip type (highest firing) between trials of each hand for each unit that was

significantly tuned (based on CBPT) during both movements. During cue and memory, not a single unit in either area switched preference between contra- and ipsilateral trials, indicating that grip tuning was preserved regardless of hand used. During the movement period, less than 5% of units in either area switched their preference between trials of each hand, implicating a shared representation regardless of hand used.

Laterality encoding

Up to this point we have only considered tuning for classical features such as grip type and handle orientation. Now, we consider tuning for the hand itself using the same analysis as above, but plotted for each monkey separately (Figure 4). For monkey S, hand tuning was virtually non-existent in AIP before movement onset, suggesting that AIP encodes task-relevant features in a hand-independent manner before the movement has started. In contrast, hand tuning in F5 seemed to ramp continuously throughout the entire trial, reaching a maximum (50% of units tuned) just before the hold period, showing that F5 is strongly dependent on hand used.

The results of monkey P with respect to hand tuning were significantly different, showing phasic spikes in hand tuning in both areas immediately after the cue. As described in the Materials and Methods, an infrared camera tracked the position of the hands during the cue and memory periods to ensure that no premature movements occurred. However, very small movements were likely still possible. To test this possibility, the sum of a measure of hand movement (Materials and Methods) that was recorded during the memory period of individual trials was used to decode task conditions offline. Using a linear classifier (n-fold cross-validated, Matlab function: *fitcdiscr*), it was possible to decode the hand used from the memory period with 52% accuracy in monkey S, where 50% is chance level, suggesting that no premature movements occurred. In contrast, the hand used could be decoded with 75% accuracy in monkey P, suggesting that small premature movements may have occurred. Additionally, grip type could never be decoded from the infrared video during the memory period (50% accuracy in both monkeys), indicating that any premature movements never represented the grip type. Taken together, this result suggests that monkey P made small premature movements during the memory that caused phasic spikes in hand tuning, but left grip tuning unaffected. Therefore, we

believe that the results of monkey S with respect to hand tuning during the cue and memory periods conservatively reflects the function of these areas.
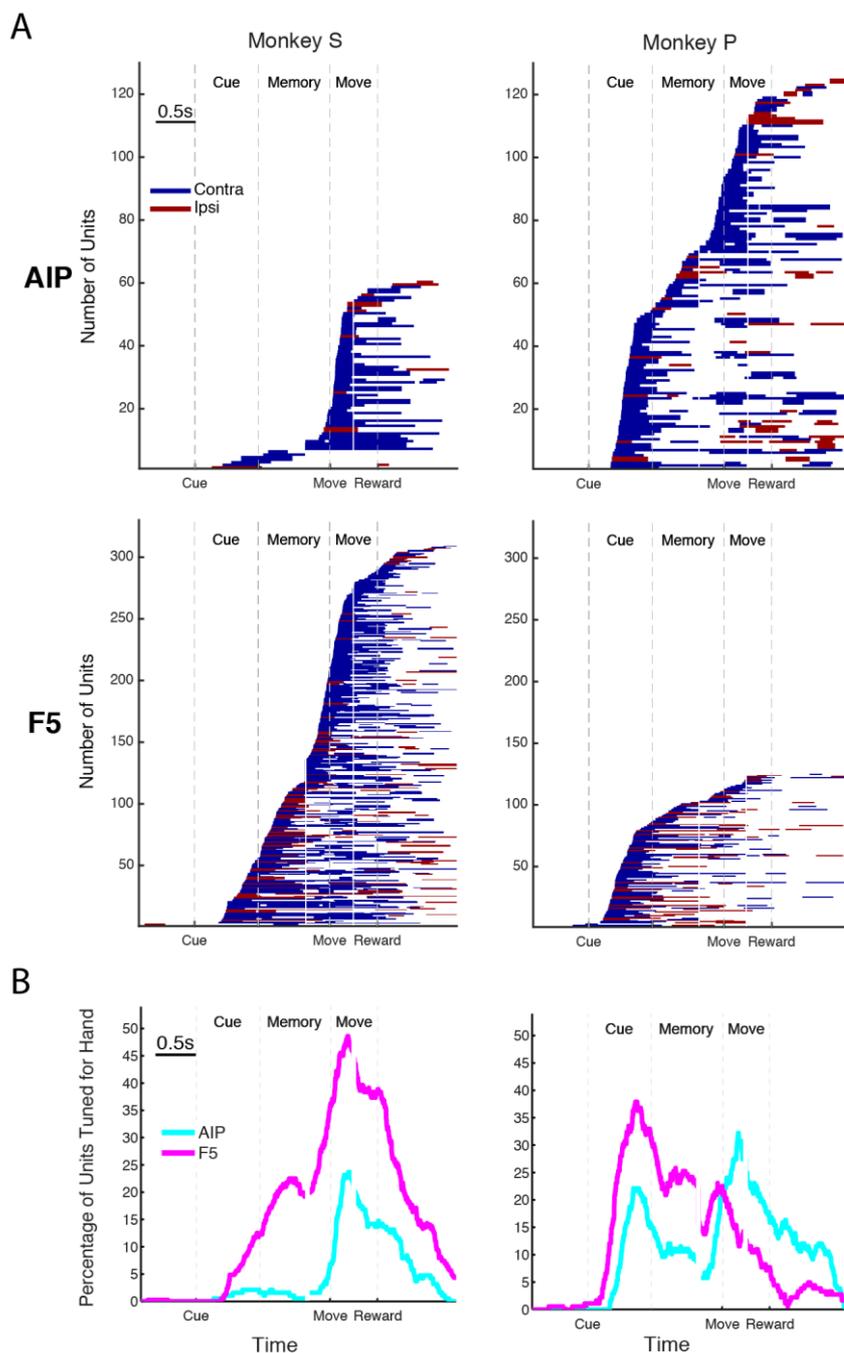
Figure 4. Hand tuning over all recordings separated by monkey. *A*, Times of significant hand tuning for all units from both areas and monkeys, aligned by onset of first tuning, as determined by cluster-based permutation test (*p* < 0.017). Labeled condition represents the one that elicited the higher firing rate. *B*, Percentage of units tuned for hand over time, obtained from *A*. Results significantly differed between monkeys and were therefore presented separately. Data were aligned to three events, cue onset, movement onset, and reward.

Now that we have seen that hand tuning exists in both areas, Figure 5 plots the hand specificity index (HSI, Materials and Methods) of all units of monkey S to identify how much firing rates differed between contra- and ipsi-lateral grasps. A value of one indicates firing purely during contralateral movements, while a value of negative one indicates firing purely during ipsilateral movements. As expected, during the cue and memory period the HSI was not significantly shifted from zero in AIP (Wilcoxon signed rank test, $p > 0.05$), while the HSI was significantly biased towards contralateral during movement and the hold period (Wilcoxon signed rank test, $p < 0.001$). Furthermore, this shift did not differ when splitting units into single- and multi-unit recordings. In F5, the HSI was significantly shifted in the positive direction during all epochs (Wilcoxon signed rank test, $p < 0.001$), suggesting that producing contralateral movements generally increased firing rates more than ipsilateral movements. However, during memory it is clear that units were tuned to both contra- and ipsi-lateral movements, as seen by the bimodal distribution of significantly tuned units, suggesting that movements of either hand were likely to produce higher firing rates during the memory period.
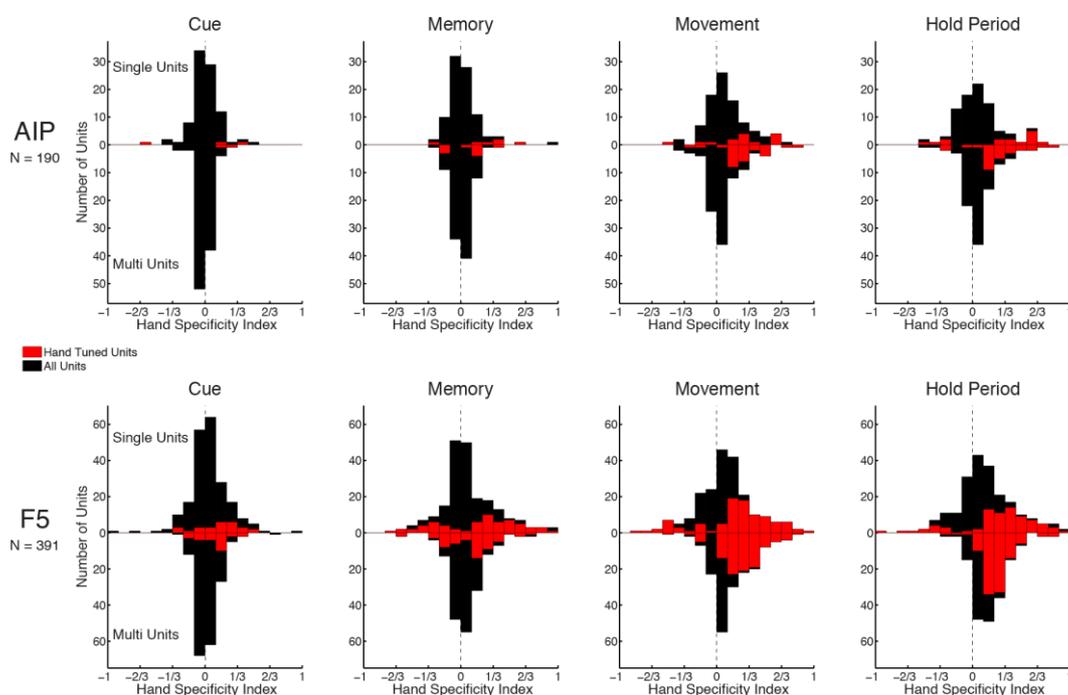

Figure 5. Hand specificity index (HSI) for monkey S. HSI was calculated as the sum of the spikes in a specific task epoch for contralateral trials, minus the same for ipsilateral trials and then divided by the sum of both values. Therefore, a value of +1 indicated purely contralateral firing while a value of -1 indicated purely ipsilateral

firing. Units were plotted separately for single- and multi-units to determine if a bias was introduced by pooling activity into multi-units. This analysis was only carried out on monkey S since hand tuning was significantly different between monkeys.

Tuning consistency

A large number of units were tuned during each epoch of the task. How often was this tuning sustained between epochs, and was there significant preference for particular conditions over time? Figure 6A shows a chart of which epochs all units in both areas were tuned for either grip or hand. For grip tuning, most units were only tuned during one epoch, although there was substantial overlap in many cases. Hand tuning was least during the cue, most during movement, and showed substantial overlap between memory and movement. However, most units were still only tuned during one epoch, suggesting that most of the time tuning to either factor is phasic and not sustained over all epochs.
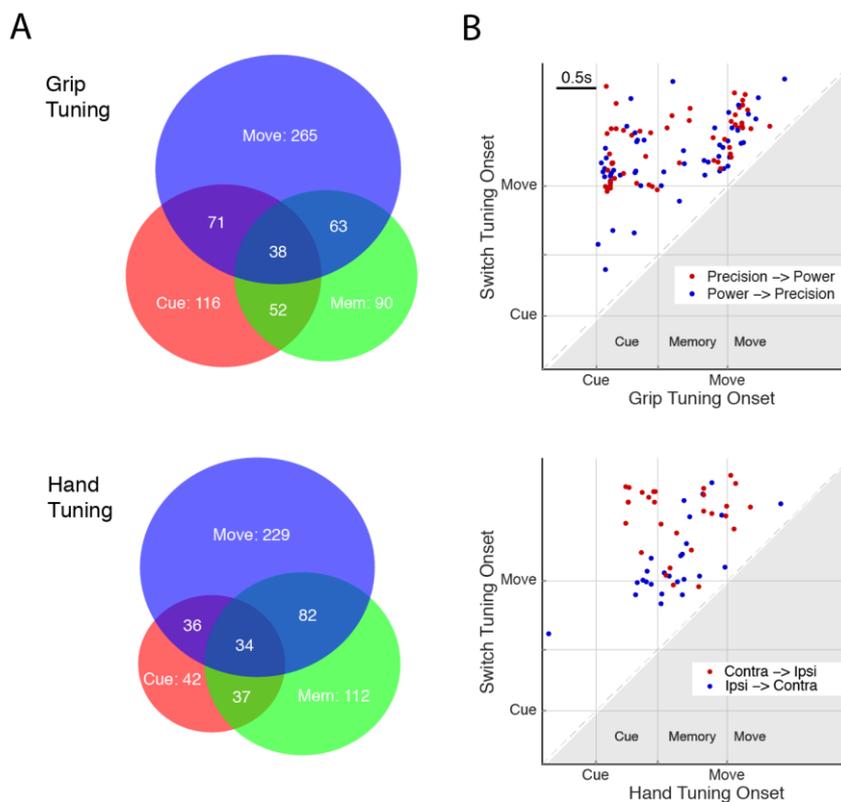
**Figure 6. Consistency of grip and hand tuning in F5.** *A*, (Top) Distribution of when units (from both monkeys) were tuned for grip in F5. (Bottom) Distribution of when units (from monkey S only) were tuned for hand in F5. Although some units remained tuned over the entire trial, many only showed a preference during specific epochs. *B*, (Top) Onset of first tuning and switch tuning for units that switched their preferred

grip type during the trial. (Bottom) Onset of first tuning and switch tuning for units that switched their preferred hand during the trial. All tuning is based on the cluster-based permutation test used in Figures 3-4.

To investigate the consistency of tuning over time, Figure 6B shows the times when units first became tuned for grip or hand, and the time at which certain units became significantly tuned for the other factor. Only units that switched tuning are plotted. Interestingly, very few units changed grip preference during the cue or memory period, yet many either changed their preference between cue and movement, or throughout the movement itself. Overall, 23% of units switched their tuning at some point during the trial, suggesting that preference varied throughout time. A similar effect was seen for hand tuning, since 18% of units switched their hand preference at some point during the trial, however almost exclusively during the movement itself.

Coordinate frame

Both areas investigated in the current study are essential parts of the visuo-motor transformation process, and therefore involved in transforming visual information into a body-centered coordinate frame so that muscle movements can be executed to the appropriate location in physical space. By examining the relationship between handle orientation preference and hand used, it is possible to test the representation of extrinsic (visual-centric) and intrinsic (body-centric) coordinate frames in all recorded units. Figure 7A illustrates the two potential preferences of an example unit for the two coordinate frames. If a unit were tuned purely in an extrinsic frame, orientation preference should be shared regardless of hand used. In contrast, in an intrinsic frame the orientation preference may shift between hands to match the correspondingly mirrored wrist rotation (i.e., pronation or supination) required between each hand. To test this, we correlated the average firing rate of each unit during all handle orientations of one hand and the other over time (200 ms spiking windows in steps of 50 ms). If the correlation exceeded zero and was significant, that indicated an extrinsic frame, as shown in Figure 7B, while units showing a significant correlation below zero indicated an intrinsic frame (permutation test, $p < 0.05$). Figure 7C shows the number of units fitting either frame over both monkeys for AIP and F5 separately. The number of units fitting the extrinsic frame was highest at the end of the cue and

decreased throughout the trial in both areas, increasing slightly in AIP during movement. Interestingly, very few units showed an intrinsic frame, suggesting that such a transformation is not represented in the units of these areas. However, the general lack of orientation tuning in F5 makes conclusions difficult and is discussed later.
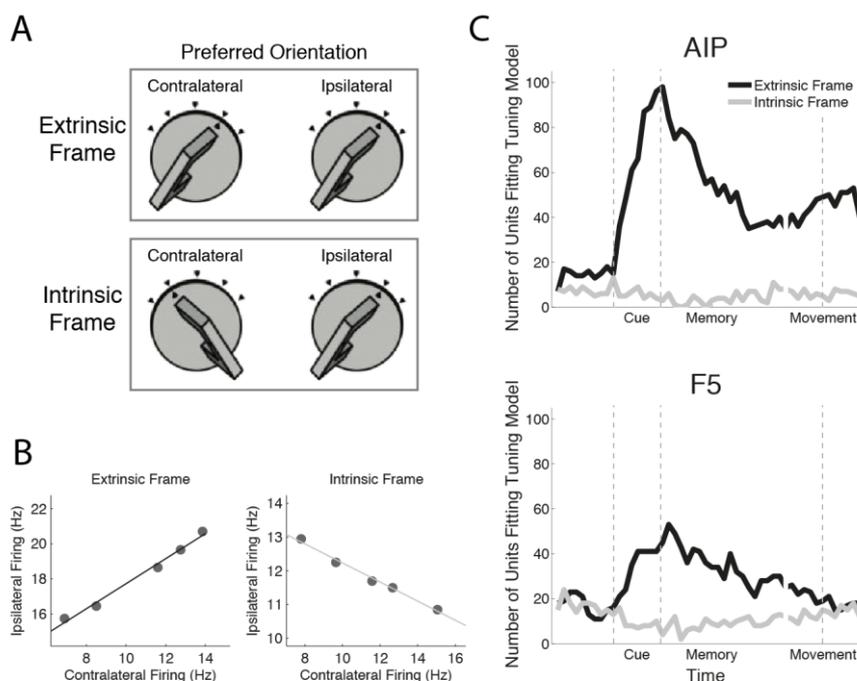
Figure 7. Extrinsic or intrinsic reference frame representation in individual units. *A*, Illustration of preferred handle orientations for example units following an extrinsic (visual) or an intrinsic (body-centered) frame. *B*, Average firing rates of example units. Each dot represents the average firing rate for one of the five handle orientations, with the line of best fit. A positive slope for the line of best fit corresponded to an extrinsic frame and a negative slope an intrinsic frame. *C*, Number of units fitting either model in AIP and F5 (200 ms spiking windows in steps of 50 ms).

Demixed principal component analysis

Both AIP and F5 are involved in the processing of a large multitude of task factors. These factors must be processed in parallel, and are distributed over many units in the population. To get a picture of how these factors are represented in the populations of AIP and F5, we implement demixed principal component analysis (dPCA), a dimensionality reduction methods for extracting low-dimensional linear combinations of a population that represent specific task features (Kobak et al., 2014). Since there was a significant difference in how hand was encoded between the two monkeys, we

developed an additional procedure for extracting only dimensions that were present independently in both monkeys (Materials and Methods). Figure 8 shows the results of dPCA over all units of both monkeys in AIP and F5. The amount of variance explained by dPCA was very similar to that extracted by classical PCA, suggesting that the method is able to properly extract relevant task features.

Intriguingly, the components taking up the most variance overall are condition independent signals. The first component in AIP is a large condition-independent signal that begins shortly before movement onset, while the second component is a large phasic response to the cue. In F5, the first component is the same, while the second component seems to progress linearly with time, especially during the memory period. The next largest component in both areas was related to the hand used, although the hand used could only be decoded from this signal in AIP during the movement period, as denoted by the black bars (Materials and Methods), while hand used could be consistently decoded from this signal in F5 starting towards the end of the cue and lasting throughout memory and movement. The same was true of grip decoding, being present in F5 throughout the trial starting in early cue, but only during movement in AIP. However, grip could be decoded at earlier time points in component 10 in AIP, as shown by the black bars marking significant decoding. Handle orientation was very well decodable from AIP throughout the trial, while barely so in F5.

Given that the task factors must be represented within each area in parallel, we tested the orthogonality of each component by taking the dot product between the coefficients of each component (i.e. the contributions of all units to each component). Interestingly, most components were not significantly non-orthogonal, suggesting that the task factors are independently represented.
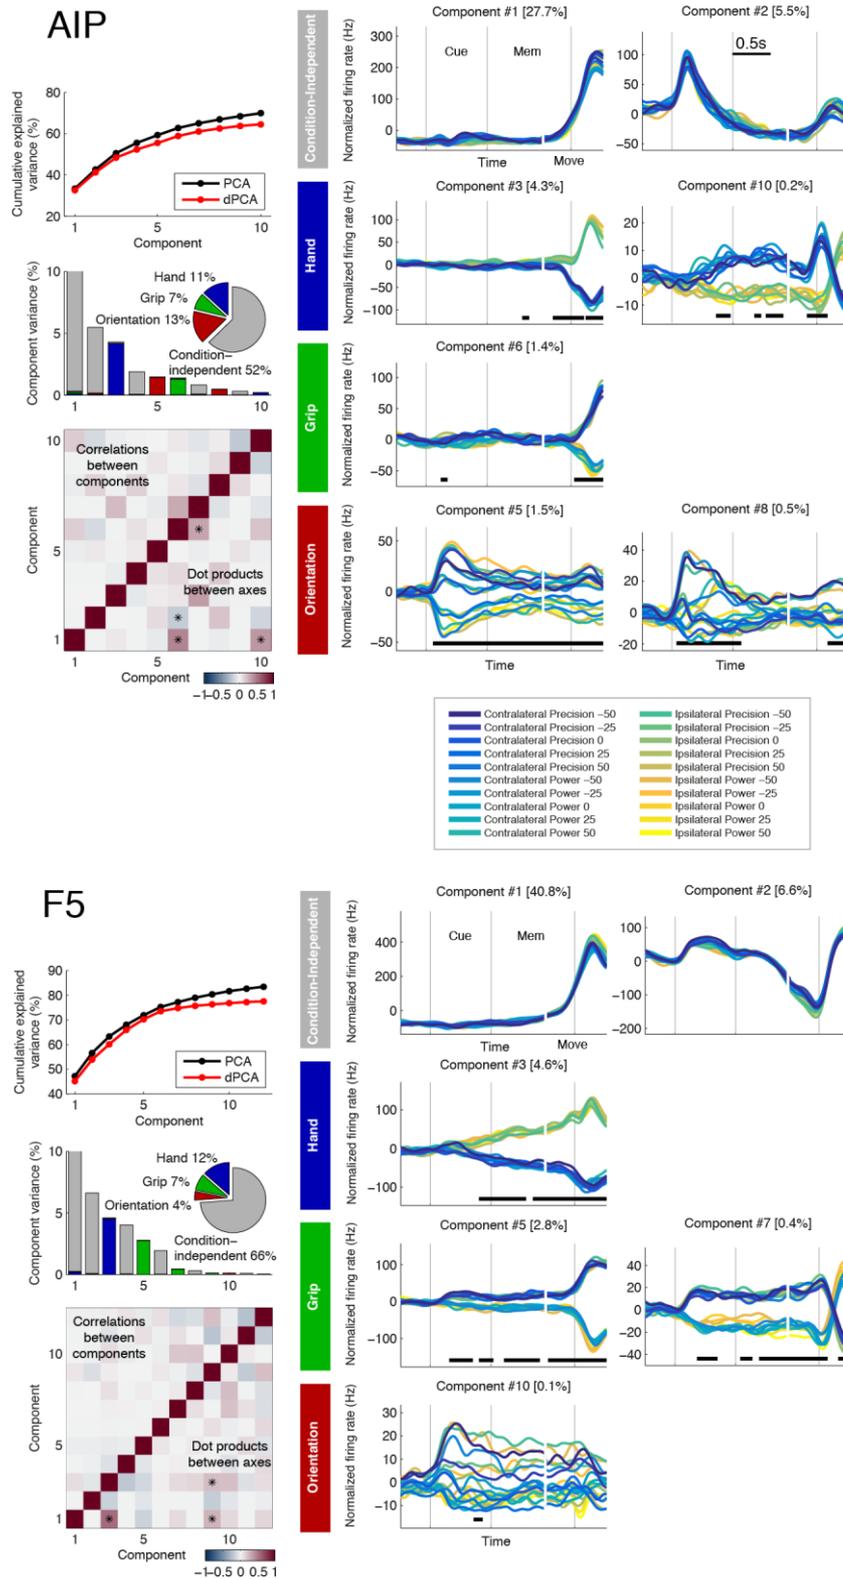
**Figure 8. Demixed PCA on AIP and F5 of both monkeys.** (Right Panels) Demixed principal components. Thick black lines show time intervals during which the respective task parameters can be reliably extracted from single-trial activity. Note that the vertical scale differs across subplots. (Left Top) Cumulative signal variance explained by PCA (black) and dPCA (red). Demixed PCA explains almost the same

amount of variance as standard PCA. (Left Middle) Variance of the individual demixed principal components. Each bar shows the proportion of total variance, and appears to be single-colored, which signifies nearly perfect demixing. Pie chart shows how the total signal variance is split between parameters. (Left Bottom) Lower-right triangle shows dot products between all pairs of demixed principal axes. Stars mark the pairs that are significantly and robustly non-orthogonal. Upper-left triangle shows correlations between all pairs of the components. Most of the correlations are very low.

Task-space correlation

Communication between AIP and F5 is required to properly prepare and execute grasping movements. Therefore, we wanted to test how functionally similar the representation of our task was in both areas at different times during the trial. To achieve this, we calculated the Euclidean distance between the average firing rates of all pairs of conditions (for each hand separately) in the space of all recorded units across both monkeys, producing distance matrices. Two example matrices are shown in Figure 9A. The unique portions (upper triangle excluding the diagonal) of these matrices were then correlated between AIP and F5 for trials of each hand separately. The resulting correlations (Spearman) are shown in Figure 9B for five epochs, and significant differences between contra- and ipsi-lateral trials are marked with a star (permutation test, $p < 0.05$, false-discovery rate correction, Materials and Methods). A high correlation would indicate that the relationship between task conditions in neural space was very similar between areas, while a low correlation would indicate differing representations of task conditions between areas. At the beginning of the cue, the correlation was reasonable and not significantly different between contra- and ipsi-lateral movements. However, by the end of the cue and during the memory period the correlation between areas was significantly higher during contralateral movements, suggesting that the representation of the task was highly similar in AIP and F5 specifically for these movements. Finally, during the movement the task representation was very similar between areas and did not differ between contra- and ipsi-lateral movements. If data are separated by monkey, the same result is obtained for monkey S, but not for monkey P, a relatively expected result given the differences in memory period behavior between monkeys. Although not a direct measure of causal communication, this analysis strongly suggests that communication between AIP and F5 is more extensive during preparation of contra- vs. ipsi-lateral grasping movements.
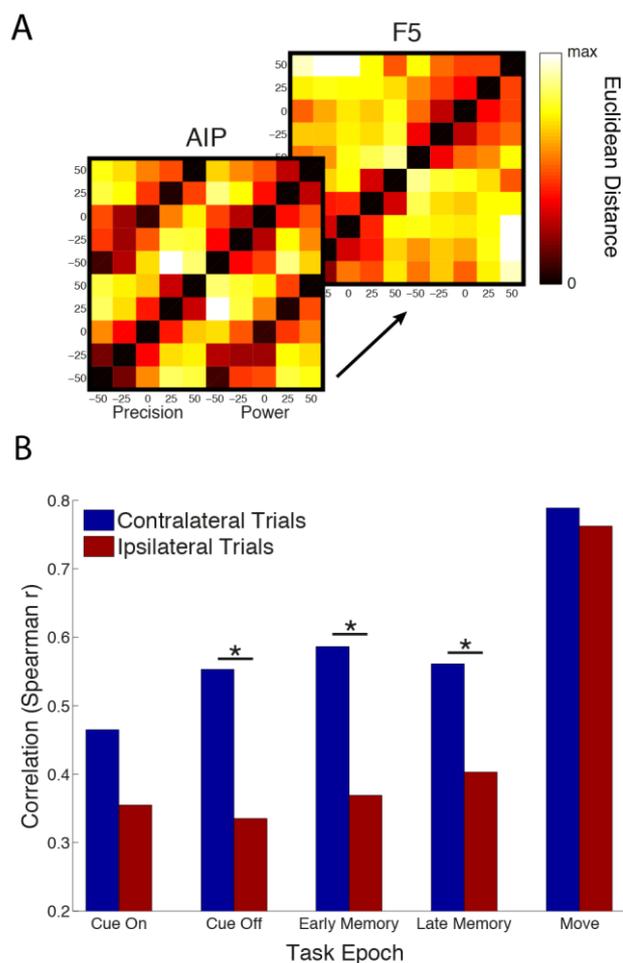
Figure 9. Task space correlation between contra- and ipsi-lateral movements. *A*, Example distance matrices showing Euclidean distance between all pairs of conditions for one hand (contralateral) in the neural space of all recorded units of both monkeys in the early memory epoch. Distance matrices produced from the activity of each area were correlated to each other to determine the similarity in functional representation between areas. *B*, Spearman correlation of distance matrices for contra- and ipsi-lateral trials separately. Stars mark epochs when the correlation was significantly higher during contra-lateral trials (permutation test, *p* < 0.05, False-discovery rate correction).

Discussion

In the current study, by recording from many units in AIP and F5 of two macaque monkeys during a delayed grasping task, we found that the laterality of hand use was not robustly encoded in AIP before movement, while hand tuning steadily increased in F5. Furthermore, while orientation tuning was prevalent in AIP, it was only found in F5 during preparation of contralateral movements.

The amount of grip tuning obtained in AIP was significantly lower than found in previous studies (Baumann et al., 2009) as was the amount of orientation tuning found in F5 (Fluet et al., 2010). We believe these differences are due to selective recording of task-related units in previous studies, while in the current study units were not evaluated for tuning online, presumably giving a more unbiased estimate of tuning percentage.

As mentioned earlier, severing all connections between hemispheres of macaques has shown that the ipsilateral hemisphere can coordinate reaches, but not properly pre-shape the fingers of the hand to grasp the food (Brinkman and Kuypers, 1973). Therefore, it's interesting that in the current study grip preference and tuning was identical regardless of hand use. It suggests that premotor cortex in either hemisphere should be successful in hand shaping; yet this was not the case in this classic study. Therefore, it is unlikely that involvement of F5 in ipsilateral movements is contributing directly to muscle activation. Supporting this, stimulation of ipsilateral M1 produces no direct corticospinal activation of the muscles (Soteropoulos et al., 2011). Therefore, F5 modulation during ipsilateral movements is likely coordinated with the contralateral F5 through the corpus callosum and likely also plays a larger role during bimanual rather than unimanual grasping movements, as is the case in M1 (Donchin et al., 1998).

Interestingly, while grip preference did not change between hands, many units changed their grip or hand preference between the preparatory periods and the movement itself (Fig. 6), in line with studies showing that individual unit tuning tends to be unstable, and that different dynamics govern preparation and movement (Churchland et al., 2012; Ames et al., 2014).

The fact that AIP showed no preparatory response to hand use is unexpected, especially since the nearby parietal reach region shows strong modulation (Chang et al., 2008). Additionally, AIP is part of the network that responds to passive auditory listening (Poremba et al., 2003), and since the current task employed an auditory cue,

it would be expected that playing multiple tones would elicit a task dependent response. It could be that the auditory stimuli in our task were not varied enough to elicit a significant effect, or rather that since the task was active rather than passive, AIP was likely dominated by visual processing demands.

During the movement itself, both areas strongly represented the hand used. Along with F5 contributing to the initiation of movement, hand tuning during movement could originate from projections from secondary sensory cortex to both F5 (Kurata, 1991; Gerbella et al., 2011) and AIP (Borra et al., 2008).

Monkeys were required to reach to the target as well as grasp. Therefore, reach planning and execution is likely a significant part of the observed activity. However, as we have argued previously (Michaels et al., 2015), previous research employing a grasp-only task (Hepp-Reymond et al., 1994) and a grasp-reach dissociation task (Lehmann and Scherberger, 2013) suggests that F5 encodes grasping quite independently of reaching, although both areas contain information about reach position. Furthermore, reversibly inactivating F5 (Fogassi et al., 2001) or AIP (Gallese et al., 1994) selectively impairs hand-shaping and not reaching, suggesting that our results are an accurate representation of the grasping network.

Although monkeys received grip cue information at the center of their visual field, the effector cue was auditory, introducing a potential confound in lateralized processing. However, it is unlikely that any lateralization effects found in the current study are a result of asymmetric processing of auditory information since only complex stimuli, such as vocalizations, evoke a lateralized response in macaque monkeys (Poremba and Mishkin, 2007; Joly et al., 2012).

Based on our analysis of infrared motion tracking, it is very likely that monkey P made small movements of the hand during the preparatory phases, biasing hand tuning during that time. However, the same analysis showed that no such movements occurred with monkey S, and grip and orientation tuning appeared unaffected. Furthermore, we were able to extract the population level preparatory and movement related signals that were shared in both monkeys (Fig. 8), revealing the commonality in data sets. The relative orthogonality of these extracted dimensions is an indicator of the coding mechanisms that may be employed by cortex in order to efficiently encode information in parallel. A number of studies support the notion that low-dimensional features of neural populations have a biological basis, including learning induced

plasticity (Sadtler et al., 2014), gating of motor output (Kaufman et al., 2014), and parallel encoding (Mante et al., 2013).

In our analysis of orientation tuning, the lack of intrinsic coordinate frame tuning observed across both areas (Fig. 7) raises two possibilities. Either, wrist orientation is not well encoded in either area, or an intrinsic frame in which wrist orientation preference is mirrored in individual units between hands used is not a sensible coordinate frame employed by these areas. Given the low amounts of orientation tuning observed in F5, it is likely that wrist orientation is simply not directly encoded in F5, and that properly testing an intrinsic frame requires considering more complex reference frames that well explain wrist and finger movements.

Given that the hand tuning observed in F5 does not likely originate from AIP, the hand tuning in F5 likely comes indirectly through prefrontal cortex, from which a number of areas project to F5 (Gerbella et al., 2011). This is also in line with the fact that hand tuning in F5 appears only towards the end of the cue and builds slowly, as observed previously (Hoshi and Tanji, 2006). An alternative explanation for large amounts of hand-invariant tuning could be that many proximal muscles are required for movements of either arm, given the large postural adjustments required in extending the arm, and could therefore explain large amounts of hand invariant tuning. However, experiments limiting movements to the distal muscles alone (Tanji et al., 1988), or controlling for postural contributions to ipsilateral control (Ganguly et al., 2009), have shown strongly bilateral representation of hand movement in premotor cortex, suggesting that postural control cannot fully explain hand-independence.

On the other hand, an extrinsic frame representation in PMv is not altogether surprising, since PMv has been shown to be very sensitive to visuo-spatial information as opposed to the dynamics of movement (Hoshi and Tanji, 2006; Xiao et al., 2006), but likely shifts its control strategy during movement (Suminski et al., 2015). A change in control strategy between preparation and movement is supported by the fact that although there was an equal amount of orientation tuning in F5 during movement as during cue, the number of units fitting an extrinsic frame dropped to baseline levels, suggesting that the reference frame used during execution in F5 does not match either tested frame. In contrast, almost all units in AIP tuned for orientation during movement also showed a significant extrinsic frame, suggesting that AIP

continues employing a largely visuo-spatial frame. Therefore, PMv is likely a site of transformation between hand-invarient and hand-specific representations, representing stimuli from both the contra- and ipsi-lateral visual hemi-field (Boussaoud and Wise, 1993a; 1993b), leading to a strong visuo-spatial dependence, while also accumulating hand-specific information before and during movement. Yet, the dependence of orientation tuning in PMv on hand use reveals a more direct representation of sensorimotor integration than posited previously (Hoshi and Tanji, 2007), and provides an interesting perspective on the function of the premotor and parietal grasping circuit.

Materials and Methods

Experimental Setup

Two female rhesus monkeys (*Macaca mulatta*) participated in this study (monkeys P and S; weight 4.5 and 5.5 kg, respectively). They were pair-housed in a spacious and enriched environment. All procedures and monkey care were conducted in accordance with the regulations set by the *Guidelines for the Care and Use of Mammals in Neuroscience and Behavioral Research* (National Research Council, 2003), and in agreement with German and European laws governing monkey care.

Monkeys were habituated to comfortably sit upright in an individually adjusted primate chair with the head rigidly fixed to the chair. A grasp target was located at a distance of 24 cm in front of the monkey. The target consisted of a handle that could be grasped with two different grip types, either with a precision grip (using index finger and thumb in opposition) or a whole-hand power grip (Baumann et al., 2009; Fluet et al., 2010). Inside the handle, two touch sensors were placed in small, visible recessions to detect the contact of the monkey's thumb and index finger during precision grips. An infrared light barrier placed inside the opening of the handle detected power grips. Grip type was instructed by two colored light emitting diodes (LEDs) that were positioned immediately above the grasping handle. The handle was rotatable and was presented in five different orientations (upright and 25° or 50° clockwise and counter-clockwise) and two spotlights could illuminate it from the left and right side in an otherwise dark experimental room. Two capacitive touch sensors (model EC3016NPAPL; Carlo Gavazzi) were placed at the level of the monkey's waist as handrest buttons. A single speaker, which produced the audio tones for cuing the appropriate arm, was positioned directly above and behind the monkey's head. The speaker was oriented such that the audio tone was equally directed into each ear. Monkeys had to fixate on a red LED that was positioned between the two cue LEDs. Eye movements were measured using an optical eye tracker (ET-49B; Thomas Recording) and custom-written software implemented in LabView Realtime (National Instruments) using a time resolution of 5 ms was used to control the behavioral task.

In additional to normal behavioral control, the stationarity of each monkey's hands on the hand rests was also tracked during the memory period of every trial with an infrared camera positioned directly over the hands. Using a separate custom-written LabView control program, the stationarity of both hands was simultaneously monitored for several criteria: (a) the total luminance of the hand, (b) the center of the

hand, i.e. the position of the weighted average of the most luminous pixels in both the x and y direction separately, and (c) the standard deviation around the center of the hand in both the x and y direction. If at any time during the memory period either hand moved more than a pre-set threshold with respect to any of these factors, the trial was aborted without reward. The thresholds beyond which a trial would be aborted were fixed during all recordings of both monkeys at: (a) ±2%, (b) ±1%, and (c) ±2%.

Additionally, during all recordings of monkey S and a small portion of the recordings from monkey P, continuous infrared hand motion information was digitally stored (500 Hz) for later offline analysis. For each hand the sum of these variables was recorded, i.e. the sum of the total hand luminance, the center of the hand in x-coordinates, and the center of the hand in y-coordinates, as these values were most valuable in controlling motion behavior.

Behavioral paradigm

Monkeys were trained in a delayed grasping task in which they were required to grasp a handle in five possible orientations with either a power grip or a precision grip using the left or right arm. This led to 20 grasp conditions that were presented in a pseudorandom order. To initiate a trial, monkeys sat in darkness and placed each hand on a handrest button. The handle was then positioned in one of the five orientations and subsequently a red fixation LED switched on. From then on, the monkey was required to fixate while keeping both hands on the handrest buttons (fixation period duration: 700–1100 ms, mean: 900 ms), as illustrated in Figure 1A. In the following cue period (cue period duration: 800 ms), the object was illuminated to reveal its orientation. The color of an additional LED presented to the left or right of the fixation LED indicated which grip type to perform, either a power grip (green light, left) or a precision grip (yellow light, right). In addition, an audio tone (1000Hz or 2000Hz), representing the left and right arms, respectively, was presented simultaneously with the grip cue and spotlights. The spotlights, audio tone, and the grip cue LED were then switched off while the fixation light remained on for a variable period (memory period duration: 700–1100 ms, mean: 900 ms) during which the monkey had to remember the trial instructions. A brief blinking of the fixation LED (130 ms) instructed the monkey to reach and grasp the object in the dark with the correct arm while maintaining eye fixation and keeping the other arm on the

handrest. After a hold period of 300 ms, each correct trial was rewarded with a fixed amount of water.

Surgical procedures and MRI scans

Details of the surgical procedures and MRI scans have been described previously (Lehmann and Scherberger, 2013). In short, a titanium head post was secured in a dental acrylic head cap and a custom made oval-shaped recording chamber [material PEEK (polyether ether ketone); outer dimensions, 40 Xئ 25 mm^2; inner dimensions, 35 X 20 mm^2] was implanted over the right or left hemisphere to provide access to parietal area AIP and premotor area F5.

Two structural magnetic resonance image (MRI) scans of the brain and skull were obtained from each monkey, one before the surgical procedures, to help guide the chamber placement, and one after chamber implantation to register the coordinates of the chamber with the cortical structures (Fig. 1C,D). AIP was defined as the rostral part of the lateral bank of parietal sulcus (Borra et al., 2008), whereas recordings in F5 were made primarily in F5ap, which is in the post-arcuate bank lateral to the tip of the principal sulcus (Belmalih et al., 2009).

Neuronal recordings

Single-unit and multi-unit (spiking) activity was recorded using quartz-glass-coated platinum/tungsten single electrodes (impedance 1–2 MΩ at 1 kHz) or tetrodes (impedance 500-800 kΩ at 1 kHz) that were positioned simultaneously in AIP and F5 by two five-channel micromanipulators (Mini-Matrix, Thomas Recording). Neural signals were amplified (400X), digitized with 16-bit resolution at 30kS/s using a Cerebus Neural Signal processor (Blackrock Microsystems), and stored on a hard drive together with the behavioral data.

Preprocessing

All data analysis was performed offline. Neural signals were band-pass filtered (forward-backward) with cutoff frequencies between 300-5000 Hz. Waveforms were extracted when the signal deflected beyond 5 standard deviations from baseline either negatively or positively. The refractory period between spikes was set at 1.5 ms. During tetrode recordings spikes that were detected on one of the electrode tips were

extracted from all 4 and aligned to the peak or valley of the first channel to cross the threshold. Units were isolated using principal component analysis techniques (Offline Sorter v3.2.2, Plexon), and sorted into single and multi units. Using Matlab (Mathworks) for further analysis, we included all units in our database that were stably recorded for at least 5 trials per condition (100–260 trials in total). Average firing rate curves were generated using a Gaussian window as a kernel (SD: 57 ms) in three alignments (cue, movement, and reward). However, all statistical tests other than the cluster-based permutation test were based on exact spike counts.

Data analysis

The preferred and non-preferred orientations were determined for each unit from the mean activity in the time interval from cue onset to reward onset. Activity was averaged across all trials of the same orientation. Of the five tested orientations, the orientation with the higher (or lower) mean firing rate was defined as the preferred (or non-preferred) orientation, as in Baumann et al. (2009).   In order to determine the degree to which each individual unit preferred, i.e. increased firing rate for, movements with the contra- or ipsi-lateral hand, a hand specificity index (HSI) was calculated for each unit,

$$HSI = \frac{Contra - Ipsi}{Contra + Ipsi}$$

where 'Contra' represents average spiking of a unit during a specific task epoch for all trials where the contralateral hand was cued, and vice versa for 'Ipsi'. Therefore, a value of +1 would indicate activation purely for contralateral trials and -1 for purely ipsilateral trials (Fig. 5).

To complement each firing rate curve, periods of significant tuning were calculated using a cluster-based permutation test (CBPT) to generate the significance bars in Figure 2 (Maris and Oostenveld, 2007). Briefly, this test evaluates the t-statistic (independent samples) between two conditions over all time points and extracts clusters (consecutive time segments) of activity whose t-statistic exceeds a predefined threshold ($\alpha = 0.01$), then the absolute t-statistics within each cluster were summed to produce cluster-level statistics. To generate a chance-level distribution from which the appropriate threshold could be determined, trials were randomly

partitioned between the two conditions and the t-test and clustering redone (1000 partitions). From every partition the largest cluster-level statistic was used to generate a largest chance cluster distribution. By comparing the real cluster-level statistic against the largest chance cluster distribution, significant clusters could be determined if the observed cluster value exceeded a set percentage of largest chance cluster values ($p = 0.05$). In this way, sensitivity to short time-scale differences is greatly reduced, but the overall false-alarm rate remains below the designated p-value. This test was carried out once for each of the three factors. Additionally, to see if grip and orientation tuning differed between contra- and ipsi-lateral trials, the CBPT was repeated for those trials separately.

To test which units were task-related, if a unit had a significant effect of any of the three factors at any tested time point of the CBPT, it was considered task-related. Crucially, all analyses only considered units that were determined to be task-related. As a control, if a 3-way ANOVA is used in place of the CBPT, approximately the same amount of significance is found overall, suggesting that the CBPT does not over-estimate the level of tuning for each unit.

Coordinate frame analysis

In order to determine which units significantly exhibited the characteristics of an extrinsic or intrinsic coordinate frame (Fig. 7) activity was first averaged over trials of each of the five orientations in sliding bins (width: 200 ms, interval: 50 ms) and correlated between contra- and ipsi-lateral trials (Fig. 9B). To test significance, trials were randomly permuted with respect to orientation (while keeping grip type and hand consistent), and the average and correlation were redone (1000 repetitions). If the absolute value of the real correlation obtained from the unshuffled data exceeded 95% of the permuted correlation distribution ($p = 0.05$), this unit was deemed to be significantly representing an extrinsic or intrinsic reference frame, corresponding to an r-value > 0 or < 0, respectively.

Dimensionality reduction

A common problem with large data sets is their inherent complexity. Principal component analysis (PCA) is commonly employed to reduce the dimensionality of such data sets by finding a low dimensional representation of the data by defining independent linear combinations, or weighted averages, of units that explain most of

the variance in firing rates. PCA finds a 'decoder', **D**, which represents a linear mapping of the full data onto a compressed read out. Using an 'encoder', **F**, data can then be approximately reconstructed by decompressing it.

To formalize this, given a matrix of data **X**, where each row contains he average firing rates of one neuron for all task conditions, PCA finds an encoder, **F**, and an equivalent decoder, **D**, which minimizes the loss function

$$L = \|\mathbf{X} - \mathbf{FDX}\|^2$$

under the constraint that the principal axes are normalized and orthogonal and therefore $\mathbf{D} = \mathbf{F}^{\mathrm{T}}$ (Hastie et al., 2009), and where the matrix norm is the Frobenius norm, i.e. $\|\mathbf{X}\|^2 = \sum_i \sum_j X_{ij}^2$. Unfortunately, data that is represented in this way heavily mixes the effect of different task parameters between latent dimensions, since no information regarding the actual task conditions is present in the loss function.

However, we would like to extract dimensions that dissociate our specific task conditions. To achieve this demixed principal component analysis (dPCA) was performed (Kobak et al., 2014) using freely available code: https://github.com/wielandbrendel/dPCA.

dPCA is similar to classical PCA in the sense that it seeks to find a rotation of the full neural space that explains most of the variance in average firing rates in a small number of latent dimensions. In contrast to PCA, dPCA seeks to explain marginalized variance with respect to our specific task variables (hand, grip type, orientation, and time), instead of merely explaining total variance. The differences between traditional PCA and dPCA can be formalized by comparing the loss functions that are minimized in each procedure. dPCA utilizes a separate encoder and decoder

$$L = \sum_{\emptyset} L_{\emptyset} = \sum_{\emptyset} (\|\mathbf{X}_{\emptyset} - \mathbf{F}_{\emptyset}\mathbf{D}_{\emptyset}\mathbf{X}\|^2 + \lambda\|\mathbf{D}_{\emptyset}\|^2)$$

where $\mathbf{X}_{\emptyset}$ is the marginalization of the full data with respect to each of our task parameters of interest and the $\lambda\|\mathbf{D}_{\emptyset}\|^2$ term is a regularization parameter, preventing overfitting. Marginalizations of **X** can be obtained by averaging over all parameters

which are not being investigated and subtracting all simpler marginalizations (then replicating matrix entries so that $\mathbf{X}_\emptyset$ has the same dimensionality as $\mathbf{X}$). In our case the marginalizations of interest were Time, Hand x Time, Grip x Time, and Orientation x Time (for more detail regarding marginalization and regularization see Brendel et al., 2011; Kobak et al., 2014). The values of λ determined via cross-validation for brain areas AIP and F5 in the pooled data were $1.3 \bullet 10^{-5}$ and $5.8 \bullet 10^{-6}$, respectively.

In contrast to Kobak et al. (2014), We did not construct our time-courses in order to extract certain effects during certain enforced task epochs. All extracted dimensions were permitted to vary along the entire time axis in addition to their respective dimension.

In addition to finding demixed latent dimensions, our goal was to find latent dimensions in the pooled data of both monkeys that accurately represented aspects of the task that were present of both monkeys. Crucially, we wanted to exclude dimensions that could only explain variance in the units taken from a single monkey. In order to achieve this dPCA was first carried out (with cross-validated regularization parameters) on the data of each monkey separately. Next, for each brain area, all pairs of dPCs between the two monkeys were correlated. If any pair of dimensions produced an absolute correlation of at least 0.45 (20% variance explained), those dimensions were considered to be robust in both monkeys, and all other dimensions were discarded. Next, each set of shared dimensions was correlated with the dimensions produced by executing dPCA on the pooled data from both monkeys. Similar to the previous computation, dimensions in the pooled dPCA were extracted if they were significantly (absolute r-value ≥ 0.45) correlated with single dimensions in both the set of dimensions from monkey S and monkey P. All other dimensions were discarded, so that only these dimensions were plotted in Figure 8.

To determine whether the encoding axes, f_1 and f_2, of two dPCA components were orthogonal, the dot product, $f_1 \cdot f_2$, between the axes, was calculated. To determine the chance level, we calculated the dot product between many pairs of randomly (uniform) generated unit vectors in N dimensions, where N is the number of units. For large N, the distribution of dot products is approximately Gaussian. Therefore, if $|f_1 \cdot f_2| > 3.3 * N^{-1/2}$, we classified them as significantly non-orthogonal ($p < 0.001$). Significantly non-orthogonal dimensions that also had a

spearman correlation above 0.2 (and $p < 0.001$) are marked with a star in Figure 9 (for more information see Kobak et al., 2014).

A decoding procedure was undertaken to determine if the dPCs provided useful decoding axes for the task conditions. We used 100 iterations of stratified Monte Carlo leave-group-out cross-validation, where on each iteration we held out one random trial for each unit in each condition forming X_{test} (as the units were not recorded simultaneously, we do not have recordings of all units in any actual trial). Remaining trials were averaged to form a training set X_{train}. We then calculated dPCA on X_{train} and used the resulting components as linear classifiers for the trials in X_{test}. We then used 100 shuffles to compute Monte Carlo distribution of classification accuracies expected by chance. For each unit and iteration we shuffled all available trials between conditions, respecting the number of trials per condition. If the real classification accuracy exceed that expected by chance on all iterations and for 10 consecutive time bins, classification was considered significant and are marked as black bars in Figure 8 (Right Panels) (for more information see Kobak et al., 2014).

Task space correlation

The objective of the task space correlation was to determine how similar the functional representation of our specific behavioral task was between areas AIP and F5. Firstly, spike rates were averaged in 400 ms bins for five epochs (cue on, cue off, early memory, late memory, and movement). For the results in Figure 9B, pairwise Euclidean distance was calculated between all task conditions, separately for each area and each hand, in the space of all task-related units across both monkeys. The unique values resulting from the distance calculation, i.e., the upper triangle of the distance matrix excluding the diagonal, were correlated (Spearman r) between areas.

In order to determine if correlations significantly differed between contra- and ipsi-lateral trials, a permutation test was carried out. The contra- vs ipsi-lateral assignment of trials was either flipped or remained unchanged, randomly, for each unit individually, while keeping the other task conditions constant. This process was repeated 1000 times and the real difference in r-value between contra- and ipsi-lateral trials was compared to the distribution of all permutations. If the real r-value difference was greater than the chance distribution ($p < 0.05$, false-discovery rate correction), then the difference was considered to be significant.

3 – GENERAL DISCUSSION

3.1 – Summary

In chapter 2.1 we saw how large-scale parallel recordings can be leveraged to make behavioral predictions on single trials. The methods used to extract single-trial predictions varied in their performance, but population-based methods provided the most consistent and meaningful interpretation of the data. In addition, the success of these behavioral predictions could be used to make inferences about how areas differ in their contribution to preparation of grasping movements. It was found that while reaction time could be predicted from the population activity of either area, performance was significantly higher using the data from premotor cortex, suggesting that population activity in premotor cortex may have a more direct effect on behavior.

In chapter 2.2 we saw how preparation and movement intermingle and interact with one another on the continuum between immediate and withheld movement. Our population-based analysis and dimensionality reduction enabled interpretation of the data, even though single neuron tuning properties were highly temporally and functionally complex. Activity in AIP stabilized during the memory period, while it continued to evolve in F5, revealing a decodable signature of time. Furthermore, activity during movement initiation clustered into two groups, movements initiated as fast as possible and movements from memory, showing how a state shift likely occurs on the border between these two types of actions.

In chapter 2.3 we saw that the question of how motor cortex controls movement is an ongoing issue in the field. We addressed some details about recent methodology used to extract rotational dynamics in motor cortex. We saw how a simple simulation could explore the limits of mathematical methodology, and how models of movement generation can be essential tools in adding perspective to empirical results.

Finally, in chapter 2.4 we saw how the specificity of hand use could be used as a tool to dissociate levels of abstraction in the visual to motor transformation in parietal and premotor cortex. While preparatory activity was mostly hand-invariant in parietal cortex, activity in premotor cortex dissociated the intended hand use well before movement. Furthermore, we saw how appropriate dimensionality reduction techniques could disentangle the effects of multiple task parameters and find latent dimensions consistent between areas and animals.

3.2 – Outlook

"… the vague question of 'how the brain works' can be meaningfully reduced to the more precise, and proximally answerable question of how do the connectivity and dynamics of distributed neural circuits give rise to specific behaviors and computations"

(Gao and Ganguli, 2015)

Understanding the kinds of computations that can be carried out by a circuit does not necessarily entail simulating such a network to the finest degree of detail, down to the cell membranes and messenger proteins. As an analogy, when presented with an algebraic formula to solve, mathematics informs us how to make inferences about the kind of solutions that satisfy our equation of interest without explicitly solving for x. The same process of analysis can be applied to neuroscience research. To formalize this way of thinking, the circuit-motifs hypothesis has been proposed in recent years (Figure 1) (Womelsdorf et al., 2014; Womelsdorf and Everling, 2015).
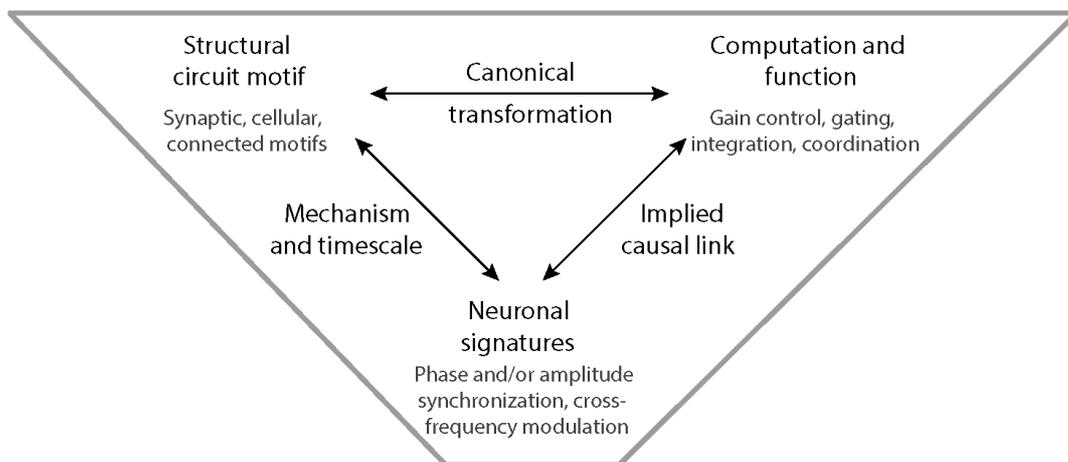
Figure 1. Dynamic motifs represent a tight linkage of three components: the (1) circuit motif describes the structural basis (synaptic, cellular, local connectivity) that gives rise to (2) a characteristic neuronal activation signature, and (3) completion of a dynamic motif requires a link to a canonical input-output transformation that serves to implement a behavioral function. Modified from Womelsdorf et al. (2014).

Although this framework has been proposed largely to aid in the understanding of rhythmicity and how the functional significance of oscillations in

differing frequency ranges can be disentangled, it can be used as a general framework of computation in the brain. Most motor neuroscience research in primates relies on relating neural activation signatures, such as firing rate, spike frequency, or synchronization, to behavior. However, the essence of this tripartite approach makes explicit the linking of circuit motifs and computations to these neural activation signatures. In this context, one must directly link synaptic properties, cellular identity, and local connectivity to specific neural signatures, as well as linking those signatures to specific computations and functions, such as selective integration or segregation. Going even further, one must investigate what kinds of computations can be undertaken using a specific neural circuit, or what kinds of circuits would be necessary to compute a given function. Then, by isolating and identifying each of these possible connections, these links can be manipulated through experiment.

Purely behavior-based neuroscience will always have a necessary place in the field. After all, the world we inhabit is not to be confused with the model realities we generate in experiments.

> *"Nothing in neuroscience makes sense except in the light of behavior"*
> (Shepherd, 1988)

We will always only have one piece of the picture. The incremental goal of neuroscience should be not to solve the whole process, but rather find the key ingredients and how they relate to one another.

Even Cajal, the first person to extensively observe and document individual neurons, was able to deduce the dynamic nature of the brain (De Carlos and Borrell, 2007). In a system where cells cannot replicate, as is the case for neurons, he held that the specific connections between neurons and their ability to change over time was the fundamental power of the brain. That is to say, he recognized that the function of neurons must not be delimited to the specific firing patterns we observe in a given experiment, but rather by the dynamic circuits in which these neurons are embedded.

4 – BIBLIOGRAPHY

Aflalo T, Kellis S, Klaes C, Lee B, Shi Y, Pejsa K, Shanfield K, Hayes-Jackson S, Aisen M, Heck C, Liu C, Andersen RA (2015) Neurophysiology. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. Science 348:906–910.

Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV (2011) Single-trial neural correlates of arm movement preparation. Neuron 71:555–564.

Ames KC, Ryu SI, Shenoy KV (2014) Neural Dynamics of Reaching following Incorrect or Absent Motor Preparation. Neuron 81:438–451.

Baumann MA, Fluet M-C, Scherberger H (2009) Context-specific grasp movement representation in the macaque anterior intraparietal area. J Neurosci 29:6436–6448.

Belmalih A, Borra E, Contini M, Gerbella M, Rozzi S, Luppino G (2009) Multimodal architectonic subdivision of the rostral part (area F5) of the macaque ventral premotor cortex. J Comp Neurol 512:183–217.

Benjamin PR, Staras K, Kemenes G (2010) What roles do tonic inhibition and disinhibition play in the control of motor programs? Front Behav Neurosci 4:30.

Binkofski F, Buccino G, Stephan KM, Rizzolatti G, Seitz RJ, Freund HJ (1999) A parieto-premotor network for object manipulation: evidence from neuroimaging. Exp Brain Res 128:210–213.

Bonini L, Ugolotti Serventi F, Bruni S, Maranesi M, Bimbi M, Simone L, Rozzi S, Ferrari PF, Fogassi L (2012) Selectivity for grip type and action goal in macaque inferior parietal and ventral premotor grasping neurons. J Neurophysiol 108:1607–1619.

Borra E, Belmalih A, Calzavara R, Gerbella M, Murata A, Rozzi S, Luppino G (2008) Cortical connections of the macaque anterior intraparietal (AIP) area. Cereb Cortex 18:1094–1111.

Boussaoud D, Wise SP (1993a) Primate frontal cortex: effects of stimulus and movement. Exp Brain Res 95:28–40.

Boussaoud D, Wise SP (1993b) Primate frontal cortex: neuronal activity following attentional versus intentional cues. Exp Brain Res 95:15–27.

Brendel W, Romo R, Machens CK (2011) Demixed principal component analysis. Advances in Neural Information Processing Systems:2654–2662.

Brinkman J, Kuypers HG (1973) Cerebral control of contralateral and ipsilateral arm, hand and finger movements in the split-brain rhesus monkey. Brain 96:653–674.

Buccino G, Binkofski F, Fink GR, Fadiga L, Fogassi L, Gallese V, Seitz RJ, Zilles K, Rizzolatti G, Freund HJ (2001) Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. Eur J Neurosci 13:400–

404.

Buonomano DV, Laje R (2010) Population clocks: motor timing with neural dynamics. Trends in Cognitive Sciences 14:520–527.

Cajal SY (1888) Estructura de los centros nerviosos de las aves. Rev Trim Hostil Norm Pat 1:1–10.

Caminiti R, Johnson PB, Urbano A (1990) Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. 10:2039–2058.

Carmena JM, Lebedev MA, Henriquez CS, Nicolelis MAL (2005) Stable ensemble performance with single-neuron variability during reaching movements in primates. J Neurosci 25:10712–10716.

Carpaneto J, Umilta MA, Fogassi L, Murata A, Gallese V, Micera S, Raos V (2011) Decoding the activity of grasping neurons recorded from the ventral premotor area F5 of the macaque monkey. Neuroscience 188:80–94.

Chang SWC, Dickinson AR, Snyder LH (2008) Limb-specific representation for reaching in the posterior parietal cortex. J Neurosci 28:6128–6140.

Chen X, Scangos KW, Stuphorn V (2010) Supplementary motor area exerts proactive and reactive control of arm movements. J Neurosci 30:14657–14675.

Chestek CA, Batista AP, Santhanam G, Yu BM, Afshar A, Cunningham JP, Gilja V, Ryu SI, Churchland MM, Shenoy KV (2007) Single-neuron stability during repeated reaching in macaque premotor cortex. J Neurosci 27:10742–10750.

Chestek CA, Gilja V, Nuyujukian P, Foster JD, Fan JM, Kaufman MT, Churchland MM, Rivera-Alvidrez Z, Cunningham JP, Ryu SI, Shenoy KV (2011) Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. J Neural Eng 8:045005.

Churchland MM et al. (2010a) Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat Neurosci 13:369–378.

Churchland MM, Abbott LF (2012) Two layers of neural variability. Nat Neurosci 15:1472–1474.

Churchland MM, Afshar A, Shenoy KV (2006a) A central source of movement variability. Neuron 52:1085–1096.

Churchland MM, Cunningham JP (2014) A Dynamical Basis Set for Generating Reaches. Cold Spring Harb Symp Quant Biol 79:67–80.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. Nature 487:51–56.

Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV (2010b) Cortical preparatory activity: representation of movement or first cog in a

dynamical machine? Neuron 68:387–400.

Churchland MM, Shenoy KV (2007) Delay of movement caused by disruption of cortical preparatory activity. J Neurophysiol 97:348–359.

Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV (2006b) Neural variability in premotor cortex provides a signature of motor preparation. J Neurosci 26:3697–3712.

Churchland MM, Yu BM, Sahani M, Shenoy KV (2007) Techniques for extracting single-trial activity patterns from large-scale neural recordings. Current Opinion in Neurobiology 17:609–618.

Cisek P (2007) Cortical mechanisms of action selection: the affordance competition hypothesis. Philosophical Transactions of the Royal Society B: Biological Sciences 362:1585–1599.

Cisek P, Crammond DJ, Kalaska JF (2003) Neural activity in primary motor and dorsal premotor cortex in reaching tasks with the contralateral versus ipsilateral arm. J Neurophysiol 89:922–942.

Cisek P, Kalaska JF (2002) Simultaneous encoding of multiple potential reach directions in dorsal premotor cortex. J Neurophysiol.

Cisek P, Kalaska JF (2010) Neural mechanisms for interacting with a world full of action choices. Annu Rev Neurosci 33:269–298.

Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJ, Velliste M, Boninger ML, Schwartz AB (2012) High-performance neuroprosthetic control by an individual with tetraplegia. The Lancet.

Cowley BR, Kaufman MT, Butler ZS, Churchland MM, Ryu SI, Shenoy KV, Yu BM (2013) DataHigh: graphical user interface for visualizing and interacting with high-dimensional neural activity. J Neural Eng 10:066012.

Crammond DJ, Kalaska JF (2000) Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. J Neurophysiol 84:986–1005.

Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings. Nat Neurosci 17:1500–1509.

Day BL, Rothwell JC, Thompson PD, Maertens de Noordhout A, Nakashima K, Shannon K, Marsden CD (1989) Delay in the execution of voluntary movement by electrical or magnetic brain stimulation in intact man. Evidence for the storage of motor programs in the brain. Brain 112 (Pt 3):649–663.

De Carlos JA, Borrell J (2007) A historical reflection of the contributions of Cajal and Golgi to the foundations of neuroscience. Brain Res Rev 55:8–16.

DiCarlo JJ, Zoccolan D, Rust NC (2012) How does the brain solve visual object

recognition? Neuron 73:415–434.

Dickey AS, Suminski A, Amit Y, Hatsopoulos NG (2009) Single-unit stability using chronically implanted multielectrode arrays. J Neurophysiol 102:1331–1339.

Donchin O, Gribova A, Steinberg O, Bergman H, Vaadia E (1998) Primary motor cortex is involved in bimanual coordination. Nature 395:274–278.

Doya K (1993) Universality of fully connected recurrent neural networks. Dept of Biology.

Dyson FJ (2012) History of science. Is science mostly driven by ideas or by tools? Science 338:1426–1427.

Edelman GM, Mountcastle VB (1978) The Mindful Brain. Cambridge, Mass: MIT Press.

Engel AK, Fries P (2010) Beta-band oscillations--signalling the status quo? Current Opinion in Neurobiology 20:156–165.

Erlich JC, Brunton BW, Duan CA, Hanks TD, Brody CD (2015) Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. eLife.

Evarts EV (1966) Pyramidal tract activity associated with a conditioned hand movement in the monkey. J Neurophysiol 29:1011–1027.

Evinger C, Kaneko CR, Fuchs AF (1982) Activity of omnipause neurons in alert cats during saccadic eye movements and visual stimuli. J Neurophysiol 47:827–844.

Fagg AH, Arbib MA (1998) Modeling parietal-premotor interactions in primate control of grasping. Neural Netw 11:1277–1303.

Fluet M-C, Baumann MA, Scherberger H (2010) Context-specific grasp movement representation in macaque ventral premotor cortex. J Neurosci 30:15175–15184.

Fogassi L, Gallese V, Buccino G, Craighero L, Fadiga L, Rizzolatti G (2001) Cortical mechanism for the visual guidance of hand grasping movements in the monkey: A reversible inactivation study. Brain 124:571–586.

Fries P (2015) Rhythms for Cognition: Communication through Coherence. Neuron 88:220–235.

Frith C, Frith U (2010) Learning from Others: Introduction to the Special Review Series on Social Neuroscience. Neuron 65:739–743.

Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. Brain 119 (Pt 2):593–609.

Gallese V, Murata A, Kaseda M, Niki N, Sakata H (1994) Deficit of hand preshaping after muscimol injection in monkey parietal cortex. Neuroreport 5:1525–1529.

Ganguly K, Secundo L, Ranade G, Orsborn A, Chang EF, Dimitrov DF, Wallis JD, Barbaro NM, Knight RT, Carmena JM (2009) Cortical representation of ipsilateral arm movements in monkey and man. J Neurosci 29:12948–12956.

Gao P, Ganguli S (2015) On simplicity and complexity in the brave new world of large-scale neuroscience. Current Opinion in Neurobiology 32:148–155.

Genovesio A, Tsujimoto S, Wise SP (2006) Neuronal activity related to elapsed time in prefrontal cortex. J Neurophysiol 95:3281–3285.

Georgopoulos AP, Carpenter AF (2015) Coding of movements in the motor cortex. Current Opinion in Neurobiology 33:34–39.

Georgopoulos AP, Naselaris T, Merchant H, Amirikian B (2007) Reply to Kurtzer and Herter. J Neurophysiol 97:4391–4392.

Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. Science 233:1416–1419.

Gerbella M, Belmalih A, Borra E, Rozzi S, Luppino G (2011) Cortical connections of the anterior (F5a) subdivision of the macaque ventral premotor area F5. Brain Struct Funct 216:43–65.

Gerits A, Farivar R, Rosen BR, Wald LL, Boyden ES, Vanduffel W (2012) Optogenetically induced behavioral and functional network changes in primates. Curr Biol 22:1722–1726.

Ghez C, Favilla M, Ghilardi MF, Gordon J, Bermejo R, Pullman S (1997) Discrete and continuous planning of hand movements and isometric force trajectories. Exp Brain Res 115:217–233.

Goel A, Buonomano DV (2014) Timing as an intrinsic property of neural networks: evidence from in vivo and in vitro experiments. Philos Trans R Soc Lond, B, Biol Sci 369:20120460.

Golgi C (1873) Sulla struttura della sostanza grigia del cervello. Lombardia: Gazzetta Medica Italiana.

Gollisch T, Meister M (2010) Eye smarter than scientists believed: neural computations in circuits of the retina. Neuron 65:150–164.

Goudar V, Buonomano DV (2014) Useful dynamic regimes emerge in recurrent networks. Nat Neurosci 17:487–489.

Gouvêa TS, Monteiro T, Motiwala A, Soares S (2015) Striatal dynamics explain duration judgments. bioRxiv.

Gozani SN, Miller JP (1994) Optimal discrimination and classification of neuronal action potential waveforms from multiunit, multichannel recordings using software-based linear filters. IEEE Trans Biomed Eng 41:358–372.

Grafton ST (2010) The cognitive neuroscience of prehension: recent developments.

Exp Brain Res 204:475–491.

Grafton ST, Arbib MA, Fadiga L, Rizzolatti G (1996) Localization of grasp representations in humans by positron emission tomography. 2. Observation compared with imagination. Exp Brain Res 112:103–111.

Gray CM, König P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature 338:334–337.

Grèzes J, Armony JL, Rowe J, Passingham RE (2003) Activations related to "mirror" and "canonical" neurones in the human brain: an fMRI study. Neuroimage 18:928–937.

Halsband U, Ito N, Tanji J, Freund HJ (1993) The role of premotor cortex and the supplementary motor area in the temporal control of movement in man. Brain 116 (Pt 1):243–266.

Hanks TD, Kopec CD, Brunton BW, Duan CA, Erlich JC, Brody CD (2015) Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature 520:220–223.

Harvey AC (1990) Forecasting, Structural Time Series Models and the Kalman Filter. Cambridge: Cambridge University Press.

Hastie T, Tibshirani R, Friedman J (2009) The Elements of Statistical Learning. Springer.

Hawkins J, Ahmad S (2015) Why Neurons Have Thousands of Synapses, A Theory of Sequence Memory in Neocortex. arXiv.

Hawkins J, Blakeslee S (2007) On Intelligence. Macmillan.

Hebb DO (1949) The Organization of Behavior. New York: Wiley.

Hennequin G, Vogels TP, Gerstner W (2014) Optimal control of transient dynamics in balanced networks supports generation of complex movements. Neuron 82:1394–1406.

Hepp-Reymond MC, Hüsler EJ, Maier MA, Qi HX (1994) Force-related neuronal activity in two regions of the primate ventral premotor cortex. Can J Physiol Pharmacol 72:571–579.

Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, Haddadin S, Liu J, Cash SS, van der Smagt P, Donoghue JP (2012) Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. Nature 485:372–375.

Hoshi E, Tanji J (2000) Integration of target and body-part information in the premotor cortex when planning action. Nature 408:466–470.

Hoshi E, Tanji J (2006) Differential involvement of neurons in the dorsal and ventral premotor cortex during processing of visual signals for action planning. J

Neurophysiol 95:3596–3616.

Hoshi E, Tanji J (2007) Distinctions between dorsal and ventral premotor areas: anatomical connectivity and functional properties. Current Opinion in Neurobiology 17:234–242.

Hubel DH (1957) Tungsten Microelectrode for Recording from Single Units. Science 125:549–550.

Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol (Lond) 160:106–154.

Iacoboni M, Dapretto M (2006) The mirror neuron system and the consequences of its dysfunction. Nat Rev Neurosci 7:942–951.

Inoue K-I, Takada M, Matsumoto M (2015) Neuronal and behavioural modulations by pathway-selective optogenetic stimulation of the primate oculomotor system. Nat Commun 6:8378.

Janssen P, Scherberger H (2015) Visual guidance in control of grasping. Annu Rev Neurosci 38:69–86.

Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. Nat Neurosci 8:234–241.

Joly O, Ramus F, Pressnitzer D, Vanduffel W, Orban GA (2012) Interhemispheric differences in auditory processing revealed by fMRI in awake rhesus monkeys. Cereb Cortex 22:838–853.

Kakei S, Hoffman DS, Strick PL (2001) Direction of action is represented in the ventral premotor cortex. Nat Neurosci 4:1020–1025.

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. Nat Neurosci 17:440–448.

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2015) Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. eLife 4:e04677.

Kaufman MT, Churchland MM, Santhanam G, Yu BM, Afshar A, Ryu SI, Shenoy KV (2010) Roles of monkey premotor neuron classes in movement preparation and execution. J Neurophysiol 104:799–810.

Kaufman MT, Churchland MM, Shenoy KV (2013) The roles of monkey M1 neuron classes in movement preparation and execution. J Neurophysiol 110:817–825.

Kermadi I, Liu Y, Tempini A, Calciati E, Rouiller EM (1998) Neuronal activity in the primate supplementary motor area and the primary motor cortex in relation to spatio-temporal bimanual coordination. Somatosens Mot Res 15:287–308.

Khanna P, Carmena JM (2015) Neural oscillations: beta band activity across motor networks. Current Opinion in Neurobiology 32:60–67.

Kipke DR, Shain W, Buzsáki G, Fetz E, Henderson JM, Hetke JF, Schalk G (2008) Advanced neurotechnologies for chronic neural interfaces: new horizons and clinical opportunities. J Neurosci 28:11830–11838.

Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, Mrsic-Flogel TD (2011) Functional specificity of local synaptic connections in neocortical networks. Nature 473:87–91.

Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Romo R, Qi X-L, Uchida N, Machens CK (2014) Demixed principal component analysis of population activity in higher cortical areas reveals independent representation of task parameters. arXiv q-bio.NC.

Kohler E, Keysers C, Umilta MA, Fogassi L, Gallese V, Rizzolatti G (2002) Hearing sounds, understanding actions: action representation in mirror neurons. Science 297:846–848.

Kraskov A, Dancause N, Quallo MM, Shepherd S, Lemon RN (2009) Corticospinal Neurons in Macaque Ventral Premotor Cortex with Mirror Properties: A Potential Mechanism for Action Suppression? Neuron 64:922–930.

Kriegeskorte N (2015) Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing. Annu Rev Vis Sci 1:417–446.

Kurata K (1991) Corticocortical inputs to the dorsal and ventral aspects of the premotor cortex of macaque monkeys. Neurosci Res 12:263–280.

Kurata K (2007) Laterality of movement-related activity reflects transformation of coordinates in ventral premotor cortex and primary motor cortex of monkeys. J Neurophysiol 98:2008–2021.

Kurata K (2010) Conditional selection of contra- and ipsilateral forelimb movements by the dorsal premotor cortex in monkeys. J Neurophysiol 103:262–277.

Kurtzer I, Herter TM (2007) Contrasting interpretations of the nonuniform distribution of preferred directions within primary motor cortex. J Neurophysiol 97:4390–authorreply4391–2.

Kutas M, Donchin E (1974) Studies of squeezing: handedness, responding hand, response force, and asymmetry of readiness potential. Science 186:545–548.

Laje R, Buonomano DV, Buonomano DV (2013) Robust timing and motor patterns by taming chaos in recurrent neural networks. Nat Neurosci 16:925–933.

Lau K-Y, Ganguli S, Tang C (2007) Function constrains network architecture and dynamics: a case study on the yeast cell cycle Boolean network. Phys Rev E 75:051907.

Lehmann SJ, Scherberger H (2013) Reach and gaze representations in macaque parietal and premotor grasp areas. J Neurosci 33:7038–7049.

Leon MI, Shadlen MN (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. Neuron 38:317–327.

Li F, Long T, Lu Y, Ouyang Q, Tang C (2004) The yeast cell-cycle network is robustly designed. Proceedings of the National Academy of Sciences 101:4781–4786.

Lillicrap TP, Scott SH (2013) Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. Neuron 77:168–179.

Litwin-Kumar A, Doiron B (2012) Slow dynamics and high variability in balanced cortical networks with clustered connections. Nat Neurosci 15:1498–1505.

Lucchetti C, Ulrici A, Bon L (2005) Dorsal premotor areas of nonhuman primate: functional flexibility in time domain. Eur J Appl Physiol 95:121–130.

Luppino G, Murata A, Govoni P, Matelli M (1999) Largely segregated parietofrontal connections linking rostral intraparietal cortex (areas AIP and VIP) and the ventral premotor cortex (areas F5 and F4). Exp Brain Res 128:181–187.

Luppino GG, Luppino G, Matelli MM, Matelli M, Camarda R, Camarda RR, Rizzolatti G, Rizzolatti GG (1993) Corticocortical connections of area F3 (SMA-proper) and area F6 (pre-SMA) in the macaque monkey. J Comp Neurol 338:114–140.

Machens CK (2010) Demixing population activity in higher cortical areas. Front Comput Neurosci 4:126.

Machens CK, Romo R, Brody CD (2010) Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. J Neurosci 30:350–360.

Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503:78–84.

Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. J Neurosci Methods 164:177–190.

Markram H et al. (2015) Reconstruction and Simulation of Neocortical Microcircuitry. Cell 163:456–492.

Martinez CA, Wang C (2015) Structural constraints on learning in the neural network. J Neurophysiol 114:2555–2557.

Mauritz KH, Wise SP (1986) Premotor cortex of the rhesus monkey: neuronal activity in anticipation of predictable environmental events. Exp Brain Res 61:229–244.

Maynard EM, Nordhausen CT, Normann RA (1997) The Utah intracortical Electrode Array: a recording structure for potential brain-computer interfaces. Electroencephalogr Clin Neurophysiol 102:228–239.

Mazzucato L, Fontanini A, La Camera G (2015) Dynamics of multistable states

during ongoing and evoked cortical activity. J Neurosci 35:8214–8231.

Menozzi P, Piazza A, Cavalli-Sforza L (1978) Synthetic maps of human gene frequencies in Europeans. Science 201:786–792.

Menz VK, Schaffelhofer S, Scherberger H (2015) Representation of continuous hand and arm movements in macaque areas M1, F5, and AIP: a comparative decoding study. J Neural Eng 12:056016.

Michaels JA, Dann B, Intveld RW, Scherberger H (2015) Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. J Neurosci 35:11415–11432.

Michaels JA, Dann B, Scherberger H (2016) Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. PLoS Comput Biol 12:e1005175.

Muakkassa KF, Strick PL (1979) Frontal lobe inputs to primate motor cortex: evidence for four somatotopically organized "premotor" areas. Brain Res 177:176–182.

Murakami M, Mainen ZF (2015) Preparing and selecting actions with neural populations: toward cortical circuit mechanisms. Current Opinion in Neurobiology 33C:40–46.

Murata A, Fadiga L, Fogassi L, Gallese V, Raos V, Rizzolatti G (1997) Object representation in the ventral premotor cortex (area F5) of the monkey. J Neurophysiol 78:2226–2230.

Murata A, Gallese V, Kaseda M, Sakata H (1996) Parietal neurons related to memory-guided hand manipulation. J Neurophysiol 75:2180–2186.

Murata A, Gallese V, Luppino G, Kaseda M, Sakata H (2000) Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP. J Neurophysiol 83:2580–2601.

Musial PG, Baker SN, Gerstein GL, King EA, Keating JG (2002) Signal-to-noise ratio improvement in multiple electrode recording. J Neurosci Methods 115:29–43.

Naselaris T, Merchant H, Amirikian B, Georgopoulos AP (2006) Large-scale organization of preferred directions in the motor cortex. I. Motor cortical hyperacuity for forward reaching. J Neurophysiol 96:3231–3236.

National Research Council (2003) Guidelines for the Care and Use of Mammals in Neuroscience and Behavioral Research. Washington (DC): National Academies Press (US).

Newman MEJ (2004) Fast algorithm for detecting community structure in networks. Phys Rev E 69:066133.

Nordhausen CT, Maynard EM, Normann RA (1996) Single unit recording capabilities

of a 100 microelectrode array. Brain Res 726:129–140.

Novembre J, Stephens M (2008) Interpreting principal component analyses of spatial population genetic variation. Nat Genet 40:646–649.

Okun M, Steinmetz NA, Cossell L, Iacaruso MF, Ko H, Barthó P, Moore T, Hofer SB, Mrsic-Flogel TD, Carandini M, Harris KD (2015) Diverse coupling of neurons to populations in sensory cortex. Nature 521:511–515.

O'Leary T, Sutton AC, Marder E (2015) Computational models in the age of large datasets. Current Opinion in Neurobiology 32:87–94.

Pagan M, Urban LS, Wohl MP, Rust NC (2013) Signals in inferotemporal and perirhinal cortex suggest an untangling of visual target information. Nat Neurosci 16:1132–1139.

Paré M, Smith AM, Rice FL (2002) Distribution and terminal arborizations of cutaneous mechanoreceptors in the glabrous finger pads of the monkey. J Comp Neurol 445:347–359.

Penfield W, Rasmussen T (1950) The cerebral cortex of man. New York: Macmillan.

Peron S, Chen T-W, Svoboda K (2015) Comprehensive imaging of cortical networks. Current Opinion in Neurobiology 32:115–123.

Petrides M, Pandya DN (1984) Projections to the frontal cortex from the posterior parietal region in the rhesus monkey. J Comp Neurol 228:105–116.

Ponce-Alvarez A, Nácher V, Luna R, Riehle A, Romo R (2012) Dynamics of cortical neuronal ensembles transit from decision making to storage for later report. J Neurosci 32:11956–11969.

Poremba A, Mishkin M (2007) Exploring the extent and function of higher-order auditory cortex in rhesus monkeys. Hearing Research 229:14–23.

Poremba A, Saunders RC, Crane AM, Cook M, Sokoloff L, Mishkin M (2003) Functional mapping of the primate auditory system. Science 299:568–572.

Porter KR, Claude A, Fullam EF (1945) A study of tissue culture cells by electron microscopy: methods and preliminary observations. J Exp Med 81:233–246.

Premereur E, Van Dromme IC, Romero MC, Vanduffel W, Janssen P (2015) Effective connectivity of depth-structure-selective patches in the lateral bank of the macaque intraparietal sulcus. Plos Biol 13:e1002072.

Pruszynski JA, Johansson RS (2014) Edge-orientation processing in first-order tactile neurons. Nat Neurosci 17:1404–1409.

Quiroga RQ, Nadasdy Z, Ben-Shaul Y (2004) Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. Neural Comput 16:1661–1687.

Raposo D, Kaufman MT, Churchland AK (2014) A category-free neural population supports evolving demands during decision-making. Nat Neurosci 17:1784–1792.

Reichardt J, Bornholdt S (2006) Statistical mechanics of community detection. Phys Rev E 74:016110.

Riehle A, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. J Neurophysiol 61:534–549.

Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. Nature 497:585–590.

Rizzolatti G, Camarda R, Fogassi L, Gentilucci M, Luppino G, Matelli M (1988) Functional organization of inferior area 6 in the macaque monkey. II. Area F5 and the control of distal movements. Exp Brain Res 71:491–507.

Rizzolatti G, Fadiga L, Gallese V, Fogassi L (1996a) Premotor cortex and the recognition of motor actions. Brain Res Cogn Brain Res 3:131–141.

Rizzolatti G, Fadiga L, Matelli M, Bettinardi V, Paulesu E, Perani D, Fazio F (1996b) Localization of grasp representations in humans by PET: 1. Observation versus execution. Exp Brain Res 111:246–252.

Rizzolatti G, Luppino G (2001) The cortical motor system. Neuron 31:889–901.

Rizzolatti G, Sinigaglia C (2010) The functional role of the parieto-frontal mirror circuit: interpretations and misinterpretations. Nat Rev Neurosci 11:264–274.

Rosenbaum DA (1980) Human movement initiation: specification of arm, direction, and extent. J Exp Psychol Gen 109:444–474.

Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP (2014) Neural constraints on learning. Nature 512:423–426.

Schaffelhofer S (2014) From vision to action: Hand representations in macaque grasping areas AIP, F5, and M1.

Schaffelhofer S, Agudelo-Toro A, Scherberger H (2015) Decoding a wide range of hand configurations from macaque motor, premotor, and parietal cortices. J Neurosci 35:1068–1081.

Schettino LF, Adamovich SV, Bagce H, Yarossi M, Tunik E (2015) Disruption of activity in the ventral premotor but not the anterior intraparietal area interferes with on-line correction to a haptic perturbation during grasping. J Neurosci 35:2112–2117.

Schwartz AB, Kettner RE, Georgopoulos AP (1988) Primate motor cortex and free arm movements to visual targets in three- dimensional space. I. Relations between single cell discharge and direction of movement. 8:2913–2927.

Scott SH (2000) Population vectors and motor cortex: neural coding or

epiphenomenon? Nat Neurosci.

Scott SH, Gribble PL, Graham KM, Cabel DW (2001) Dissociation between hand motion and population vectors from neural activity in motor cortex. Nature 413:161–165.

Sejnowski TJ, Churchland PS, Movshon JA (2014) Putting big data to good use in neuroscience. Nat Neurosci 17:1440–1441.

Shenoy KV, Sahani M, Churchland MM (2013) Cortical Control of Arm Movements: A Dynamical Systems Perspective. Annu Rev Neurosci 36:337–359.

Shepherd GM (1988) Neurobiology. Oxford University Press.

Sherrington CS (1906) Observations on the scratch-reflex in the spinal dog. J Physiol (Lond) 34:1–50.

Snyder LH, Dickinson AR, Calton JL (2006) Preparatory delay activity in the monkey parietal reach region predicts reach reaction times. J Neurosci 26:10091–10099.

Soteropoulos DS, Edgley SA, Baker SN (2011) Lack of evidence for direct corticospinal contributions to control of the ipsilateral forelimb in monkey. J Neurosci 31:11208–11219.

Stevenson IH, Kording KP (2011) How advances in neural recording affect data analysis. Nat Neurosci 14:139–142.

Suminski AJ, Mardoum P, Lillicrap TP, Hatsopoulos NG (2015) Temporal evolution of both premotor and motor cortical tuning properties reflect changes in limb biomechanics. J Neurophysiol 113:2812–2823.

Sussillo D (2014) Neural circuits as computational dynamical systems. Current Opinion in Neurobiology 25:156–163.

Sussillo D, Barak O (2013) Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. Neural Comput 25:626–649.

Sussillo D, Churchland MM, Kaufman MT, Shenoy KV (2015) A neural network that finds a naturalistic solution for the production of muscle activity. Nat Neurosci 18:1025–1033.

Tanaka H (2015) Modeling the motor cortex: Optimality, recurrent neural networks, and spatial dynamics. Neurosci Res 104:64–71.

Tanji J, Evarts EV (1976) Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. J Neurophysiol 39:1062–1068.

Tanji J, Okano K, Sato KC (1988) Neuronal activity in cortical motor areas related to ipsilateral, contralateral, and bilateral digit movements of the monkey. J Neurophysiol 60:325–343.

Townsend BR, Subasi E, Scherberger H (2011) Grasp movement decoding from

premotor and parietal cortex. J Neurosci 31:14386–14398.

Tunik E, Frey SH, Grafton ST (2005) Virtual lesions of the anterior intraparietal area disrupt goal-dependent on-line adjustments of grasp. Nat Neurosci 8:505–511.

Umiltà MAM, Kohler EE, Gallese VV, Fogassi LL, Fadiga LL, Keysers CC, Rizzolatti GG (2001) I Know What You Are Doing - A Neurophysiological Study. Neuron 31:11–11.

Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB (2008) Cortical control of a prosthetic arm for self-feeding. Nature 453:1098–1101.

Venkataramani S, Taylor WR (2010) Orientation selectivity in rabbit retinal ganglion cells is mediated by presynaptic inhibition. J Neurosci 30:15664–15676.

Vigneswaran G, Kraskov A, Lemon RN (2011) Large identified pyramidal cells in macaque motor and premotor cortex exhibit "thin spikes": implications for cell type classification. J Neurosci 31:14235–14242.

Webster MJ, Bachevalier J, Ungerleider LG (1994) Connections of inferior temporal areas TEO and TE with parietal and frontal cortex in macaque monkeys. Cerebral Cortex 4:470–483.

Wise SP (1985) The primate premotor cortex: past, present, and preparatory. Annu Rev Neurosci 8:1–19.

Wise SP, Kurata K (1989) Set-Related Activity in the Premotor Cortex of Rhesus Monkeys: Effect of Triggering Cues and Relatively Long Delay Intervals. Somatosens Mot Res 6:455–476.

Womelsdorf T, Everling S (2015) Long-Range Attention Networks: Circuit Motifs Underlying Endogenously Controlled Stimulus Selection. Trends Neurosci 38:682–700.

Womelsdorf T, Valiante TA, Sahin NT, Miller KJ, Tiesinga P (2014) Dynamic circuit motifs underlying rhythmic gain control, gating and integration. Nat Neurosci 17:1031–1039.

Xiao J, Padoa-Schioppa C, Bizzi E (2006) Neuronal correlates of movement dynamics in the dorsal and ventral premotor area in the monkey. Exp Brain Res 168:106–119.

Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M (2009) Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. J Neurophysiol 102:614–635.

Yuste R (2015) From the neuron doctrine to neural networks. Nat Rev Neurosci 16:487–497.

Zhang Y, Chen Y, Bressler SL, Ding M (2008) Response preparation and inhibition: the role of the cortical sensorimotor beta rhythm. Neuroscience 156:238–246.

5 – CURRICULUM VITAE

Jonathan A. Michaels

Primate Neurobiology

German Primate Center

Kellnerweg 4

37077 Göttingen, Germany

+49 551 3851 484

JMichaels@dpz.eu

http://JMichaels.me

@JonAMichaels

Education: Dr. rer. nat. Systems Neuroscience (summa cum laude), GGNB, Georg-August-Universität Göttingen (2017). Dissertation: *"Towards population coding principles in the primate premotor and parietal grasping network"*

Bachelor of Science (Honours), Queen's University (2011), Kingston, Canada. Dissertation: *"Influence of water maze learning on low-frequency-induced synaptic potentiation in the rat hippocampus"*

Positions:

2017 – Present — Postdoctoral Fellow, Neurobiology Lab of Hansjörg Scherberger, German Primate Center, Göttingen, Germany.

2011 – 2017 — Graduate Student, Neurobiology Lab of Hansjörg Scherberger, German Primate Center, Göttingen, Germany.

2010 – 2011 — Bachelor Student, Neuroplasticity Lab of Hans C. Dringenberg, Queen's University, Kingston, Canada.

2009 – 2011 — Undergraduate Researcher, Integrative Motor Behaviour Lab of Stephen H. Scott, Queen's University, Kingston, Canada.

2008 – 2009 — Research Assistant, Language and Cognition Lab of Stanka A. Fitneva, Queen's University, Kingston, Canada.

Research:

Interests: Neural population dynamics of grasping movements in primary motor, premotor, and parietal cortex of monkeys using single-electrode and multi-electrode array electrophysiology. Recurrent neural network modeling and training. Dimensionality reduction methods. Flexible feedback control.

Skills: Animal care, training, and surgery (monkey and rat). Neural data acquisition and processing (experimental design, recording, spike sorting, etc.). EMG recording.

Awards and Honours:

Sloan-Swartz Travel Fellowship (2016).

Neural Control of Movement Travel Scholarship (2016).

Doctoral Thesis awarded "summa cum laude" (2016).

Dean's Honour List, Queen's University, Kingston, Canada (2009 – 2011).

Peer-reviewed publications:

Michaels JA, Dann B, Scherberger H (2016). Neural population dynamics during reaching are better explained by a dynamical system than representational tuning. *PLOS Computational Biology*, *12*(11), e1005175. doi:10.1371/journal.pcbi.1005175.

Michaels JA, Scherberger H (2016). hebbRNN: A reward-modulated Hebbian learning rule for recurrent neural networks. *The Journal of Open Source Software*. doi:10.21105/joss.00060.

Dann B, **Michaels JA**, Schaffelhofer S, Scherberger H (2016). Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates. *eLife*. doi:10.7554/eLife.15719.

Dann B, **Michaels JA**, Scherberger H (2016). Separable decoding of cue, intention, and movement information from the fronto-parietal grasping-network. *Proceedings of the Sixth International Brain-Computer Interface Meeting: BCI Past, Present, and Future,* 218. doi:10.3217/978-3-85125-467-9.

Michaels JA, Dann B, Intveld RW, Scherberger H (2015). Predicting reaction time from the neural state space of the premotor and parietal grasping network. *Journal of Neuroscience*, *35*(32), 11415–11432. doi:10.1523/JNEUROSCI.1714-15.2015.

Yang L, **Michaels JA**, Pruszynski JA, Scott SH (2011). Rapid motor responses quickly integrate visuospatial task constraints. *Experimental Brain Research, 211*(2): 231-242. doi:10.1007/s00221-011-2674-3.

In-progress publications:

Michaels JA, Dann B, Intveld RW, Scherberger H (in revision). Probing the continuum of immediate to withheld grasping movements in the macaque fronto-parietal network.

Michaels JA, Scherberger H (in revision). Hand-invariant to hand-specific representation of grasping movements in the macaque parieto-frontal network.

Michaels JA, Scherberger H (in prep). Uncovering neural population coding mechanisms in the cortical motor system.

Schaffelhofer S, **Michaels JA**, Kratzenberger A, Scherberger H (in prep). Partner neurons: a neural substrate coding the actions of others.

Michaels JA, Dann B, Intveld RW, Agudelo-Toro A, Scherberger H (in prep). Generating complex grasping movements with a modular recurrent neural network model of the primate dorsal stream.

Dann B, **Michaels JA**, Scherberger H (in prep). Separable decoding of cue, intention, and movement information from the fronto-parietal grasping-network.

Intveld RW, Dann B, **Michaels JA**, Scherberger H (in prep). Strong coding of grasp force planning and execution in macaque areas F5, M1, and AIP.

Talks:

A recurrent neural network model of the visuomotor grasp generation circuit (nanosymposium). *45th Annual Meeting of the Society for Neuroscience*. San Diego, November 16th, 2016.

Continuous decoding of hand grips with a high dimensional brain computer interface (nanosymposium, presenting in place of Andres Agudelo-Toro). *45th Annual Meeting of the Society for Neuroscience*. San Diego, November 16th, 2016.

A recurrent neural network model of the visuomotor grasp generation circuit. *Sloan-Swartz Centers for Theoretical Neurobiology Annual Meeting*. Pasadena, August 3rd, 2016. Travel grant awarded.

Probing and modeling the continuum of immediate to withheld grasping movements in the macaque fronto-parietal network (invited talk). *Neural-Prosthetic Systems Laboratory*. Stanford, May 24th, 2016.

Probing the continuum of immediate to withheld grasping movements in the macaque fronto-parietal network. *26th Neural Control of Movement Conference*. Montego Bay, Jamaica. April 26th, 2016.

Grasping with and without motor preparation (nanosymposium). *45th Annual Meeting of the Society for Neuroscience*. Chicago, October 20th, 2015. Travel grant awarded.

Laterality of grasp-related activity in macaque areas AIP and F5. *8th Primate Neurobiology Meeting*. Göttingen, March 18th, 2015.

Single trial neural correlates of grasping movement preparation in macaque areas AIP and F5. *24th Neural Control of Movement Conference*. Amsterdam, April 24th, 2014.

Abstracts:

Dann B, **Michaels JA**, Scherberger H (2016). Disentangling cue, intention, and movement information from the fronto-parietal network. *9th Primate Neurobiology Meeting*. Tübingen, Germany.

Michaels JA, Dann B, Scherberger H (2016). Emergent properties in a dynamical model of movement generation. *9th Primate Neurobiology Meeting*. Tübingen, Germany.

Dann B, **Michaels JA**, Schaffelhofer S, Scherberger H (2015). The single unit network for hand grasping has a small-world and rich-club topology with oscillators as hubs. *Neurizons*. Göttingen, Germany.

Michaels JA, Scherberger H. Laterality of grasp-related activity in macaque areas AIP and F5 (2015). *11th Göttingen Meeting of the German Neuroscience Society*. Göttingen, Germany.

Michaels JA, Scherberger H. Laterality of grasp-related activity in macaque areas AIP and F5 (2014). *44th Annual Meeting of the Society for Neuroscience*. Washington, DC, USA. Travel grant awarded.

Dann B, **Michaels JA**, Schaffelhofer S, Scherberger H (2014). Role of beta and low frequency oscillations in functional network connectivity of single units in the primate motor system. *10th Bernstein Conference*. Göttingen, Germany.

Michaels JA, Wellner B, Scherberger H (2014). Single trial neural correlates of grasping movement preparation in macaque areas AIP and F5. *7th Primate Neurobiology Meeting*. Tübingen, Germany.

Michaels JA, Wellner B, Scherberger H (2013). Single trial neural correlates of grasping movement preparation in macaque areas AIP and F5. *43rd Annual Meeting of the Society for Neuroscience*. San Diego, CA, USA. Travel grant awarded.

Michaels JA, Wellner B, Scherberger H (2013). Single trial neural correlates of grasping movement preparation in macaque areas AIP and F5. *EPFL Life Science Symposium (LSS)*. Lausanne, Switzerland.

Wellner B, **Michaels JA**, Wellner AW, Scherberger H (2013). Single trial neuronal correlates of decision-making for hand grasping in macaque area F5 and AIP. *10^{th} Göttingen Meeting of the German Neuroscience Society*. Göttingen, Germany.

Teaching experience, skills, and professional membership:

Supervision:
- Bachelor Thesis of Roman Eppinger*, Title: State space analysis and visualization of neuronal data in a delayed grasping task* (2015).
- Master Thesis of Luis Ángel Pardo Sánchez, *Title: A recurrent neural network model of bimanual coordination and interference* (2017).

Teaching assistant, Motor Systems Lecture & Seminar of Hansjörg Scherberger (2013). University of Göttingen, Germany.

Programming: Matlab, Labview, Simulink, Java, Haskell, Prolog, html/css, Python, C/C++

Software: Plexon Spike Sorting Tool, Wave_Clus, SPSS (IBM), Adobe Photoshop/Illustrator, Microsoft Office

Journals reviewed for: *J Neuroscience*.

Member of: Bernstein Association for Computational Neuroscience (2014-), Society for Neural Control of Movement (2014-), Society for Neuroscience (2013-).

Languages: English (Native), German (Fluent), Hebrew (Advanced), French (Basic).

Workshop participation:

FieldTrip Workshop (2014). Göttingen, Germany.

Tübingen International Summer School for Neuroenhancement (2013). Cloister Heiligkreuztal, Germany

NWG Practical Course – Transcranial Magnetic and Electrical Stimulation (2013). Göttingen, Germany.

BBCI Summer School: Brain-Computer Interfacing and Neurotechnology (2012). Berlin, Germany.

Nerve Cell Culture and Patch-Clamp Recording (2012). Göttingen, Germany.

Bernstein R&D Workshop Cellular Electrophysiology (2012). Heidelberg, Germany.

Rejections and failures:

Academic:
DPZ Sponsorship Award (2016). Not awarded.
GGNB Excellence Stipend (2011). Finalist, not awarded.
NSERC Undergraduate Research Award (2010). Not awarded.
Applied to Neuroscience PhD programs at eight North American institutions. No acceptances.