

Encoding, coordination, and decision making in the primate fronto-parietal grasping network

Dissertation

For the awarding of the degree

“Doctor rerum naturalium”

of Georg-August-Universität Göttingen

within the doctoral program Biology

of the Georg-August University School of Science (GAUSS)

submitted by

Benjamin Dann

from Darmstadt, Germany

Göttingen, 2017

Doctoral Thesis Committee

Prof. Dr. med Hansjörg Scherberger (First Referee, Supervisor), Neurobiology Laboratory, German Primate Center, Kellnerweg 4, 37077 Göttingen

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

Members of the examination board

Examiner: Prof. Dr. med Hansjörg Scherberger (First Referee, Supervisor), Neurobiology Laboratory, German Primate Center, Kellnerweg 4, 37077 Göttingen

Co-examiner: Prof. Dr. Alexander Gail (Second Referee), Cognitive Neuroscience Laboratory, German Primate Center, Kellnerweg 4, 37077 Göttingen

Additional Members of the examination board

Prof. Dr. Martin Göpfert, Department of Cellular Neurobiology, Schwann-Schleiden, Research Centre, Georg-August University of Göttingen, Julia-Lermontowa-Weg 3, 37077 Göttingen

Prof. Dr. Hannelore Ehrenreich, Max-Planck-Institute for experimental Medicine, Clinical Neuroscience, Hermann-Rein-Str. 3, 37075 Göttingen

Prof. Dr. Tim Gollisch, Department of Ophthalmology, School of Medicine, Georg-August University of Göttingen, Waldweg 33, 37073 Göttingen

Dr. Igor Kagan, Decision and Awareness Group, Cognitive Neuroscience Laboratory, German Primate Center, Kellnerweg 4, 37077 Göttingen

Date of oral examination: August 7th, 2017

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

Göttingen, 26th of June, 2017

Benjamin Dann

To my late mother, Dagmar

Acknowledgments

First of all, I like to thank Hans Scherberger for giving me the opportunity to work in his laboratory and supporting me with my project and analytical ideas. The scientific results which led to this doctoral thesis and several publications would not have been possible without his openness for new ideas and trust in me to successfully conduct the necessary experiments and analyses. Thanks also to Alexander Gail for the fruitful discussions with him and for his constructive criticism over the years.

My special thanks goes to Jonathan Michaels who whom I have successfully collaborated for years and finished and published several studies. Working with him was so far the best collaboration of my scientific career.

I also want to thank Matthias Dörge, Natalie Bobb, Ricarda Lbik and Sabine Borchert for their technical support and their support in animal training. It is a pleasure to work with such a friendly and supportive team.

Thank you to my former trainee Steve Suway for countless constructive discussions and debates, a nice year of cooperation, which led us become friends and helpful comments on my doctoral thesis.

Thank you to Thomas Wunderle and Kentaroh Takagaki, both my former supervisors and close friends, for many useful discussions and their constructive scientific support everytime I asked.

Special thanks also go to the three monkeys I worked with: Sally, Zara, and Bart. It is a great privilege to get the opportunity to record and analyze the brain activity of such intelligent species. This opportunity and their sacrifice as a laboratory animal motivated me to respect their needs and to treat them as good as possible.

Thank you to my mum and dad, who always believed in me and my abilities. I would have possibly not become a scientist without their support and the countless discussions with them about politics, psychology and philosophy.

My final thanks go to my wife Tanja and my daughter Julia for their love, which in particular gave me strength during the stressful phases of my doctoral thesis.

Table of Contents

1. General Introduction	11
1.1. The neuronal signal	12
1.2. Decision making	16
1.3. Encoding of information	25
1.3.2. Dynamical system perspective	29
1.4. Functional connectivity	32
1.4.1. Oscillatory synchrony.....	36
1.4.2. Network topology	39
1.5. The fronto-parietal grasping network	42
2. Original Articles and Manuscripts	45
Additional coauthored published articles	46
2.1 Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates	47
2.2 Three information subspaces explain the category-free population dynamics in the macaque fronto-parietal network	101
2.3 Probing the continuum of immediate to withheld grasping movements in the macaque fronto-parietal network	138
3. General Discussion	177
3.1. Summary	177
3.2. Outlook	180
4. Bibliography	183
5. Curriculum Vitae	193
Appendix A. Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning	201
Appendix B. Predicting reaction time from the neural state space of the premotor and parietal grasping network	224

1. General Introduction

In every higher being, the brain is the source of all higher perception, emotion, and cognition, and is the generator of behavior. In particular, the cortico-thalamic system of mammals is capable of generating highly complex cognitive and behavioral processes, up to the complexity of human cognition and behavior. The most important cellular unit of the brain is the neuron, and all brain function is thought to be generated by the greater network of neurons. However, exactly how function is generated by the network has not yet been fully understood. Donald Hebb was among the first thinkers who explicitly stated that the brain's ability to generate coherent thoughts derives from the spatiotemporal orchestration of neuronal activity (Hebb, 1949; Buzsáki, 2010). His idea was that not the neurons themselves but groups of strongly interconnected "cell assemblies" generate emergent function (Figure 1). A sequence of cell assembly activations would then in turn generate complex perceptual and cognitive processes, decisions, and, if required, behavioral output.




Figure 1: Schematic view of Hebb's neuronal "cell-assembly" idea. Intersections between arrows represent nodes and arrows represent directed links between the nodes, while the whole network represent a schematic ensemble. The number next to the links represents the order of activation within the cell assembly. Based on Hebb's writing it is unclear if nodes represent single neurons or groups of neurons. Adapted from Hebb (1949).

1. General Introduction

Today, the idea of simple linear sequences is known to be too simplified and has been extended to also take into account parallel activations and higher order interactions (Buzsáki, 2010; Cunningham and Yu, 2014), yet the basic idea of neuronal ensembles remains relevant in the field. Unfortunately, the experimental identification of cell assemblies has proven highly difficult. This is in part due to practical reasons, such as limitations in recoding neuronal activity, as well as conceptual reasons, such as our limited understanding of the computations and transformations taking place in the brain.

1.1. The neuronal signal

A wide range of techniques to record neuronal activity has been developed. These techniques can be roughly classified into three groups: 1) electrophysiological techniques such as patch clamp, which measure the direct electric currents caused by the depolarizations of single neurons, up to recordings via microelectrodes and electroencephalography (EEG), which measure the cumulative population activity of large parts of the brain, (Buzsáki et al., 2011); 2) optical techniques where neuronal activity with single cell or larger resolution is filmed through a microscope, such as two-photon calcium-imaging and voltage sensitive dye imaging (VSD) (Tsodyks et al., 1999; Harvey et al., 2012); and 3) functional imaging techniques, where brain activity is measured indirectly, such as positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) (Logothetis et al., 2001). Aside from differences in the recording techniques themselves, they also differ a lot in terms of spatial and temporal resolution (Figure 2) (Sejnowski et al., 2014). Beyond these differences, two other factors must be considered, including the spacial coverage of neuronal tissue and the tissue damage caused by the recording technique. For a comprehensive assessment of neuronal activity, the ideal recording technique would have a high spatial and temporal resolution, paired with a high spatial coverage and a minimum damage caused by the technique. Unfortunately, the ideal recording technique does not yet exist and all mentioned techniques have their advantages and disadvantages. Patch clamp measuring of the intracellular membrane current allows for single neuron recordings with maximum temporal resolution, yet patching of several neurons in parallel is difficult and mainly performed in brain slices (Perin et al., 2011). EEG recordings have a high temporal resolution, cover the whole surface of the skull, and are non-invasive, but the spatial resolution is in the range of centimeters (Buzsáki et al., 2011). Although fMRI recordings

1. General Introduction

allow for a complete three dimensional measurement of the brain, still, their temporal and spatial resolution is poor. In addition, this method is only an indirect measurement of neuronal activity (Logothetis et al., 2001). VSD imaging, despite its high spatial and temporal resolution, strongly suffers from bleaching and photo-toxicity effects (Takagaki et al., 2008).




Figure 2: The spatiotemporal resolution of neurophysiological recording techniques of the main methods available in neuroscience as of 2014. Adapted from Sejnowski et al. (2014).

However, recent developments of new optical, electrophysiological, and computational tools have made it possible to record large populations of neurons with high temporal and single cell resolution, with an acceptable amount of damage to the neuronal tissue (Buzsáki, 2004; Sejnowski et al., 2014; Yuste, 2015; Pnevmatikakis et al., 2016). In the field of optical methods, technical advantages have made it possible to even recode the whole brain of zebrafish with cellular resolution (Ahrens and Keller, 2013), albeit with low sampling rates. Nevertheless, new faster microscopes combined with deconvolution algorithms, which approximate the spiking activity from the calcium signal of individual neurons, seem promising to overcome the limitations in temporal resolution of calcium-imaging, at least to a certain extent.

Similarly, improvements to the classical microelectrode recording technique (Hubel, 1957) have also made it possible to isolate and record large populations of neurons in parallel. This is due to the development of modern computers and amplification systems,

1. General Introduction

which allow investigators to amplify and store the signals from many microelectrodes simultaneously with a good signal-to-noise ratio, and the development of multielectrode arrays to overcome the spatial coverage limitation of classical microelectrode recording (Nicolelis et al., 2003). Multielectrode arrays are simply many recording contacts combined either on one shank (Buzsáki, 2004) or in the form of many microelectrodes exiting a small plastic clip (Rousche and Normann, 1998). One problem especially for chronically implanted microelectrodes is that the brain is constantly moving relative to the skull, which precludes fixing the electrodes or arrays to the skull; this configuration is prone to creating microlesions in the neural tissue. This problem was overcome through the development of floating arrays (Rousche and Normann, 1998; Musallam et al., 2007) (Figure 3). Floating arrays are only attached to the brain and move freely with it. Electrical signals are transmitted via a small, flexible goldwire-bundle leading to a plug outside of the skull. The advantage of these arrays is that they allow for long-term, robust recording even in awake animals performing a task (Barrese et al., 2013; 2016), making them one of the preferable methods to record neuronal ensembles.




Figure 3: Picture of a floating microelectrode array with 36 electrodes, the goldwire-bundle and the plug to pick up the signals, manufactured by the company Microprobes. Adapted from <https://www.microprobes.com/>.

1. General Introduction

The signal recorded by extracellular electrodes is not necessarily straightforward to interpret, since any excitable membrane including dendrites, somas, and axons around the electrode contributes to the recorded signal (Buzsáki et al., 2011). The amplitude and the frequency of the measured voltage change depend on the superimposed activations of all the surrounding neuronal compartments with decaying influence over distance. Still, due to differences in the temporal dynamics of pre- and postsynaptic processes (initial segment and axonal potentials, and dendritic and soma potentials, respectively), these two components can be extracted from the signal. The postsynaptic integration processes, called local field potentials (LFP), occur at slower time scales in the range of $<100\text{Hz}$, while the presynaptic spiking activity is thought to be in the range of $>300\text{Hz}$, which allows for a clean separation of these two components by band-pass filtering (Figure 4).

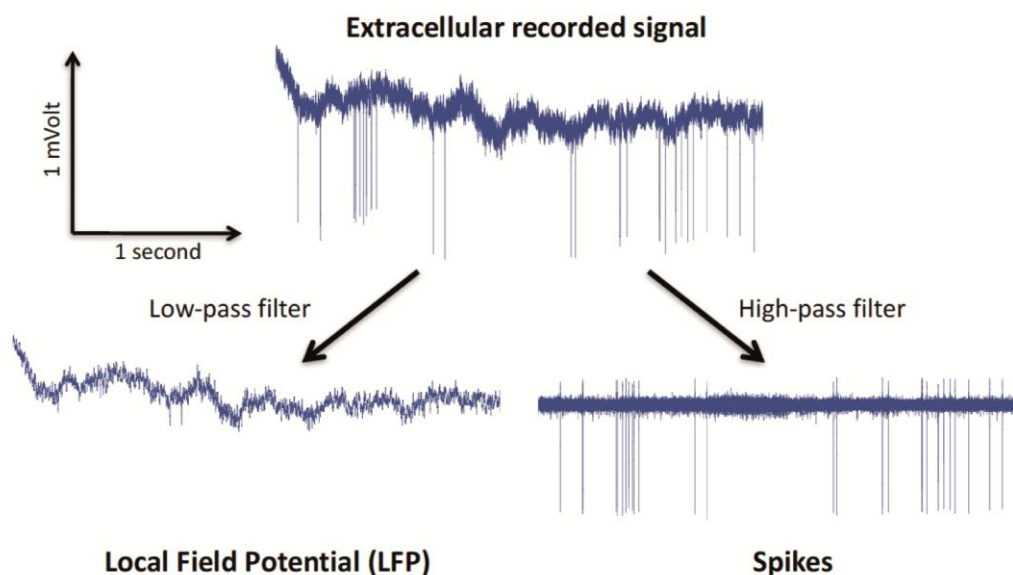


Figure 4: Extracellular recorded signal from ventral premotor cortex. The signal was low-pass filtered with a 100Hz Butterworth filter (4th order, non-causal) to extract LFP activity and high-pass filtered with a 300Hz Butterworth filter (4th order, non-causal) to extract spiking activity.

However, recent studies have shown that the energy of spikes leaks into the LFP down to 20Hz (Waldert et al., 2013; Schomburg, 2015), making the interpretation of the LFP more difficult. Even ignoring the bias from concurrent spiking, the LFP is difficult to interpret, since it reflects a nonlinear mixture of the surrounding postsynaptic processing, which in turn is dependent on the level of network synchrony, the cellular architecture, and volume conductance effects (Buzsáki et al., 2011). In contrast, the spiking activity of individual neurons is thought to be an all or nothing potential, which even allows the isolation of individual neurons via spike-sorting algorithms (Quiroga et al., 2004; Rossant et al., 2016).

1. General Introduction

Spike-sorters make use of the temporal and, in the case of multitrodes (several close recording sites which can pick up extracellular spiking activity of the same neuron), also of the spatial differences of spikes from different neurons recorded from the same site. Since the voltage deflection caused by spikes of distinct neurons is similar across occurrences, spikes from the same neuron should cluster together based on extracted features from individual spikes such as wavelet coefficients or principle components (Figure 5a). Yet, small changes of the electrode position to the nearby neurons have large nonlinear effects on amplitude and shape of the recorded extracellular spikes (Gold, 2006) (Figure 5b), which among other things makes proper and careful spike-sorting very important.

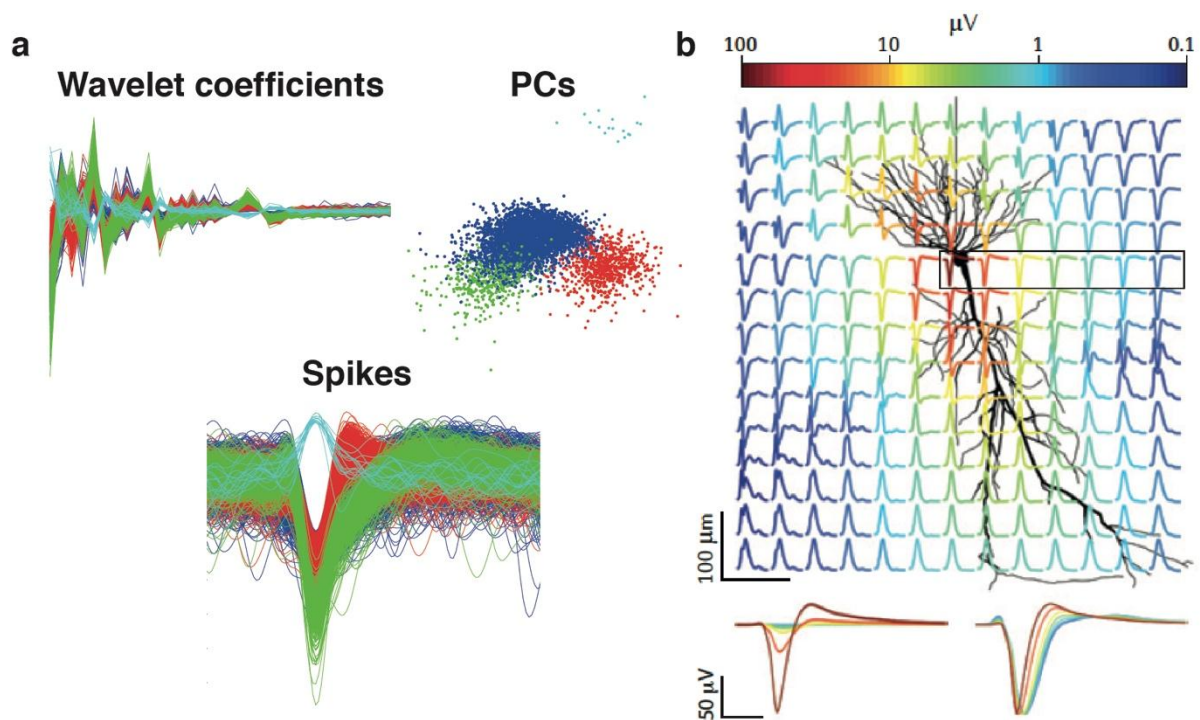


Figure 5: (a) All recorded spikes from one channel aligned on their maximum peak or trough and either shown as decomposed into Wavelet coefficients by Wavelet transform, projected onto the first three principle components (PCs) estimated by principle component analyses or as individual spike waveforms over time. The different colors reflect the four units extracted by spike-sorting. (b) Shape and amplitude of the extracellular recorded spike waveform is dependent on the recording side. The magnitude of the spike is normalized to its minimum and maximum. The peak-to-peak voltage range is indicated by the colour of the traces. Note that the spike amplitude decreases rapidly with distance from the soma. Adapted from Buzsaki et al. (2012)

1.2. Decision making

The possibility to extract spikes of individual neurons even while animals perform a behavioral task has led to large number of studies correlating activity of individual neurons with behavior. One intensively studied behavior is decision making, since whether or not we react to a stimulus or intention involves a decision process. Decision making is regarded as the process of flexibly selecting or reacting to external sensory inputs or to internal drives

1. General Introduction

(Freedman and Assad, 2016). In the field of systems neuroscience, the goals in regard to decision making are: to find neuronal correlates of decision making, to develop ideas about the mechanism of the underlying decision process, to develop models resembling the decision process, and ideally being able to causally influence the process. In the last decades, many different processes have been suggested for many different types of decision making. However, which aspects such as rules, rewards, goals, and certainty are included in the decision process is still a matter of debate (Miller, 2000; Andersen and Cui, 2009; Shadlen and Kiani, 2013; Freedman and Assad, 2016). Classical studies suggested the prefrontal cortex (PFC) as the center of decision making (Miller, 2000), yet many different areas were determined to be involved in decision making, including many parietal areas (Andersen and Cui, 2009), even V4 and middle temporal cortex (MT) from the visual system (Shadlen and Kiani, 2013; Siegel et al., 2015), as well as subcortical structures such as the superior colliculus, the basal ganglia, the thalamus and the cerebellum (Andersen and Cui, 2009; Shadlen and Kiani, 2013). One useful classification to better understand decision processes is to distinguish different kinds of decision making, such as perceptual decision making and internally driven decision making.

Perceptual decision making means that a decision has to be made about the perception of a stimulus or a property of a stimulus. One classic paradigm for perceptual decision making is the random dot motion task, where monkeys have to distinguish the direction of motion in a cloud of moving dots and signal their choice by making a saccade to the left or to the right (Newsome et al., 1989). Crucially, the percentage of dots moving in one direction (called the level of coherent motion) was varied from full up to zero percent coherent motion. The smaller the percentage of coherent motion, the more difficult it was for the monkey to choose the right direction. The firing rate of individual neurons recorded in area MT matched the corresponding psychometric function, which quantifies the ratio of choice in one direction relative to the other as a function of coherent motion. Even a weak but reliable correlation with the trial-to-trial variability was found (Sugrue et al., 2005). The activity of neurons in the lateral intraparietal area (LIP) was later found to reflect mainly the decision to make a saccade to the left or right by a ramping increase in firing rate for the neuron's preferred target. This observation led to the idea that evidence is accumulated up to some threshold in the brain (Shadlen and Kiani, 2013). It was posited that once the threshold is reached the movement is elicited. Since evidence of a fixed threshold could not

1. General Introduction

be found, the model was later extended to a more dynamical threshold, which depends on the level of certainty of the monkey about the direction of motion (Kiani and Shadlen, 2009). However, even based on the assumption that the thresholds change over time, this theory seems to be too simplified and conflicts with other findings. Neurons in LIP were found to display strong, prolonged responses related to working memory, and even a saccade could be elicited while keeping another movement target in memory (Rishel et al., 2013). Both of these findings are difficult to unite with the accumulated evidence model.

Another classic paradigm for perceptual decision making is the somatosensory flutter task, in which monkeys had to report which of two temporally separated vibration stimuli delivered to one finger was higher in frequency (Romo and Salinas, 2003). This task involved several processing steps. The initial stimulus had to first be perceived, then kept in memory until the second stimulus was given and perceived, and finally the information could be combined to form a decision. Single neuron activity was recorded from many different areas for this task, including S1, S2, PFC, and medial premotor cortex (MPC). By using a multivariate linear regression model, it was possible to relate the activity of each recorded neuron to either the first stimulus, the second stimulus, and the overall decision if modulation for either parameter was strong enough (Hernandez et al., 2010). Interestingly, while neurons in S1 were only significantly modulated for the two stimuli during their presentation, neurons in the other three regions additionally showed significant modulation related to working memory and the decision. Memory related modulation was strongest in PFC and decision related activity was strongest for PFC and MPC. This clear overlap and similarity of neuronal responses across areas suggests a graded and not area-specific representation of task parameters and the decision process. Nevertheless, the presumptions made by this model could lead to a strong preselection of neurons and as a consequence a false interpretation of the data. In particular, the often described mixed selectivity of neurons (Mante et al., 2013; Rigotti et al., 2013; Womelsdorf and Everling, 2015) for many task parameters in PFC is not accounted for by this model. Further, these findings are rather descriptive and do not offer a mechanism for decision making.

Another interesting group of perceptual decision making tasks are delayed match to category tasks (DMC). In DMC tasks, monkeys were trained to group a large, continuously varying set of visual stimuli into two categories and report their decision by a saccade to the corresponding target (Freedman and Assad, 2016). These tasks allow the dissociation of

1. General Introduction

neurons modulated by stimulus features from neurons modulated by categorical membership. In one version of the task, monkeys were trained to group images of continuous mixtures of cats and dogs into two arbitrary categorical groups (Freedman et al., 2001). Neurons in PFC were predominantly category selective, while neurons recorded in inferior temporal cortex were predominantly visual feature selective. Yet, small numbers of neurons in the two areas also showed opposite selectivity, respectively.

In summary, perceptual decision making tasks have provided insight into the temporal integration of ambiguous or noisy stimuli, which led to the accumulated evidence model, the implication of different areas across cortex in representing stimulus features as well as decision related activity, and the idea that a set of continuously varying stimuli can be arbitrary categorized based on behavioral demands. However, it can be argued that perceptual decisions are merely happening on the level of sensory perception. This would mean that in case of an ambiguous stimulus, noise either from the stimulus source itself, or in early sensory perception processing (e.g. in the retina, in the somatosensory receptor cells, or early on in the cortical sensory processing) can cause a bias towards one of the two sensory categories. Thus, from that point on, the whole transformation up to a final movement would be identical to that associated with an unambiguous stimulus (Andersen and Cui, 2009).

In contrast, internally driven decisions, which are also referred to as “free choice,” are decisions where the sensory evidence is not in question. For such decisions, the choice of a final action is assumed to be based on the integration of different factors such as rules, goals, rewards, costs or others (Andersen and Cui, 2009; Cisek, 2012). Yet, the integration of many behaviorally relevant factors raises several important questions: How and where are these relevant factors represented and integrated? And which factors are really represented in the brain? In asking the second question, we also ask the central question of whether there is truly a representation of an abstract decision variable in the brain. Or, can the whole decision process be explained by a stimulus selection process directly being transferred into movement preparation (Andersen and Cui, 2009)? Several models have been proposed for internally driven decision processes based on the empirical evidence of single neuron recordings from different areas, while monkeys performed different kinds of decision making tasks. Three prominent models of these processes are the good-based model, the action-based model, and the distributed consensus model (Cisek, 2012). Note that the

1. General Introduction

distributed consensus model is an extension of the action-based model, and both are based on the idea of biased competition between potential movement or action plans (Figure 6).

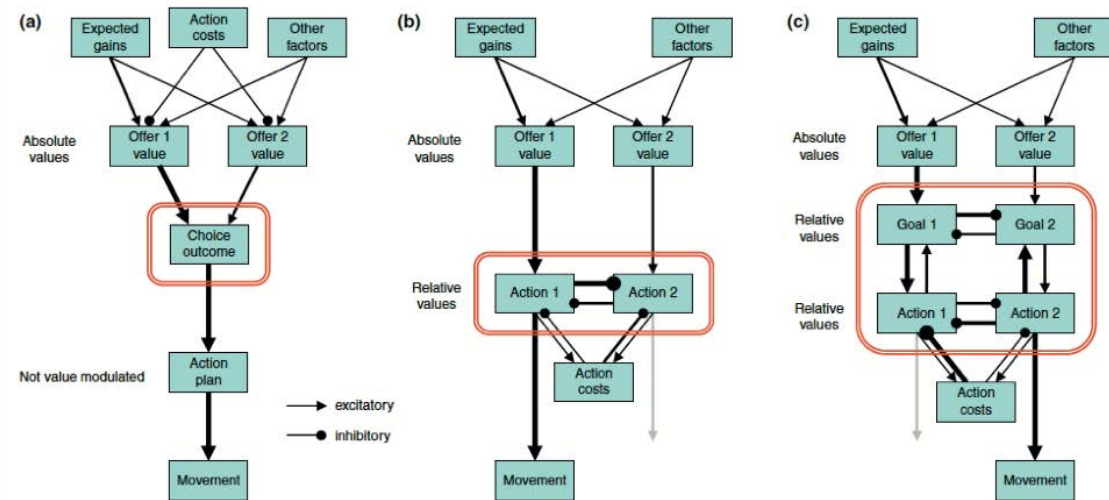


Figure 6: Three schemes for three different internal decision models. The red box highlights where and how the decisions are made. Arrows represent transformations and competitions with their strength indicated by line thickness. (A) A good-based model, in which decisions are made by comparing representations of offer values and only afterwards transformed into an action plan. (B) An action-based model, in which decisions are made through a biased competition between action plans. (C) A distributed consensus model, in which decisions are made through competition at multiple levels representing different factors such as goals and actions. Adapted from Cisek (2012).

The good-based model is derived from economic theory and experimental psychology (Padoa-Schioppa, 2011). It suggests that all relevant factors for a decision such as action cost and expected reward are separately integrated into abstract absolute values for each possible option. The comparison of the absolute values determines the decision outcome, which is then transformed into a movement plan. The activity of neurons in orbitofrontal cortex and ventromedial prefrontal cortex was found to be correlated with this suggested absolute value (Padoa-Schioppa, 2011). However, several studies have shown neuronal activity in frontal and parietal areas represents movement plans before the final decision is made (Cisek and Kalaska, 2005; Scherberger and Andersen, 2007; Klaes et al., 2011). The first evidence for a neuronal representation of motor plans preceding a final decision was provided by a study where monkeys were trained to perform a delayed center out reaching task while neuronal activity in dorsal premotor cortex (PMd) and M1 was recorded (Cisek and Kalaska, 2005). The final goal was to reach out for one of 8 cued targets after a certain delay. However, the cue for the correct reach direction was split into a spatial cue indicating two opposite targets, which was given first, and a color cue indicating the target to choose, given at a later time point of the task. Interestingly, after the ambiguous spatial cue was given, neurons spatially tuned for both targets became active, and only after

1. General Introduction

the color cue was presented just neurons tuned for the final direction remained active. These findings can hardly be explained by a good-based model, which does not include the coexistence of several action plans; these findings led to the idea of a competition between action or motor plans taking place within the sensorimotor network (Cisek, 2012). The action-based model reflects these findings by suggesting that the value of possible actions is immediately translated into parallel existing motor plans and the decision is made as a biased competition between them (Figure 6). Further evidence for this model was given by a study where only one spatial cue was displayed and at a later time point of the task a rule cue was given instructing the monkey to reach for the target or in the opposite direction (Klaes et al., 2011). This task allowed to disentangle the neuronal representation of the visual target from neuronal activity representing movement plans. After the spatial cue was presented, neurons recorded in PMd and the parietal reach region (PRR) tuned for both movement directions became active, while in case of rule dependent motor planning only neurons tuned for the spatial target location should have become active.

Nevertheless, the action-based model fails to explain choices which do not result in movements, while the good-based model seems to be better suited for that. The distributed consensus model offers a possible solution by extending the biased competition of the action-based model into two or more levels (Figure 6) (Cisek, 2012). Instead of just having a competition between motor plans in sensorimotor areas, an additional competition takes place at the level of abstract goals in presumably anterior portions of the PFC. Due to the strong reciprocal connectivity of sensorimotor areas and more anterior parts of the PFC, a common distributed consensus resulting in a decision could be made with different influences at all levels. The biased competition occurring at each different level does not need to agree, since only one common decision is made as a result of the processing over all levels.

However, the distributed consensus model is rather abstract and does neither explain how information is exactly encoded nor transformed into the final decision. This is in contrast to the accumulated evidence model, which nevertheless is too simple for many of the required transformations and observed results, as mentioned before. Yet, most of the assumptions and results on which the model is based rely on analyses of single neuron tuning analyses, which in all of the mentioned studies only explain a fraction of the neuronal population activity. Further, the assumed tuning function often only roughly matches the

1. General Introduction

neuron's response (Churchland and Shenoy, 2007; Churchland et al., 2010). Even if we assume that tuning properly reflects the encoding of information, it does not tell us much about the generation of the encoding and necessary transformations that result in a decision.

A recent study recording single neuron activity of PFC while monkeys performed a mixed rule based perceptual decision task proposed a new model for decision making taking all the previously mentioned points into account (Mante et al., 2013). Monkeys first received a rule in the form of a visual cue instructing them to decide based on either color or motion, immediately followed by a random dot motion pattern, as mentioned above, with different, independent levels of color and motion coherence. They found that the activity of many individual neurons was remarkably complex and their contribution to behavior could not be directly understood. Instead of analyzing the activity of individual neurons separately, they considered them as one interconnected assembly or population. On the level of the population, the complex response of individual neurons unfolded as one dynamic process evolving through independent subspaces for rule, motion, color and the decision variable. Interestingly, they found each kind of information represented simultaneously in the same neuronal population. Additionally, both kinds of sensory information (color and motion) were present regardless of which rule cue was given, arguing against any preselection mechanism. Instead, the different kinds of information were independent on the level of the population readout. Intriguingly, a trained recurrent neuronal network (RNN) model could reproduce the population dynamics (Figure 7). These findings suggest that the encoding and transformation of information for this task can be well characterized as a dynamical evolution of the whole neuronal population. However, the four subspaces were not derived from the neuronal activity, but assumed *a priori* based on the task design, which risks not properly capturing the flexibility of the population response. Still, this criticism also holds for analyses based on individual neuron tuning.

1. General Introduction




Figure 7: A RNN model for decision making including context dependent input selection and integration. The RNN model receives independent motion, colour and contextual inputs and generates a decision variable resembling the recorded neuronal activity of PFC. The network is fully recurrently connected, and each neuron receives all three inputs. The network output resembling the decision is generated as a linear, weighted sum over the responses of all neurons (red arrows). The network was trained with back-propagation to make a binary choice and initialized with random synaptic strengths. Adapted from Mante et al. (2013).

Another study analyzing the neuronal population of the posterior parietal cortex (PPC) of rats performing a multisensory perceptual decision task was well in line with the previously described study (Raposo et al., 2014). They also found that neuronal activity could be best explained as a dynamic process evolving through independent subspaces for modality and decision variables, which in this case can be assumed to be identical to movement preparation or planning. Representation of different kinds of information was intermingled not only in the activity of individual neurons, but also randomly distributed across the whole neuronal population. Additionally, they also found neuronal activity to span a different subspace during active movement. Active movement control is another important aspect which must be considered when analyzing decision related activity in sensorimotor areas in order to form a complete picture of the underlying processes, since there is growing evidence that these areas are also involved in active movement control (Churchland et al., 2010; 2012; Menz et al., 2015; Elsayed et al., 2016). The mixed selectivity of neurons for many kinds of information was even confirmed across 7 different cortical areas (Siegel et al., 2015) for a nearly identical task to Mante et al. 2013. Interestingly, the information for task, motion, color, and choice was present in a graded manner in all 7 areas, including the visual areas V4, and inferior temporal cortex, the lateral intraparietal area (LIP), PFC and the frontal eye field, strongly arguing against any preselection mechanism and in favor of a flexible, distributed decision process. Choice information was highest at the time point of movement initiation in all areas including FEF, which is known to

1. General Introduction

be involved in movement generation. This finding suggests that the observed choice information is at least to a certain degree movement related and not representing an abstract choice variable.

The results described up to this point have mainly focused on the types of information encoded in neuronal populations, as well as the temporal development and transformation of this information (with the exception of the RNN model). However, another important aspect of study is the selective communication and coordination of information that takes place between neurons and different brain areas during decision processes (Pesaran, 2010). One way to investigate close-range selective communication processes is to analyze the LFP. As discussed, the LFP mainly represents a nonlinear mixture of the surrounding postsynaptic processing, yet it also reflects the level of synchrony in the nearby neuronal population, since any nonsynchronous component would simply average out.

A study in which monkeys were trained on the flutter task (described above) while LFP activity in S1, S2, PFC, MPC, and M1 was measured, showed that oscillatory synchronization in the beta-band (15-30Hz) reflected the dynamics of decision making (Haegens et al., 2011). Local beta-band synchronization during the decision period of the task was reflective of the decision outcome and not the stimulus information, with the strongest effects found in MPC and weaker effects present in all other areas. In a study where monkeys had to reach to three targets either in a clear instructed order or in an order chosen by free choice showed increased long range beta synchronization between PRR and PMd for the free-choice condition (Pesaran et al., 2008). In another study analyzing cross area synchronization based on LFP activity in S1, S2, PFC, MPC, and M1 while monkeys performed the same flutter task, strong delta-band (1-4Hz) synchronization during the decision process was observed (Nácher et al., 2013). Similar to the findings for beta-band synchronization, delta-band synchronization was modulated by the decision process across nearly the whole network, indicating long range delta-band synchronization as an important communication mechanism during decision making.

The research summarized in this section demonstrates that many different cortical areas are involved in many kinds of decision paradigms. The areas involved range from early sensory areas such as S1, secondary sensory areas such as S2, V4, and MT, up to many parietal and frontal areas, where the strongest correlates of decision processes were found.

1. General Introduction

Interestingly, similar areas across cortex were identified to be involved in internal and perceptual decision making, although perceptual decision making could possibly be explained by a noise based visual selection mechanism early on in processing. This speaks in favor of one common distributed decision network as suggested by the distributed consensus model. Still, the distributed consensus model is rather abstract, while some evidence was found that decision-related neuronal activity can be understood as a dynamical process on the population level evolving through different subspaces, at least within distinct areas. Additionally, synchronization of neuronal populations in the beta and delta range within and between areas seems to be important for the selective communication underlying decision processes. However, a comprehensive picture of decision making does not yet exist. Crucial reasons for this are that it is still unclear: (1) how and which information is encoded in the neuronal population, (2) how the information flow is coordinated in the neuronal population within and between areas and, based on that, (3) how information is transformed.

1.3. Encoding of information

The way in which information is thought to be encoded by the neuronal population cannot be uncoupled from the history of neuroscience. The idea that the neuron is the functional and structural unit of the brain, called the neuron doctrine, is credited to Cajal and Sherrington (Yuste, 2015). While Cajal was the first anatomist who identified individual neurons, proposing them as the structural unit of the brain, Sherrington was the first to suggest the neuron is also the functional unit by finding receptive fields on the skin. Analysis of single neuron properties was significantly advanced by the invention of the microelectrode (Hubel, 1957). The microelectrode allowed for the isolation of single neurons, as mentioned before, yet until the development of newer recording techniques only a few neurons could be recorded simultaneously. The responses of individual neurons were found to be correlated with many visual features, as well as of other sensory modalities. Even behavioral features including overt movement parameters were found to be correlated with individual neuronal responses, which led to the idea that individual neurons represent information about perception, cognition, and behavior. These findings formed the foundation of the representational framework (Buzsáki, 2010; Yuste, 2015).

1. General Introduction

In contrast to the representational framework, neuronal network models assume that groups of strongly interconnected neuronal assemblies generate emergent function (Hebb, 1949). Although this idea was proposed as early as the 1940s and the first neuronal network models were developed soon after, the implementation of complex network models has only recently become possible through the development of modern computers (Sussillo, 2014; Yamins and DiCarlo, 2016). Furthermore, the ability to record large populations of neurons in parallel either with newly developed optical or electrophysiological tools (Sejnowski et al., 2014; Pnevmatikakis et al., 2016, Rousche and Normann, 1998) has enabled the development of novel population analyses such as dimensionality reduction methods (Cunningham and Yu, 2014). These new methods have led to a paradigm shift from single neuron to population analyses, allowing for exploratory analyses of the population structure even on the single trial level. The new insights based on neuronal network models as well as on population analyses have resulted in the proposal of the dynamical system perspective for neuronal population activity (Shenoy et al., 2013; Yuste, 2015).

It is important to state that the representational and dynamical system view are compatible to a certain extent. While the former describes the information encoded by individual neurons in terms of physical parameters of the environment, the latter assesses the population activity from the perspective of its output and the need to interact with the environment, which allows for a high degree of similarity between the two frameworks.

1.3.1. Representational view

According to the classical representational framework, the firing rate of each neuron is described as a function of correlation with (or “tuning” to) various parameters. Tuning is defined as a systematic modulation of the firing rate of an individual neuron in relation to the systematic variation of a perceptual, cognitive or behavioral parameter. Based on this framework, the neuronal correlates of certain parameters of objects or movements are presumed to have clear boundaries between them in agreement with the neuronal substrate (Buzsáki, 2010). The idea is that elementary parameters of objects or movements are bound together by the network of neurons in a meaningful way to perform the required cognition or movement. However, an unsolved problem associated with this idea is that the

1. General Introduction

elementary parameters for this process are not universal properties of movements or the perceived world but created by the interaction with the environment.

Historically, the idea that individual cortical neurons are tuned to external parameters gained prominence due to recordings in the primary visual cortex (V1) (Hubel and Wiesel, 1968). In early studies performed by Hubel and Wiesel, a huge variety of visual stimuli were presented to anesthetized cats. One groundbreaking discovery attributed to this work was that the firing rate of neurons in V1 was only enhanced when the stimulus was presented at a certain spot in space relative to the eye. This observation marked the discovery of visual receptive fields in V1. Shortly thereafter, they found that neuronal responses systematically varied with the orientation of a presented bar of light, which was the first discovery of tuned neurons in the cortex (Figure 8).

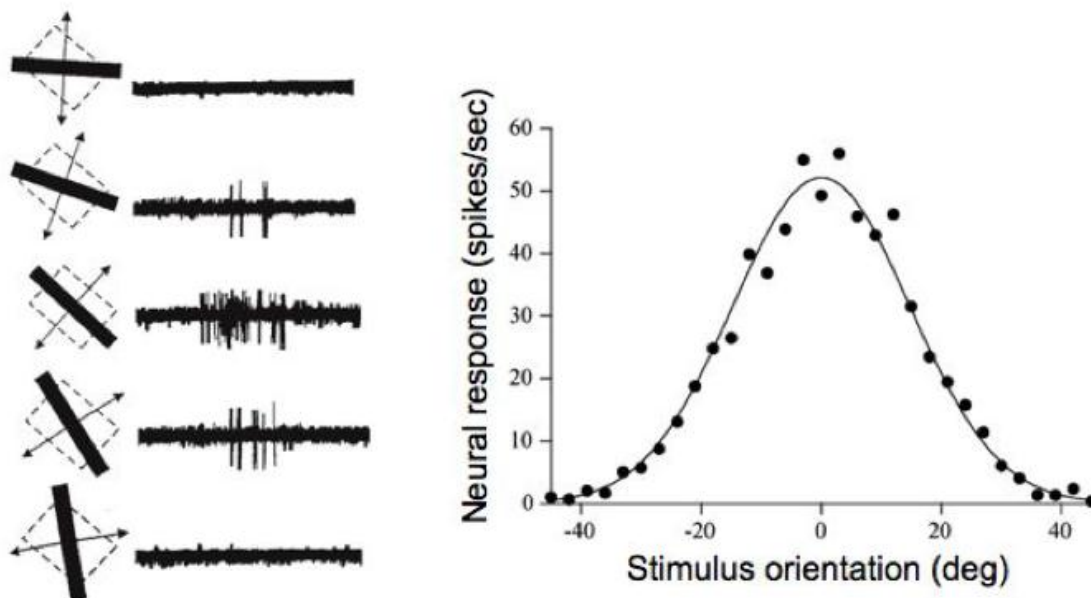


Figure 8: Orientation selectivity of a simple cell recorded in area 17 of the cat (corresponds to area V1 of the macaque monkey). Depending on the orientation of a light bar projected on a screen and moved through the receptive field of the neuron, the neuron responds with different firing rates. The orientation-dependent modulation of firing rate can be described with a with a Gaussian or cosine fit as shown on the right. Adapted from Hubel and Wiesel al. (1968).

The finding of the receptive field together with orientation tuning became the corner stone of many studies describing the activity of individual neurons from the representational view. The representational framework remains the basis of many studies today. This framework has been especially successful in describing visual processing, but has also been useful in describing movement related activity of individual neurons. The most famous example is the activity of neurons in M1, PMd, and PRR while monkeys perform a center out reaching task,

1. General Introduction

which can be well described by tuning for the reach direction (Georgopoulos et al., 1982; Klaes et al., 2011), and is similar to the orientation tuning observed in V1. Further, neurons of the fronto-parietal network were classified into visual, visuomotor, and motor related based on which aspect of a reach-to-grasp task their activity was modulated by (Murata et al., 2000). More recent studies, conducted on monkeys passively viewing many different pictures, showed that individual neurons in several locations of the temporal cortex (the so-called “face patches”) only increased their firing rate in response to pictures of faces (Freiwald et al., 2009). Individual neurons of the different face patches were found to be tuned for facial features ranging from simple orientation up to a complete generalized viewpoint in the face patch highest in the hierarchy (Freiwald and Tsao, 2010). In a similar experiment on human patients, neuronal activity was recorded in the medial temporal lobe and neurons were found which fired selectively for pictures of specific movie actors (Quiroga et al., 2005). In the motor system, including premotor and motor cortex, more modern approaches have tried to describe the activity of individual neurons as a combination of many parameters such as velocity, position, acceleration, and occasionally jerk (Todorov, 2000). Yet, even these “complex kinematic” models only coarsely matched the observed complexity of individual neuron responses (Churchland et al., 2012). Additionally, there remains little agreement regarding even the basic parameters relevant to responses in the motor system (Shenoy et al., 2013).

A common problem of the representation framework is that individual neuron tuning analyses often only explain a fraction of the recorded neuronal population and the assumed tuning function often only roughly matches the neuron’s response (Churchland et al., 2010; Mante et al., 2013). This leaves a large proportion of neuronal variance unexplained and calls into question the validity of this framework for explaining neuronal processing. Assuming that the tuning information of individual neurons is indeed behaviorally relevant, and that neurons exist that are tuned e.g. for particular individuals, how is this information linked to an equally complex network of movement related neurons? Since this would involve the unlikely necessity of dynamically linking or unlinking millions of different neurons within different networks, an encoding of information on the population level seems to be much more likely (Yuste, 2015).

1. General Introduction

1.3.2. Dynamical system perspective

Considered from the dynamical systems perspective, neural circuit function is assumed to arise from the activation of the whole network of neurons to generation an output, which cannot be understood by studying one neuron at a time (Yuste, 2015). Instead, the firing rate of each recorded neuron is considered as one dimension in a state-space, and the population firing rates over time form a trajectory through this space (Shenoy et al., 2013). The evolution of neural population activity should be best captured in terms of dynamical rules by which the current state, its input, and possibly some noise cause the next state. Dynamical rules can be attractors, such as fixed points in state-space to which the population activity moves towards, converges, or rotates around, either across time or even across conditions. Attractors can also be seen as emergent states or subspaces guiding the evolution of the population activity. The emergent states may not be recognisable by looking at responses of individual neurons, since they arise from the interaction of the whole neuronal population. However, there are at least two reasons to assume that the number of subspaces through which the population activity evolves is smaller than, and distributed across, the number of neurons present in one area or even across areas. The first reason is the tight but widespread recurrent connectivity within areas and across cortex (Smith and Kohn, 2008; Markov et al., 2014) and the second reason is the need for a representation that is robust against any kind of external or internal distortion of the conducted neuronal process (Shenoy et al., 2013). In this sense, it is essential to find the underlying dimensionality or number of subspaces of the population response which governs the trajectory through state space. Many dimensionality reduction methods for large-scale neuronal recordings have been introduced with different advantages and disadvantages (Cunningham and Yu, 2014). One of the most frequently used methods is principle component analyses (PCA), which is an unsupervised method based on the covariance between all pairs of neurons. PCA can be used to project the full neuronal state space into a lower number of orthogonal dimensions which explain most of the covariance in the data in descending order. However, since PCA is based on covariance, it captures neuronal variance of all kinds, including firing rate differences between neurons and, even worse, probabilistic spiking variability between single trials (Cunningham and Yu, 2014). For this reason, PCA is usually applied to normalized conditionwise trail-averaged data (Churchland et al., 2010; 2012; Elsayed et al., 2016). One unsupervised covariance based dimensionality reduction

1. General Introduction

method suitable for single trial analyses is Gaussian process factor analysis (GPFA) (Yu et al., 2009). Yet, GPFA assumes an explicit noise model, which could be inaccurate and could result in arbitrarily broad temporal smoothing. Another problem with covariance based dimensionality reduction methods such as PCA and GPFA is that time-shifts in firing between neurons as well as graded transitions in the activation of neurons across the population can result in an artificially high number of estimated dimensions (Novembre and Stephens, 2008; Kobak et al., 2016). In contrast, supervised methods exist that are more robust to time-shifts and graded transitions in the population activity, and some of them are applicable to single trial activity and rely on dependent variables. In most cases, the dependent variables are parameters of the performed task such as stimulus color or the final decision (Mante et al., 2013; Kaufman et al., 2015). Due to this constraint, there is the danger that the estimated dimensions do not explain a meaningful part of the neuronal population variance or miss important dimensions. Three commonly used supervised methods are support vector machines (SVM), linear discriminant analyses (LDA), and multivariate linear regression (Mante et al., 2013; Cunningham and Yu, 2014; Raposo et al., 2014). Basically, the first two find the projection which best separates the predefined groups of points from each other, while the third method estimates a linear fit of the activity of all neurons onto the dependent variable. Despite the pitfalls and restrictions of neuronal population dimensionality reduction methods, they hold potential for providing many new insights into the encoding and transformation of information in the cortical neuronal population.

In a few relevant studies, monkeys were trained to perform a large variety of different straight and curved reach movements following a delay period, while populations of neurons were recorded in M1 and PMd. These studies have helped better understand movement preparation and movement generation. The first finding was that activity of individual neurons was complex and multiphasic during the movement epoch and heterogeneously distributed across the neuronal population, which could not easily be explained by the representational framework (Churchland and Shenoy, 2007). Surprisingly, by estimating the directional tuning of the whole population of neurons during the preparatory and movement periods, it was shown that tuning was only weakly correlated between these epochs, speaking in favor of an independent population encoding of information for the two periods (Churchland et al., 2010). However, using 10 PCA-based dimensions of the population preparatory activity, movement activity could be better

1. General Introduction

predicted than with any of the tested representational models. These findings suggested that preparatory activity could be an initial state of a dynamical system whose evolution controls movement. Through the development of a method called jPCA, which is an extension of PCA that projects the population response onto planes that capture rotational variance, it was possible to show that only a few dimensions could capture a considerable amount of neuronal population variance in the form of rotational population dynamics, with the preparatory activity as an initial state as suggested before (Churchland et al., 2012). Analyses of the single trial trajectories of the preparatory activity using GPFA revealed that the closer the single trail trajectory was to the “ideal” initial subspace, the faster a movement was initiated (Afshar et al., 2011). A recent study showed that neuronal population activity during preparatory and movement period evolves through independent but linked subspaces (Elsayed et al., 2016). Independent subspaces for different stimulus features and choice or preparatory activity were also found in rat PPC and monkey PFC as described before (Mante et al., 2013; Raposo et al., 2014). It is important to mention here again that the neuronal contributions to the different subspaces were randomly distributed across the whole recorded population of neurons in rat PPC.

A strong indication that the network of neurons is well described as a dynamical system could be found by generating a model where we know by definition that it is a dynamical system. In order to be a valid model, given the same inputs, we should observe outputs closely resembling the recorded neuronal responses. Trained RNNs were shown to be suitable models for this approach, and were found to resemble the dynamics of PFC on the population level for a decision task (Mante et al., 2013), as described before in the decision making section. Intriguingly, in a recent study where a RNN received recorded preparatory activity as input, and was trained to produce the subsequently recorded muscle activity, and was additionally regularized, the dynamics of the RNN during movement resembled the dynamics seen in the recorded neuronal population at both the single-neuron and population levels (Sussillo et al., 2015). The results further strengthen the idea that motor cortex can be well described as a dynamical system generating muscle patterns.

The notion of a global, rather than local, encoding and transformation of information raises the question of how these processes are coordinated across brain structures. For a limited network with a limited number of conditions, a dynamical system, modeled by a trained and regularized RNN receiving the same inputs and generating muscle patterns,

1. General Introduction

offers a remarkably parsimonious solution for the coordination of information. However, the coordination of information across many brain areas for a nearly endless number of different behaviors presumably requires a more complex coordination mechanism, which makes it essential to study the communication structure of neuronal populations within and across areas using functional connectivity measures.

1.4. Functional connectivity

Functional connectivity is defined as any quantifiable interaction estimated based on the synchronization of parallel recorded neuronal signals (Bastos and Schoffelen, 2016). In most cases, functional connectivity is estimated for the same condition to exclude influences from the average conditionwise activation profile, which is assumed to not reflect synchronization processes (Cohen and Kohn, 2011). Functional connectivity, in contrast to anatomical connectivity, does not necessarily imply direct synaptic connections since it can also reflect synchronization processes of distant neuronal populations.

Many different metrics have been introduced to estimate functional connectivity, which can be categorized into groups on various levels. Two ways to categorize them are based on whether the method quantifies the direction of interaction or is undirected, and if the method is based on model assumptions of interaction or is model-free (Figure 9).

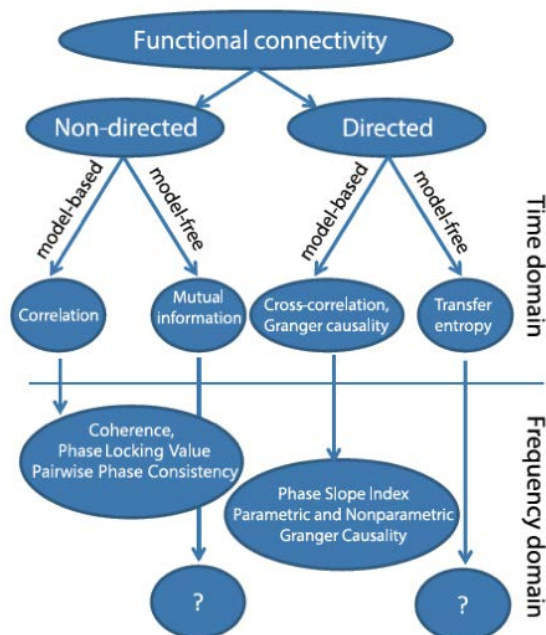


Figure 9 A taxonomy of popular methods for quantifying functional connectivity. The methods are grouped first based on whether they quantify the direction of interaction or not and secondly whether they are model based or not. The classification is done for functional connectivity methods in the time domain. In case a frequency domain adaptation of the time domain method exists, it is shown below and underlies the same categorization into groups as the equivalent time domain method. Adapted from Bastos et al. (2016)

1. General Introduction

The most well known and commonly used method is the Pearson correlation coefficient (Bullmore and Sporns, 2009; Cohen and Kohn, 2011), which measures the non-directed interaction between pairs of neuronal signals under the model assumption of a linear interaction. The advantage of the Pearson correlation coefficient is that it is fast and easy to calculate and gives very robust results. However, the strength and even the detectability of functional interactions estimated with the Pearson correlation coefficient varies considerably with the chosen window size of temporally averaged neuronal signal, and entails the danger of missing fast interactions that switch sign over time (König et al., 1995; Cohen and Kohn, 2011). A model-free alternative to the Pearson correlation coefficient is mutual information, which in comparison takes longer to calculate and is more vulnerable to noise, since no assumption is made about the noise of the signal (Kraskov et al., 2004). Still, the Pearson correlation coefficient can also be used to calculate directional connectivity by calculating cross-correlation histograms (CCHs) between pairs of neuronal signals, which also solves the problem of window size (Kohn and Smith, 2005). CCHs are estimated by incrementally shifting the time series of the neuronal signals with respect to one another and calculating the Pearson correlation coefficient for every time lag. The timing of the peaks and troughs of the CCHs give information about the directionality of the interaction. However, the interpretation of peaks or troughs at each time lag between the two signals is difficult. Originally it was thought that no time shift in correlation between the two signals could indicate common input from another source (Ts'o et al., 1986), yet recent studies on complex network models have suggested that zero-lag peaks instead reflect reciprocal connectivity (Vicente et al., 2008; Gollo et al., 2014). Another group of methods to estimate directed functional connectivity is based on linear auto-regressive models such as generalized linear models (Okatan et al., 2005) and Granger causality (Dhamala et al., 2008; Seth, 2010). Granger causality allows for the separate estimation of functional interactions from signal x to signal y and vice versa, yet it can only be estimated properly with a high signal-to-noise level. Finally, a model-free method to estimate directed functional connectivity is transfer entropy (Lindner et al., 2011), but similar to mutual information it takes longer to calculate and it is more vulnerable to noise than the linear methods.

Another important aspect of functional connectivity estimation is the kind of signals between which the interaction is calculated. Depending on the signal, different assumptions have to be made with clear implications for the results and which method is best to choose.

1. General Introduction

Crucial factors are spatial and temporal resolution and the signal-to-noise ratio of the different signals (Bastos and Schoffelen, 2016), as well as whether the signal is continuous such as the BOLD signal, the signal from EEG and MEG, and the LFP, or binary signals such as spike trains of single neurons or a group of neurons (Cohen and Kohn, 2011; Bastos and Schoffelen, 2016).

In general, the estimation of functional connectivity of continuous signals is much easier because most metrics can be directly applied. However, since the LFP as well as the EEG and MEG signal on a coarser scale predominantly represents a nonlinear sum of synchronization of the surrounding postsynaptic processing, as mentioned before, the strength of interaction of these signals is difficult to interpret. Nevertheless, LFP recordings in particular have a comparably good signal-to-noise ratio, which makes them a preferred choice of signal for many interaction studies, especially Granger causality estimations of directed interaction (Roelfsema et al., 1997; Fries et al., 2001; Womelsdorf et al., 2007; Salazar et al., 2012; Bastos et al., 2015). For BOLD signals measured with fMRI, which normally have a temporal resolution in the range of seconds, directional measures do not make much sense, since most neuronal interactions are known to take place in the millisecond range (Buzsáki, 2010). One disadvantage of all continuous signals mentioned so far is that they each represent an average signal of a neuronal population. Assuming that the interactions within the population are heterogeneous they could be averaged out or be strongly distorted on the population level resulting in an inaccurate estimation of functional connectivity. This makes single neuron functional connectivity analyses essential to understand the coordination taking place within a neuronal assembly (Yu et al., 2008; Nigam et al., 2016).

The most common single neuron signals are spike trains recorded from microelectrodes and isolated by spike sorting algorithms. Spike trains are binary signals (Okatan et al., 2005; Cohen and Kohn, 2011) for which the estimation of interaction is more complicated compared to continuous signals. Single neuron calcium-signals recorded with optical methods are also basically binary signals, since the calcium-signals have to be considered as low-pass filtered spike trains, which can be recovered by deconvolution methods (Pnevmatikakis et al., 2016). One disadvantage of analyzing spike trains is that only the spike events are known states while all other time points are hidden states of the neuronal activity, resulting in a low signal-to-noise ratio (Cohen and Kohn, 2011). Due to this

1. General Introduction

reason, Pearson correlation analyses of very short time windows can lead to an underestimation of the interaction strength between neurons. Additionally, average neuronal spike rates for different species, independent of the behavior or area where they were recorded from, were shown to be log-normally distributed across the population, spanning around three orders of magnitude (Buzsáki and Mizuseki, 2014). The large differences in firing rate result in large differences in signal-to-noise ratio between neurons. Unfortunately, this in turn results in an underestimation of interaction strength for neurons with low average firing rates (la Rocha et al., 2007; Cohen and Kohn, 2011). This bias cannot be corrected for and results in an average firing rate dependency of all functional connectivity measures between spike trains. The estimation of directed interaction by Granger causality is strongly dependent on the signal-to-noise ratio, as already mentioned before. Although implementations of Granger causality for spike train interactions do exist, their usage is questionable due to the extreme heterogeneity of firing rates and the modeled data they were tested on had unrealistically high and homogeneous average firing rates (Kim et al., 2011; Quinn et al., 2011). Pearson correlation analyses are comparably robust to rate differences, which makes this classical method still one of the best choices for single neuron functional connectivity estimations. In particular, CCHs are still commonly used for spike train based interaction estimates (Fujisawa et al., 2008; Smith and Kohn, 2008; Ecker et al., 2010; Ramalingam et al., 2013). However, CCHs only allow for pairwise interaction estimates of spike trains, which cannot reveal more complex multivariate interactions of several spike trains. Nevertheless, this possibility was tested in a study where many neurons were recorded in parallel, and the multivariate Ising model (based on the principle of maximum entropy) as well as CCHs were applied to estimate neuronal interactions (Yu et al., 2008). The direct comparison of both methods showed that nearly all interactions could be reliably captured by pairwise CCHs.

Out of the many studies estimating functional connectivity, studies focused on the coordination of information can be roughly grouped into studies analyzing the kind of synchronization between areas or neurons, suggesting oscillatory synchrony as a crucial mechanism for dynamic network coordination (Fries, 2005; 2015) and studies analyzing the network topology of the interactions of many areas or neurons (Bullmore and Sporns, 2009; Schröter et al., 2017). A review of studies in each category follows.

1. General Introduction

1.4.1. Oscillatory synchrony

In the last decades, oscillatory synchronization between single neurons as well as neuronal populations has been found in many studies, suggesting oscillatory synchronization as an important mechanism involved in dynamic network coordination (Engel et al., 2001; Fries, 2009; Engel and Fries, 2010; Buzsáki and Wang, 2012). Oscillatory synchronization in neuronal populations has been described in different distinct frequency bands, such as delta (1-4 Hz), theta (4-8 Hz), alpha (8-15Hz), beta (18-35Hz), and gamma (40-100Hz) (Engel and Fries, 2010), raising the question of whether these different frequency bands are coupled to distinct perceptual, cognitive, or motoric functions and whether they have different anatomical origins.

The first specific oscillatory synchronization processes in the gamma-band (40-120Hz) were described in a series of anesthetized cat experiments, while animals were passively observing different visual stimuli and neuronal activity was recorded in V1. Gamma-band synchronizations between neurons as well as neuronal populations were found to be stimulus specific (Gray and Singer, 1989). A few years later, long range synchronizations (>2mm) between neurons in V1 of one hemisphere as well as between the two hemispheres were found to be almost always in the gamma-band (König et al., 1995). Experiments conducted on awake monkeys that had to attend one of two visual stimuli on a monitor showed that neurons recorded in V4 within the receptive field of the attended stimulus showed increased gamma-band synchronization with their surrounding population (Fries et al., 2001). Interestingly, lower frequency synchronizations (< 17Hz) were also present, showing modulation in the opposite direction. In another study, the same modulation of gamma-band synchrony was found between FEF and V4 (Gregoriou et al., 2009). Recent experiments with monkeys performing a similar task revealed that populations of neurons in V1 within the receptive field of the attended stimulus were selectively synchronized in the gamma-band with populations in V4, while populations in V1 within the receptive field of the non-attended were not synchronized with V4 (Bosman et al., 2012). By using Granger causality, they could show that the direction of the synchronization was mainly from V1 to V4 and not the other way around, suggesting gamma-band synchronization as a bottom-up coordination mechanism in the visual system. The electrocorticogram grid arrays used in this study actually spanned large parts of the visual system, including parietal and frontal areas, allowing for a more systematic assessment of the information coordination across cortex.

1. General Introduction

Granger causality analyses of the directed functional connectivity across 8 areas revealed that bottom-up information flow is coordinated by gamma-band and theta-band synchronization, while the top-down information flow is coordinated by beta-band synchronization, with V1 as the lowest and parietal area 7a as the highest in the hierarchy (Bastos et al., 2015).

The importance of beta-band synchronization originating from parietal areas was in fact shown many years earlier in a study where monkeys had to press and hold a hand lever for variable amounts of time. The investigators showed directed functional connectivity via Granger causality from several parietal areas to motor areas (Brovelli et al., 2004). In a study where monkeys had to perform a mixed delayed center-out reach and saccade task while neuronal activity was recorded from PRR, two important findings regarding beta-band synchronization were established (Scherberger et al., 2005). First, the level of beta-band synchronization of neurons with their surrounding population was selective for the preparation of reach compared to saccade movements and, secondly, the level of beta-band synchronization was predictive of the task period, suggesting beta-band synchrony to be involved in intention or movement preparation coordination. These findings are well in line with the described results in the decision making section showing that beta-band synchronization of neuronal populations selectively reflected the decision outcome or intention, which is presumably the same as movement preparation (Pesaran et al., 2008; Haegens et al., 2011), as mentioned before. Also, findings from more recent studies where monkeys had to perform coordinated reach and saccade movements while single neuron and LFP activity were recorded simultaneously in PRR and LIP are in accordance with the idea that beta-band synchrony is involved in the coordination of movement intention or preparation (Dean et al., 2012; Wong et al., 2016). They found that only neurons synchronized with the larger populations in both areas were predictive of the movement initiation of coordinated reach and saccade movements. However, low frequency synchronizations of populations across areas seem to be involved in movement intention coordination as well (Nácher et al., 2013), as described before in the decision section.

It is important to state that there are many more studies describing selective coordination mechanisms by oscillatory synchronization. Many experiments have been conducted on rats performing a vast assortment of different tasks while activity in the hippocampus, the entorhinal cortex, and different cortical regions was recorded, with

1. General Introduction

findings corroborating a coordinative role of gamma-band and theta-band oscillatory synchrony (Buzsáki, 2010; Buzsáki and Wang, 2012; Schomburg, 2015). The studies presented here were selected with regard to coordination of information across cortex spanning perceptual processing, decision making, and behavior generation. A possible interpretation of all presented results is that gamma-band synchronization and possibly theta-band synchronization coordinate the bottom-up attention control originating from the visual areas. In contrast, beta-band and possibly low-frequency synchronization could serve as coordinative mechanisms for intention or top-down control of the information flow with beta-band synchronizations originating from parietal areas. Additionally, beta-band synchronization could possibly be the coordinative mechanism of a putative distributed consensus across cortex, as suggested for decision making (Cisek, 2012).

Two important questions remain unanswered. What is the advantage of oscillatory synchronization as a coordinate mechanism? And, how is the information flow coordinated by this synchronization? It is important to stress that so far, no common agreement or causal proof exists to answer these two questions. However, a convincing answer to the first question is the idea of feedforward coincidence detection (Fries, 2009). The number of synaptic inputs to a neuron is large (1000- 10000) and the postsynaptic potentials triggered by spikes are known to decrease rapidly after initiation, which effectively leaves only a few milliseconds for arriving spikes to be integrated to elicit a spike from the target neuron. If neurons are oscillatory synchronized to each other, then their spikes have on average a greater impact on their targets. The advantage of such a mechanism is not only a reduction of energy cost and an increase in spike efficiency, but also a rhythmic gain modulation. A rhythmic, synchronized activation of a population of neurons results in phases of high excitability when all neurons fire and phases of low excitability in between. As a consequence, the amount of excitation necessary to elicit spikes from the target neurons is rhythmically modulated or, in other words, the gain is modulated. This allows for a selective amplification of inputs from one group of neurons to another group of neurons, by simply changing the phase of synchrony of the target neural population to be in phase with one group of neurons and out of phase with the other group. This highly flexible mechanism of selective communication, which results in a coordination of information flow, is called communication through coherence and is a possible answer to the second question posed above (Fries, 2005).

1. General Introduction

Most studies to date have analyzed the kind of synchronization thought to coordinate information flow between pairs of neurons, areas, or local populations. Yet, the brain or brain areas are a strong interconnected network on the anatomical as well as functional scale (Berger et al., 2007; Bullmore and Sporns, 2009; Markov et al., 2014), which makes it essential to analyze the functional network structure to understand the coordination of information flow. However, due to the above-mentioned possibility that aspects of the communication can average out at the level of population signals, it is essential to analyze functional interaction on the level of single neurons to understand the formation of potential ensembles.

1.4.2. Network topology

The ability to record many neuronal signals simultaneously (e.g. with recently developed optical, electrophysiological, and computational tools (Buzsáki, 2004; Sejnowski et al., 2014; Yuste, 2015; Pnevmatikakis et al., 2016), see above), has allowed investigators to estimate functional networks using functional connectivity measures (Bastos and Schoffelen, 2016). However, identifying the functional connectivity of hundreds or thousands of neuronal signals presents a problem for analyzing these networks in terms of their structure and their organizational principles, referred to as network topology. Many useful analyses for this purpose were developed by mathematicians from the field of network science, which was only recently established in the late 1990s based on graph theory (Watts and Strogatz, 1998; Bullmore and Sporns, 2009).

In the first study of the field of network science (Watts and Strogatz, 1998), three important network measures were defined. The first two are the cluster coefficient, which measures interconnectivity between direct neighbors of one node of a network, and the shortest path length, which measures the minimum number of nodes which have to be passed to get to another node. A simple regular network where each node is connected to its four spatial neighbors has a high average cluster coefficient but a long average shortest path length. In contrast, a random network has a small average cluster coefficient and a short average path length. An interesting finding of this study was that, by randomly switching pairs of connections of a regular network, an intermediate state of high average cluster coefficient and small average shortest path length was present before the network became random. Networks that combine both are referred to as small-world, which was the

1. General Introduction

third defined network measure. Interestingly, the anatomical single neuron network of *C. elegans* and who-played-with-whom network of Hollywood actors both turned out to be small-world. After this striking finding, many more topological principles were found and described which seem to be common principles of many natural networks and led to the definition of complex networks (Barabási and Oltvai, 2004; Barabási, 2009; Bullmore and Sporns, 2009). Natural networks were found to have a modular topology, which means that groups of nodes within a network are more strongly interconnected with each other than with the rest of the network (Ravasz et al., 2002). The importance of individual nodes for the network communication or the network coherence of natural networks can be measured by centrality metrics. Natural networks were shown to have heavy-tailed centrality distributions, with a small number of nodes connecting the network and coordinating the network function (these nodes are called “hubs”), while the majority of nodes are only of minor importance for the overall network function (van den Heuvel and Sporns, 2013). The first described and simplest measure of centrality is degree centrality, which is defined as the number of connections per node (Barabási et al., 1999; Jeong et al., 2000). A more global aspect of centrality is captured by betweenness centrality, an index of the number of shortest paths from all single units to all others that pass through that node (Freeman, 1977). In some networks, hubs exhibit a strong tendency to link to each other, forming a so-called rich-club (Colizza et al., 2006). This property can be measured by a rich-club coefficient that expresses the tendency of highly connected hub nodes to show above-random levels of interconnectivity.

Network analyses of anatomical and functional inter-area brain networks measured with tracers, EEG, MEG, or fMRI also revealed them to be topologically organized, as with complex networks (Bullmore and Sporns, 2009). The regional brain networks of humans and monkeys were found to have a modular and small-world topology (Hilgetag et al., 2000; Stephan et al., 2000; Bullmore and Sporns, 2009) Further, the centrality distributions of areas were found to be heavy-tailed with hub areas (Achard et al., 2006; Honey and Kötter, 2007; Honey et al., 2007), which were strongly interconnected as a rich-club coordinating global brain communication (Harriger et al., 2012; van den Heuvel et al., 2012).

However, functional network topology analyses of more localized neuronal signals of mammalian brains are lacking in the literature. Three studies analyzing the single neuron functional network of organotypic slices of rat brain showed that the single neuron

1. General Introduction

functional connectivity topology was modular, with functional hub neurons organized as a rich-club coordinating the network communication (Bonifazi et al., 2009; Shimono and Beggs, 2014; Schroeter et al., 2015). Nevertheless, the neuronal activity of oranotypic slices is altered compared to the intact brain (Steriade, 2001). Many of the original connections and many neurons in the slice die due to the slicing procedure, no natural sensory inputs are received by the neuronal network, and plasticity effects after the extraction of the slice even further change the neuronal connectivity. Only three studies analyzed the functional network topology of single neurons recorded in the intact brain. The first study was performed on anesthetized cats passively viewing visual stimuli while many neurons were recorded in parallel in V1, showing a small-world topology of functional connectivity (Yu et al., 2008). The second study was performed on awake monkeys also viewing visual stimuli, while neurons were recorded in parallel in V1. In contrast to the first study, these investigators suggested that single neuron functional small-world topology is an artifact of distance-dependent functional connectivity (Gerhard et al., 2011). However, the number of recorded neurons was small, and even that small number was most likely due to massive oversorting, questioning the validity of the results from this study. The last and most recent study was performed on awake rats under uncontrolled behavior while neurons were recorded in medial to lateral orbitofrontal cortex. It was reported that the functional single neuron topology could be described as a rich-club (Nigam et al., 2016). Yet, the uncontrolled behavior utilized in that study did not allow for a separation of behaviorally driven common neuronal network activations, such as those triggered by different movements or from synchronization processes reflecting the coordination of network interaction. In summary, it remains unclear how the functional network of local neuronal populations or single neurons is topologically organized within and across areas in order to coordinate information flow.

Since it is so far not feasible to record the majority of neurons in the brain in parallel or of high numbers of areas, an important question is: what is an interesting cortical network from which to record many neurons in parallel? The network should be suitable for analyzing single neuron functional network topology and oscillatory synchronization process in regard to coordination of information flow, as well as the encoding and transformation of information from perception to behavior.

1.5. The fronto-parietal grasping network

One potential way to select a cortical network to record from is based on the behavior which is generated and controlled by it. Ideally the behavior is measurable and quantifiable, such as overt motor movements. Grasp movements are one of the most important for interacting with our environment on an everyday basis. The cortical network which generates and controls grasp movements includes as some of its core areas of the anterior intraparietal area (AIP), the ventral premotor cortex (F5), and the primary motor cortex (M1). AIP and F5 are part of the fronto-parietal network and are known to be strongly reciprocally connected, as are F5 and M1 (Luppino et al., 1999). Inactivation studies of area AIP and F5 showed deficits in pre-shaping of the hand during grasping, confirming them to be involved in grasp generation and control (Gallese et al., 1994; Fogassi et al., 2001). Several studies have been conducted on monkeys trained to do visual fixation tasks as well as visually guided delayed or non-delayed grasping tasks while single neuron activity was recorded in AIP and F5. These studies showed that neurons of both areas were modulated for visual object discrimination (Murata et al., 2000; Janssen and Scherberger, 2015), movement preparation (Baumann et al., 2009; Fluet et al., 2010), and movement related processing (Menz et al., 2015). These findings are well in line with the information representation of neurons recorded from the fronto-parietal networks for saccadic eye movements (LIP and FEF) (Freedman and Assad, 2006; Siegel et al., 2015) and for reach moments (PRR and PMd) (Gail, 2006; Churchland et al., 2010; 2012). The presence of visual and preparatory activity within the same network led to the assumption that AIP and F5 play an important role in visuo-motor transformation (Janssen and Scherberger, 2015), also well in line with findings from studies of the fronto-parietal saccadic eye movement and reaching network. Strong evidence for this idea was provided by two studies showing that, in the fronto-parietal grasping network including M1, visual information was found to be most strongly represented in AIP, followed by F5, and movement related information was most strongly represented in M1, followed by F5, and most weakly in AIP (Schaffelhofer et al., 2015; Schaffelhofer and Scherberger, 2016). These findings suggest a graded representation and transformation of neuronal information across the areas, again in agreement with studies of the fronto-parietal saccadic eye movement network (Siegel et al., 2015). Interestingly, information relevant to reach and eye position was found to be encoded by the population of neurons in F5 and AIP (Lehmann and Scherberger, 2013), further suggesting a graded representation for the controlled motor

1. General Introduction

moment across the whole fronto-parietal network. A potential reason for this overlapping representation, which was also found for eye and reach representation and coordination in LIP and PRR (Andersen and Cui, 2009), could be that flexible coordination of eye, reach, and grasp movements is necessary in everyday life. The high degree of similarity between the fronto-parietal saccadic eye movement, reaching, and grasping networks, as well as the overlap between them, suggests AIP and F5 are also involved in decision making.

Taken together, neurons in the fronto-parietal grasping network are selective for visual, preparatory, and grasp movement related information and are involved in the transformation from visual to preparatory activity, from preparatory to movement activity, and very likely also in the decision making process that are part of these transformations. Conveniently, the involvement of this network in grasp movement preparation and generation allows for the direct measurement and quantification of the output of the system. Furthermore, there is evidence for beta-band synchronization originating from parietal regions such as AIP, which potentially is an important coordinative mechanism involved in decision making and movement intentions, as mentioned before. However, the exact interplay of all these processes is currently not well understood (Janssen and Scherberger, 2015), positioning the fronto-parietal grasping network of macaque monkeys as a suitable structure to study the encoding, transformation, and coordination of information and decision making. Such studies will provide the characterizations needed to better understand the formation of functional neuronal ensembles.

In order to explore these processes leading to clearer comprehension of functional neuronal ensembles within the fronto-parietal grasping network large populations of neurons of this network were recorded in parallel as a databasis of this thesis, while monkeys performed different delayed grasping tasks. Four monkeys were trained on two different tasks and were chronically implanted with four to six floating microelectrode arrays with 36 electrodes (Figure 3) in AIP, F5 and in one case M1 (two per area). The signal of all electrodes were recorded in parallel and as a basis of all performed analyses large populations of neurons were extracted via spike-sorting algorithm (Figure 5).

In chapter 2.1 the coordination of the information flow across the fronto-parietal single neuron network was analyzed by estimating the the directed functional connectivity between all pairs of single neurons. The kind of synchronization process was analyzed together with the functional network topology allowing for a unifying view of both aspects.

1. General Introduction

In chapter 2.2 the encoding of information across the neuronal population of AIP and F5 was analyzed, while two monkeys performed a mixed instructed and free-choice delayed grasping task. Analyses of the classical representational framework were contrasted with population analyses in line with the dynamical system perspective. Furthermore, a regularized RNN model was trained for the same conditions to produce muscle activity for the performed grip types. This model offered a biological plausible explanation for decision related transformation of information within the fronto-parietal grasping network.

In chapter 2.3 the neuronal population dynamics across AIP and F5 of two monkeys were analyzed of the transition between immediate and delayed grasp movements. Population analyses by using dimensionality reduction techniques revealed how dynamical as well as static aspects of movement preparation can be encoded simultaneously in different dimensions in the same neuronal state space.

2. Original Articles and Manuscripts

This chapter contains the following research articles and manuscripts:

2.1. Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates

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.

Published in: Elife 5:2870. Doi: 10.7554/elife.15719

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

2.2. Three information subspaces explain the category-free population dynamics in the fronto-parietal network

Dann B*, Michaels JA*, Agudelo-Toro A, Scherberger H *Equal contribution

Manuscript in preparation

Author contributions: B.D., and A.A.-T. collected the data. B.D., J.A.M., and H.S. designed and performed research. B.D., and J.A.M. analysed the data. B.D., and J.A.M. wrote the manuscript. All authors revised the manuscript.

2.3. 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). *Equal contribution

Manuscript in preparation

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

Please note that a previous version of this chapter was already published by Jonathan A. Michaels as part of his dissertation with the title: Towards population coding principles in the primate premotor and parietal grasping network. However, in collaboration of Jonathan A. Michaels with me substantial changes have been made in the manuscript, Jonathan A. Michaels agrees to the usage of this chapter in my dissertation.

Additional coauthored published articles

Appendix A 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.

Published in: PLoS Comput Biol 12:e1005175–22. Doi: 10.1371/journal.pcbi.1005175

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

Appendix B 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.

Published in: Journal of Neuroscience 35:11415–11432. Doi: 10.1523/JNEUROSCI.1714-15.2015

Author contributions: B.D., and R.W.I. collected the data. J.A.M., B.D., 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.

Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates

Benjamin Dann¹, Jonathan A. Michaels¹, Stefan Schaffelhofer¹, & Hansjörg Scherberger^{1,2}

¹Neurobiology Lab, German Primate Center, Kellnerweg 4, Göttingen 37077, Germany

²Faculty of Biology, Georg-August University Göttingen, Germany

Acknowledgements: We thank R. Lbik, N. Bobb, and S. Borchert for animal support, M. Dörge for technical support, and S. Suway and T. Wunderle for help with the development of data analysis and fruitful discussions. This work was supported by the Deutsche Forschungs Gemeinschaft (SCHE 1575/1-1 & 3-1 and SFB889, C9), the Bundesministerium für Bildung und Forschung (BCCN-II, 01GQ1005C) and the European Union (FP7-611687, NEBIAS).

2.1. Uniting functional network topology and oscillations

Abstract

The functional communication of neurons in cortical networks underlies higher cognitive processes. Yet, little is known about the organization of the single neuron network or its relationship to the synchronization processes that are essential for its formation. Here, we show that the functional single neuron network of three fronto-parietal areas during active behavior of macaque monkeys is highly complex. The network was closely connected (small-world) and consisted of functional modules spanning these areas. Surprisingly, the importance of different neurons to the network was highly heterogeneous with a small number of neurons contributing strongly to the network function (hubs), which were in turn strongly inter-connected (rich-club). Examination of the network synchronization revealed that the identified rich-club consisted of neurons that were synchronized in the beta or low frequency range, whereas other neurons were mostly non-oscillatory synchronized. Therefore, oscillatory synchrony may be a central communication mechanism for highly organized functional spiking networks.

2.1. Uniting functional network topology and oscillations

Introduction

Perception, cognition, and movement are generated by the functional interaction of neuronal circuits. In order to understand the basis of these processes, especially in highly complex networks such as the primate brain, it is essential to know their network structure, termed topology. Graph theoretical approaches have enabled analysis of the brain's network topology (Watts and Strogatz, 1998; Bullmore and Sporns, 2009). Using such approaches in EEG, MEG, DTI or fMRI studies, anatomical regions have been grouped into functional and anatomically strongly connected modules, which are segregated from each other (Bullmore and Sporns, 2009). Still, every region can be reached by bypassing a few others (small-world), a topology which is robust and allows efficient information processing (Hilgetag et al., 2000; Stephan et al., 2000; Bullmore and Sporns, 2009). A few regions of the brain are highly connected and centrally located within the network (van den Heuvel and Sporns, 2013a) (hubs) as well as strongly connected to each other (van den Heuvel et al., 2012) (rich-club). This rich-club forms a global communication pathway across the network, thereby cross-linking segregated modules (van den Heuvel and Sporns, 2013b).

However, single neurons and their functional network topology are the fundamental computational structure of the primate brain. While neuronal modules, hubs, and rich-club organization has been shown in organotypic slices of rats (Bonifazi et al., 2009; Shimono and Beggs, 2014; Schroeter et al., 2015), hardly anything is known about single neuron network topology in the intact brain during behavior. Limitations in recording high number of single neurons in parallel, incorporating distance-dependent connectivity, and addressing subsampling and firing rate biases makes it difficult to assess these networks. Only small-world topology has been debated (Yu et al., 2008; Gerhard et al., 2011) and rich-club topology has been shown recently in mice (Nigam et al., 2016).

Equally important to topology is the mechanism which coordinates and synchronizes neurons during cognitive or perceptual processes. Previous research has revealed oscillatory synchrony in time as a crucial feature of functional coordination (Fries, 2009; Buzsáki and Wang, 2012; Womelsdorf et al., 2014). Different distinct frequency bands for information transmission and functional network coordination have been identified, such as gamma (40-100Hz) and theta (4-8 Hz) in the visual areas and up to frontal cortex for coordinated attention selection (Roelfsema et al., 1997; Bosman et al., 2012; Gregoriou et al., 2012), and beta (18-35Hz) and delta (1-4 Hz) in fronto-parietal regions for network coordination during

2.1. Uniting functional network topology and oscillations

decision and working memory processes (Brovelli et al., 2004; Pesaran et al., 2008a; Haegens et al., 2011; Salazar et al., 2012; Nácher et al., 2013). Recently, gamma and theta oscillations have been proposed as feedforward communication frequencies across large parts of the visual network, while beta oscillations has been proposed for feedback communication (Bastos et al., 2015). However, firing rate correlations have also been found, independent of oscillatory synchronization, to be of importance for communication in the behaving brain (Fujisawa et al., 2008; Smith and Kohn, 2008). Yet, how functional network topology, described by graph theoretical approaches, relates to oscillatory and non-oscillatory synchronization remains unclear. This question must be answered at the level of single neurons, where oscillatory synchrony can be distinguished from non-oscillatory synchrony.

Here, we recorded in parallel and assessed functional connectivity and network topology from a large number of single neurons (48 to 149 per session) from the primate grasping circuit (Luppino et al., 1999), including the ventral premotor (F5), primary motor (M1), and anterior intraparietal (AIP) cortex of three behaving macaque monkeys. Across the three cortical areas we found modular, small-world topology with a clear presence of hubs that were organized as a rich-club. Moreover, rich-club hub neurons predominantly spiked and communicated by oscillatory synchrony in the beta and low frequency range, while the remainder of the network predominately communicated by non-oscillatory synchrony, suggesting that oscillatory synchrony is a central coordination mechanism for functional network topology.

Results

The current study includes 12 recording sessions from three macaque monkeys (M: 3, S: 6 and Z: 3). We recorded from the grasping motor network, including part of the ventral premotor (F5), anterior intraparietal (AIP), and additionally from primary motor (M1) cortex area for monkey M (Schaffelhofer et al., 2016) (Supplementary Table 1). To engage the grasping motor network, monkeys performed a visually-cued delayed grasping task in which the monkey grasped a handle with one of two different grasp types (Michaels et al., 2015) (Figure 1A,B; see Materials and Methods). An average number of 570 trials (SD: 177) were recorded in each session.

2.1. Uniting functional network topology and oscillations

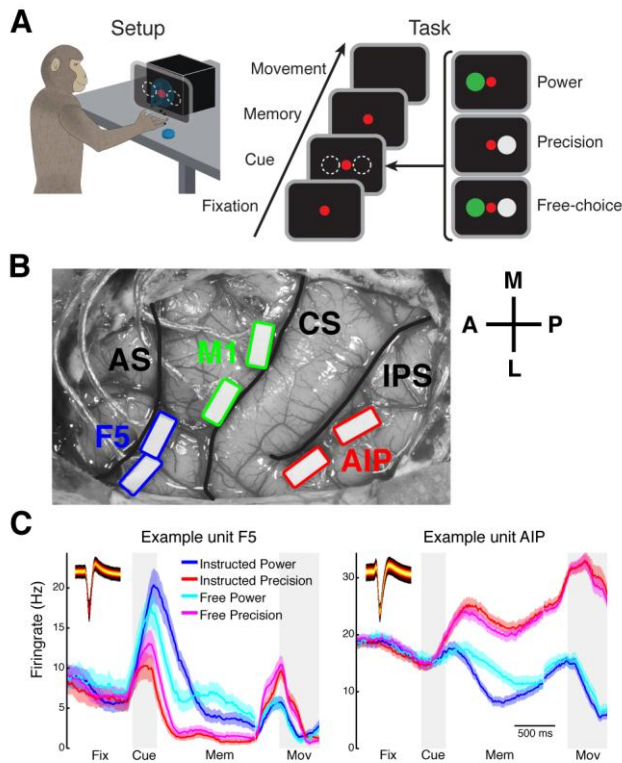


Figure 1. Task design and array implantation. (A) Choice/no-choice task. Setup: Monkeys were cued to grasp a target (handle) with one of two different grip types displayed on a monitor appearing superimposed on the handle. Task: Monkeys had to fixate a red disk for 600-1000ms (Fixation), followed by a cue period of 300ms (Cue). Then, either ('Power') a green disk was presented on the left indicating a power grip, ('Precision') a grey disk on the right indicating precision grip, or ('Free-choice') both disks were presented indicating a free-choice between both grips. After the cue a memory period followed (duration: 1100-1500ms) before the fixation dot was turned off (go-signal) indicating the monkey to execute the grasp movement (maximum duration:1000ms). (B) Electrode array implantation of monkey M with 6 floating microelectrode arrays (FMAs) in areas AIP, F5, and M1. Arrays were implanted at the lateral end of the intraparietal sulcus (IPS) in AIP, in the posterior bank of the arcuate sulcus (AS) in area F5, and in the anterior bank of the central sulcus (CS) in the hand area of M1. (C) Average firing rate across trials of two example units from area

F5 (left) and AIP (right). Each colored line corresponds to the mean activity of one condition. Line shadings represent standard error. Inlays shows the corresponding waveforms displayed as density plots.

In each area, recordings were obtained from two floating microelectrode arrays (FMAs), for a total of 64 channels (32 per microarray) per area (Figure 1B; see Materials and Methods) from which an average of 88 single units (SD: 32) were recorded in parallel. All recorded single units were modulated by the epochs of the task or the grasp types, clearly indicating the behavioral relevance of the performed task to the detected single units (Figure 1C). Nevertheless, in agreement with previous findings (Buzsáki and Mizuseki, 2014), firing rates of individual units were relatively stable for different behavioral states of the task following an approximate log-normal distribution (Figure 1-Figure Supplement 1).

Functional connectivity

The functional connectivity between all simultaneously recorded units of the grasping network was estimated by calculating cross-correlation histograms (CCHs) (Figure 2A, Figure 2 - Figure Supplement 1,2; see Materials and Methods), one of the few methods also allowing analyses of the frequency domain (Bastos and Schoffelen, 2016) (see below). It is important to stress that the functional connections we describe here do not necessarily represent monosynaptic connections, but merely the influence of one unit onto another. For

2.1. Uniting functional network topology and oscillations

each neuron pairing one single CCH was estimated over all task epochs and grasp types, since we were interested in the general network interaction and not grasp type or time specific modulations of the network. A general problem of all connectivity measures is common drive to the network, such as stimulus- or movement-locked, but not pairwise, correlations, causing an overestimations of connections. We corrected these biases by subtracting surrogate CCHs (Figure 2 - Figure Supplement 1A).

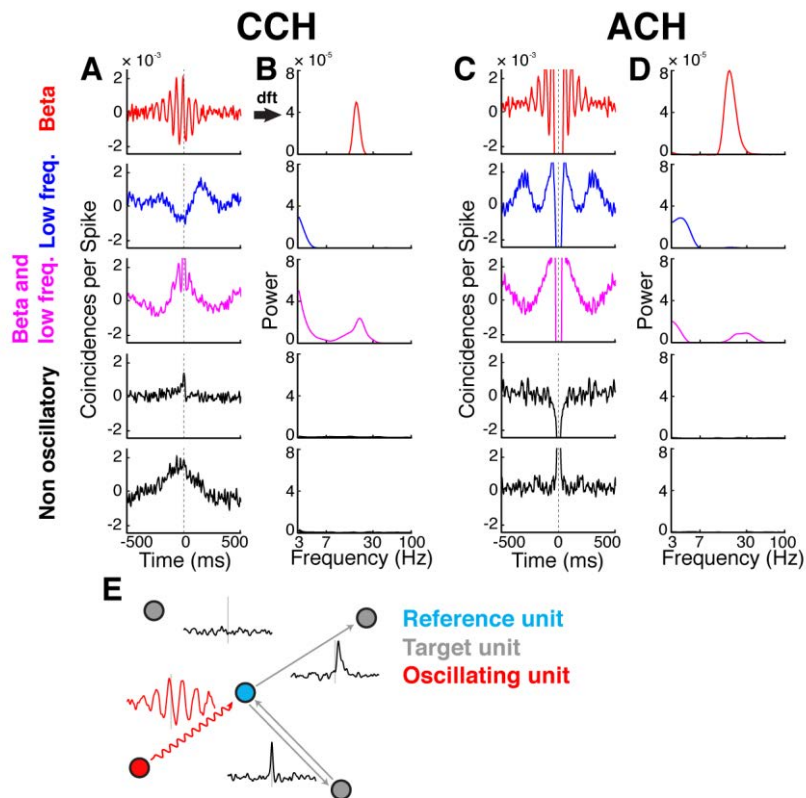


Figure 2. Cross- and auto-correlation histograms and frequency spectra. (A) Example crosscorrelation histograms (CCHs) for five example neuron pairs. Displayed amplitude is limited to $\pm 2.5 \times 10^{-3}$ coincidences per spike for better comparison. CCHs are color-coded based on their oscillatory synchronization frequency (red: beta band; blue: low frequencies; magenta: beta and low frequencies; black: no underlying frequency). (B) Corresponding frequency spectra of CCHs in a, frequency displayed on logarithmic scale (for better comparison limited to a power of 8×10^{-5}) and color-coded as in A. (C) Same as in A, but for auto-correlation histograms (ACHs). (d) Same as in B, but for the frequency spectra of the ACHs in C. (E) Illustration of different kinds of CCHs to a reference unit and the inferred connectivity.

Upper left: No peak is present in the CCH so the unit is not connected to the reference unit. Upper right: A peak at positive time lags indicates a connection from the reference to the target unit. Lower right: A peak is present straddling the 0 time lag with a maximum peak at 0, indicating a bidirectional connection. Lower left: Several peaks and troughs are present with a clear underlying frequency and a maximum peak at a negative time lag, indicating an oscillatory connection from the target to the reference unit.

Connections indicated by significant peaks or troughs in CCHs were identified by a cluster-based surrogate test (Maris et al., 2007) to all CCHs (see Materials and Methods), testing against surrogate CCHs. To control the family-wise error for the entire network, false discovery rate (FDR) correction was applied across all significant connections (Benjamini and Hochberg, 1995). For later topological analyses of oscillatory synchrony in the network, we applied Fourier transformations (Figure 2B-D; see Materials and Methods) to all CCHs and auto-correlation histograms (ACHs). The latter detected periodicity in the spiking of

2.1. Uniting functional network topology and oscillations

individual units, (Figure 2C), allowing classifying them as oscillators or non-oscillators.

Directional interaction between pairs of units was inferred from the time delay of significant peaks or troughs in the CCHs (Figure 2E). In early studies, a peak or trough in a CCH with a non-zero time lag was classified as a unidirectional connection from one neuron to another while a peak or trough with a zero time lag was classified as common drive to both neurons (Moore et al., 1970). However, recent studies based on complex models rather suggest that zero time lag peaks or troughs in CCHs mainly represent bidirectional connections, which can be explained by the dynamical relaying mechanism, and only rarely reflect a common drive (Vicente et al., 2008; Gollo et al., 2014). For this reason, we defined zero time lag peaks and troughs in the CCHs as bidirectional connections.

For additional validation of how well we could recover directed functional connectivity, we modeled two sets of “ground truth” networks with the same distribution of firing rates as recorded single units, one simple network (SN) and one complex network (CN) set (Equal rate model, see Materials and Methods). We could detect directed functional connections reasonably well (hits: 62% for SN, and 69% for CN) and hardly detected any false connections (correct rejections (CR) > 99% for SN and CN), independent of the underlying topology (Figure 2 - Figure Supplement 3B). To clarify if the missed connections were due to not detecting an existing interaction of a pair of neurons, or due to incorrect classification of directionality, we analyzed the detectability of connections independent of their direction (Figure 2 - Figure Supplement 3C), revealing similar results to the detect directed functional connections (hits: 58% for SN, and 69% for CN; CR: >99% for both). These findings suggest that the missed connections were due to not detecting an existing connection, in accordance with a high accuracy for extracting directionality of only detected connections (Figure 2 - Figure Supplement 3D; hits: 97% for SN, and 90% for CN; CR: 75% for SN, and 73% for CN).

Our simulated networks also allowed for a closer evaluation of zero time lag peaks as a result of either common drive or bidirectional connections. In direct comparison, the average common drive CCH as well as the average bidirectional CCH had a maximum at the zero time lag, but with the average bidirectional CCH having a 24 times higher peak (10.89 SD surrogate for bidirectional connections, and 0.45 SD surrogate for common drive; Figure 2 - Figure Supplement 3E), which is well in line with around 1% of all common drive pairs were detected as significant. When analyzing the distribution of maximum peaks in more detail, we found more than 7 times more bidirectional connections having a peak at the 0

2.1. Uniting functional network topology and oscillations

time lag than common drive pairs (Figure 2 - Figure Supplement 3F), in line with the results from the models described above (Vicente et al., 2008; Gollo et al., 2014). Taken together, all results from the modeled networks show an accurate detectability of directed functional interactions estimated from CCHs.

For a physiological classification of all significantly detected connections, we also analyzed their maximum peak or trough time lag distribution (Figure 2 - Figure Supplement 4A). Interestingly, the maximum peak or trough time lag distribution showed an exponential decay, with most of the peaks or troughs having a very short time lag (45.67% < 10ms, and 85.12% < 100ms), indicating predominantly direct influences of the units on each other. In case of oscillatory synchronized single units, as strongly present in the data, the classification of the maximum peak or trough time lags was more complex. Given that the maximum peak or trough time lag could be greater than half a cycle of the underlying frequency, it became unclear which unit is leading and which lagging, due to the presence of side lobes (e.g., see Figure 2A top panel). Since we found high numbers of oscillatory synchronized single units, predominantly in the beta (20Hz) and in the low frequency range (4Hz), as described in detail below, we analyzed the distribution of maximum peaks or troughs phase with respect to the underlying oscillatory frequency (Figure 2 - Figure Supplement 4B), and also found an exponential decay, similar to the maximum time lag peak or trough distribution. The majority of phase lags were within half a cycle around the zero time lag for both frequencies (beta connections: 77.70% < π , low frequency connections: 87.66% < π), suggesting that for most oscillatory synchronized connections we could accurately determine which unit was leading and which unit was lagging.

For analyzing the functional network topology, all units not connected to the largest inter-connected component were first discarded (mean number of units dropped: 17.75, SD: 9.56; mean percentage: 23.5%, SD: 13.3%; Supplementary Table 1) and binary directional connectivity matrices were created for every dataset (Figure 3A). We did not quantify the connection strength, since it has been shown to be biased by different firing rates (Cohen and Kohn, 2011).

Inter-area modular and small-world topology

First, we tested if the networks could be subdivided into modules, such that the number of connections was maximized within and minimized between modules. To properly evaluate modular topology, the fact that connectivity decays with distance has to be considered

2.1. Uniting functional network topology and oscillations

(Smith and Kohn, 2008; Gerhard et al., 2011). Figure 3B shows the distance-dependent decay of connectivity of our networks according to different subgroups: on the same electrode, on the same array, in the same area, between AIP and F5, between F5 and M1, and between AIP and M1. Connection density was not significantly different within all subgroups (Kruskal-Wallis test, $p > 0.05$).

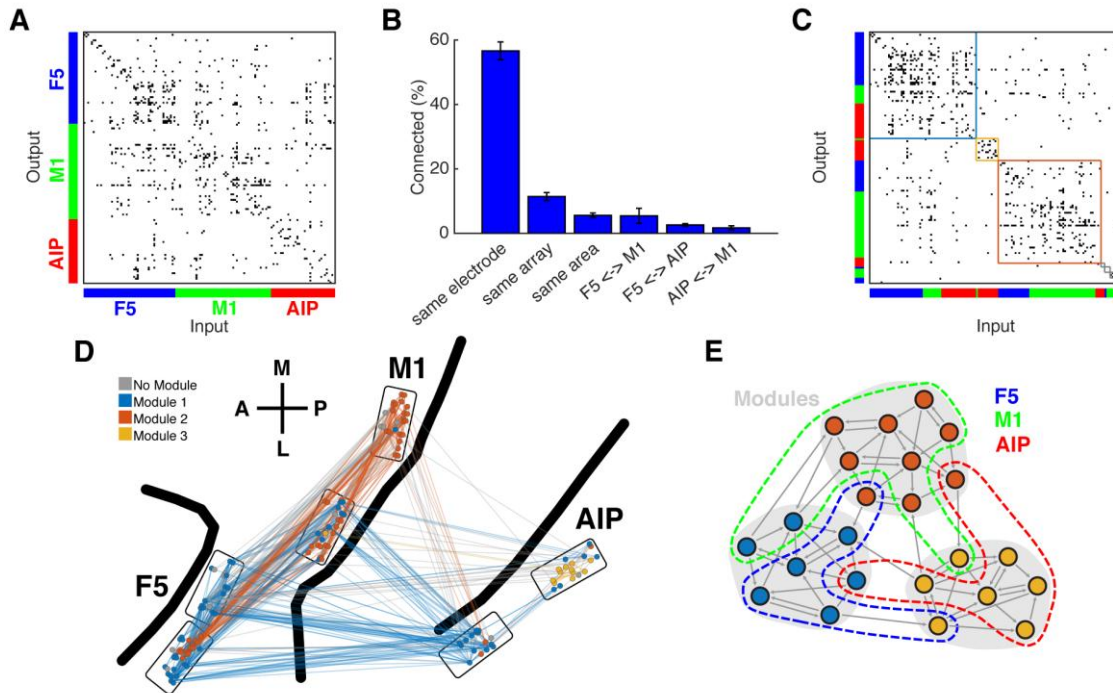


Figure 3. Connectivity characteristics and modular topology. (A) Connectivity matrix of one dataset from monkey M. Each dot represents a significant connection (Online Methods). Units are ordered by channel number of the recording system. (B) Distance dependent connectivity. From left to right: 56,7%, 11,5%, 5,6%, 5,5%,2,6%, and 1,7%. Note the clear distance dependent decay. (C) The same matrix as in A, but with nodes ordered according to an optimal modularity partition. Colored rectangles surround different network modules. (D) Anatomical network representation of the connectivity matrix in A. The brain is viewed as in Figure 1B. Single units and connections are color coded by module. (E) Schematic illustration of modular topology. Modules (dashed regions) consist mainly of single units of one cortical area, but also include small fractions of units from other areas.

Modular topology can be quantified by the modularity index Q . If a network can be completely subdivided into modules, Q will be 1. In contrast, if there is no modular structure present at all, Q will be close to 0. We found significant modular topology present in most of the networks (Mean Q : 0.405, SD: 0.087; permutation test, $p < 0.05$, sig. 10/12 datasets), taking the distance-dependent decay of connectivity into account. Modules were significantly predominated by units from a single area (mean largest proportion: 81.4%, SD: 14%; permutation test, $p < 0.001$), but 84% of all modules also included units from other areas, as became apparent when visualized as anatomical networks (Figure 3D, and Figure 3 - Figure Supplement 1A) or when displayed as a web where the locations of all units is

2.1. Uniting functional network topology and oscillations

determined by visualization of similarities (VOS) (Van Eck and Waltman, 2007) (Figure 3 - Figure Supplement 2A,B). These results reveal a functional modular topology partially not related to the anatomical boundaries between the different areas (Figure 3E).

Having shown that a modular topology is present, what is the detailed structure of how individual units are connected within the network? For this, we calculated the cluster coefficient C (with $C=1$ corresponding to every neighbor of every unit being interconnected, and $C=0$ indicating no interconnections between neighbors) and the average path length, L (defined as the average minimum number of units connecting one unit with another, across all pairs of nodes of the network; see Methods section). If units have dense local clustering (large cluster coefficient C) and can be reached from all other units via a short average path length, L , similar to random networks, the network is considered small-world (SW) (Watts and Strogatz, 1998; Bullmore and Sporns, 2009). Here, a value of $SW \gg 1$ indicates a small-world topology, whereas $SW=1$ corresponds to no small-world effect.

We found significantly higher average cluster coefficients C in comparison to surrogate networks (mean: 0.266, SD: 0.068; permutation test, $p < 0.001$, sig. 12/12 datasets) and on average similar path lengths L (mean: 3.451, SD: 0.823; mean difference to surrogate networks: -0.007; permutation test, $p < 0.05$, sig. higher 5/12, sig. smaller 5/12 datasets). Consequently, all networks had a significant SW-coefficient (mean: 3.05, SD: 0.66; permutation test, $p < 0.001$, sig. 12/12 datasets), suggesting that despite a modular structure the neuronal network is efficiently processing and transmitting information (Watts and Strogatz, 1998).

Degree centrality, betweenness centrality, and hubs

Some networks, have been shown to exhibit heavy-tailed centrality distributions, with a small number of nodes strongly embedded in the network (hubs), which make a strong contribution to the network function (van den Heuvel and Sporns, 2013a). A simple and robust measure of centrality is degree centrality (k), which is the number of connections per unit. On average 6.27% (SD: 2.29%) of all possible connections were realized. The degree distribution (Figure 4A) was heavy-tailed and best described by an exponential truncated power law model ($P(k) \sim k^{\gamma-1}e^{-k/k_c}$, $\gamma = 0.6839$; cutoff degree of $k_c = 8.657$; EXPTPL: adjusted $R^2 = 0.9891$, including a penalty for number of fitted variables), compared to a power law ($P(k) \sim k^{-\gamma}$; PL: adjusted $R^2 = 0.9177$), exponential (EXP: adjusted $R^2 = 0.9742$), or Gaussian (GAUS:

2.1. Uniting functional network topology and oscillations

adjusted $R^2 = 0.6826$) model. In contrast, surrogate networks with the same distance-dependent connectivity were not heavy-tailed and were best described by a GAUSS model (GAUS: adjusted $R^2 = 0.9655$; PL: adjusted $R^2 = 0.3061$; EXPTPL: adjusted $R^2 = 0.5006$; EXP: adjusted $R^2 = 0.6419$). In agreement with the EXPTPL model, networks had significantly more single units within the low, less within the intermediate, and especially more in the high degree range, than surrogate networks (cluster-based permutation test, $p < 0.05$), clear evidence of hubs, independent of distance-dependent connectivity.

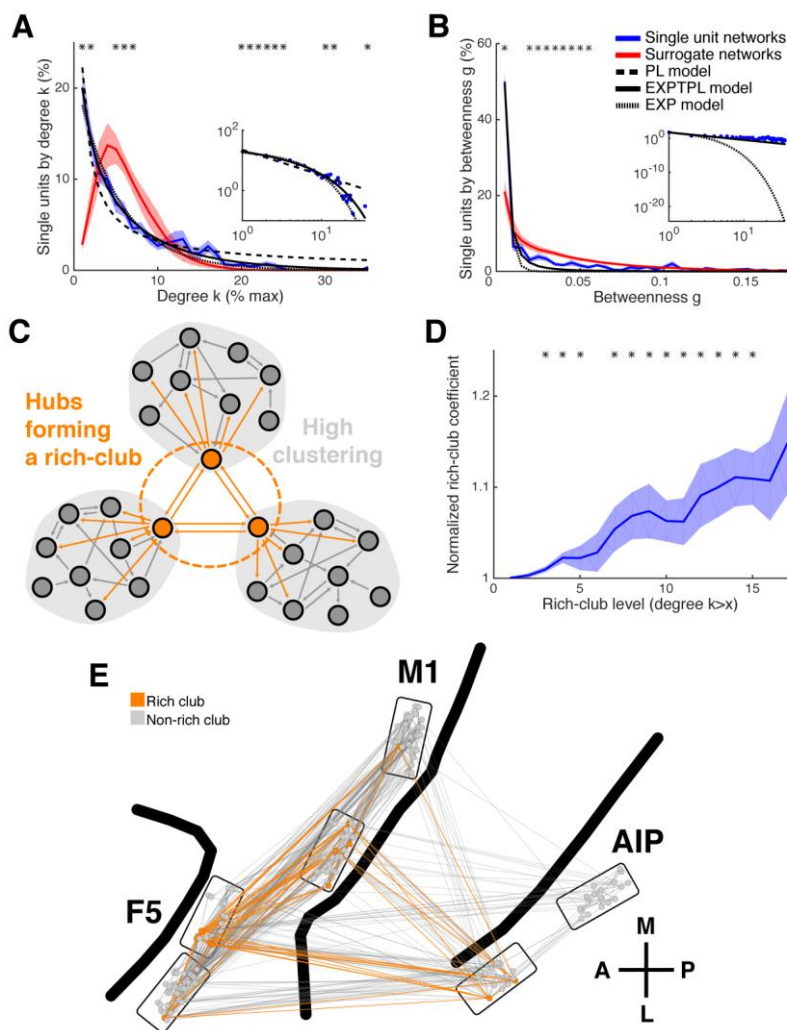


Figure 4. Centrality measures, hubs, and rich-club topology. (A) Average degree centrality distribution of all networks (blue) and corresponding surrogate networks (red). Black lines reflect different models fitted to the data (see legend in B). The degree distribution of each dataset was normalized to the possible maximum number of connections per network. The area under the curve was normalized to 100% before averaging. Line shadings show standard error across datasets. Asterisks represent significant differences to surrogate networks. Inlay shows the same distribution and models on a log-log scale. (B) Same as in A, but for the betweenness centrality distribution. Note that the slopes for the EXPTPL and PL model are identical, since the exponential coefficient of the EXPTPL model was zero. (C) Schematic view of a rich-club

topology connecting highly clustered modules. (D) Average rich-club level of all datasets relative to surrogate datasets. Asterisks represent significant differences of rich-club level to surrogate networks. (E) Anatomical network representation, as in Figure 3D, with connections and units color-coded based on rich-club membership (orange).

A more global aspect of centrality is captured by betweenness centrality (g), an index of the number of shortest paths from all single units to all others that pass through that single unit, normalized by the number of all shortest paths (van den Heuvel and Sporns, 2013a). Similar to degree centrality, the betweenness centrality distribution (Figure 4B) was

2.1. Uniting functional network topology and oscillations

heavy-tailed and best described by a PL model, with an estimated exponent of $\gamma = 2.212$ (PL: adjusted $R^2 = 0.9753$; EXPTPL: adjusted $R^2 = 0.9745$; EXP: adjusted $R^2 = 0.9593$; GAUS: adjusted $R^2 = -0.1509$). The betweenness centrality distribution of surrogate networks was also heavy-tailed and was best described by an EXPTPL model (EXPTPL: adjusted $R^2 = 0.99$; PL: adjusted $R^2 = 0.9771$; EXP: adjusted $R^2 = 0.9061$; GAUS: $R^2 = -0.5511$). Still, in contrast to the PL model, the EXPTPL model had smaller values in the high and low betweenness centrality range. Statistically networks showed a significantly higher number of single units in the low and fewer units in the intermediate betweenness range than surrogate networks (cluster-based permutation test, $p < 0.05$). These findings confirm the presence of hub neurons for betweenness centrality. Units with high degree centrality also tended to have high betweenness centrality ($r = 0.75$, $p < 0.001$, Spearman correlation), suggesting a coherent group of hub units. We found no significant differences in number of hubs per area (normalized $k \geq 9$, $g \geq 0.03$; Tukey's honest significant difference test on average group ranks, $p < 0.05$), indicating a distributed hub topology with no area acting as a network center. Together, we have shown that centrality of single units is strongly heterogeneous in the network, with a large group of units being marginally involved in the network and a small group of spatial distributed hub units being extremely central. The presence of hubs provides further evidence of a complex network topology at the single unit level.

However, it has been shown that detectability of functional connections decreases with lower firing rates (Cohen and Kohn, 2011). Since the detected firing rates varied approximately across two orders of magnitude (Figure 1 - Figure Supplement 1B), this could lead to an underestimation of degree for low spiking units and an overrepresentation of high firing units as hubs. Therefore, we performed a careful examination of the influence of firing rates on degree and betweenness centrality based on our equal rate model (see Materials and Methods). Two sets of networks were tested, simple networks (SNs) and complex networks (CNs), as mentioned previously. SNs had normally distributed connectivity based on the best fitting Gaussian model for the surrogate network degree centrality distribution, while connectivity for CNs were set to precisely resemble the EXPTPL model for the average degree centrality distribution of the measured networks. CNs additionally had a small-world and rich-club topology, as described in the following section.

Differences in firing rate and any possible biases due to the applied method to estimate directed functional connectivity had no effect on the shape of the degree centrality

2.1. Uniting functional network topology and oscillations

distribution for both kind of networks (Figure 4 - Figure Supplement 1A). The betweenness centrality distribution for CNs was also unchanged and only slightly impaired for the SN (Figure 4 - Figure Supplement 1B). Nevertheless, the best fitting model for the betweenness centrality distribution of SNs was in neither case (modeled or detected) a PL, as it were for the measured data and the CNs, suggesting no distorting effect by differences in firing rate and the applied method to estimate directed functional connectivity. Importantly, also the average C, average L, and SW-coefficient were correctly detected for both kind of networks.

It is also possible that subsampling, a natural limitation in electrophysiological recordings, could artificially cause a heavy tailed degree centrality distribution even if the underlying connectivity is random (Han et al., 2005; Gerhard et al., 2011). We simulated a neuronal layer of 32,000 neurons with the same distance-dependent connectivity density as detected in our data (Figure 3B), but with Poisson distributed connectivity (Figure 4 - Figure Supplement 2A; see Materials and Methods). Subsampling was performed in correspondence with our array configuration down to the number of neurons we recorded for real datasets, showing no change to the shape of the degree distribution (Figure 4 - Figure Supplement 2B). Only when we decreased the connection density of the model below the detected connectivity in our data was a false heavy-tailed degree distribution apparent (Figure 4 - Figure Supplement 2C), which was highly correlated with the networks breaking apart into unconnected components ($R^2 = 0.93$). Additionally, this effect could not be present in our analyzed data since we only analyzed the largest component of the single unit networks. These controls suggest that the existence of hubs can neither be explained by distance-dependent connectivity, differences in firing rates, or subsampling.

Rich-club topology

In some networks hubs exhibit a strong tendency to link to each other, forming a rich-club (Colizza et al., 2006), which can be measured by a rich-club coefficient that expresses the tendency of highly connected hub nodes to show above-random levels of interconnectivity (Figure 4C). Hub units showed a significantly higher level of interconnectivity than surrogate networks, with up to 15% more connections (Figure 4D; cluster-based permutation test, $p < 0.05$).

For our equal rate model, we tested if differences in firing rate and the applied method to estimate directed functional connectivity could cause a false rich-club effect. The

2.1. Uniting functional network topology and oscillations

present rich-club topology of CNs could be correctly detected, as well as no false rich-club topology was detected for SNs (Figure 4 - Figure Supplement 1C). Although the slope of the rich-club coefficient was changed for CNs, rich-club topology was only significant if present (cluster-based permutation test, $p < 0.05$), suggesting a correct representation of rich-club topology for the measured networks.

The rich-club contained neurons from all areas with a rich-club level set to $k \geq 9\%$ (Figure 4E, Figure 3 - Figure Supplement 1B, and Figure 3 - Figure Supplement 2C; mean rich-club neurons: 27%, SD: 18%; similar results with k set to other levels). A rich-club that spans multiple areas, as described here, has been proposed as a robust structure facilitating efficient communication (van den Heuvel and Sporns, 2013a).

Network topology of oscillatory synchrony

Oscillatory synchronization has been proposed as a mechanism for efficient communication (Fries, 2009). As demonstrated above, oscillatory and non-oscillatory synchronized spike patterns for communication could be identified (Figure 2, Figure 2 - Figure Supplement 1B,2). We therefore investigated if specific relationships between distinct frequencies and network topology emerged. Frequency spectra of ACHs of all units and of CCHs that had a significant connection were tested for significant frequency bins above chance (cluster-based surrogate test, $p < 0.05$). We found beta (18-35 Hz) and low frequency (3-7 Hz) oscillations predominantly present in the spiking patterns of all datasets (Figure 5A, and Figure 5 - Figure Supplement 1C-E). Oscillatory synchrony in both frequency ranges was present more often in CCHs (mean beta: 38.3%, low: 44.3%) than in ACHs (mean beta: 22.5%, low: 31.7%), suggesting that the group of oscillating single units (oscillators; Supplementary Table 2) communicates in their underlying frequency to a larger group of units.

Interestingly, there was also a significant group of oscillating single units present in the gamma range (45-80Hz), which was not mirrored in the CCHs. One possible explanation could be that these units communicate via long-range gamma synchronization with topographically distant areas we did not record, such as the visual cortex (Gregoriou et al., 2009).

2.1. Uniting functional network topology and oscillations

Oscillators and oscillatory connections were widely distributed and seemed to be very central across all areas (Figure 5B, Figure 3 - Figure Supplement 1C, and Figure 3 - Figure Supplement 2D), giving rise to the idea that oscillators could be the hubs of the networks.

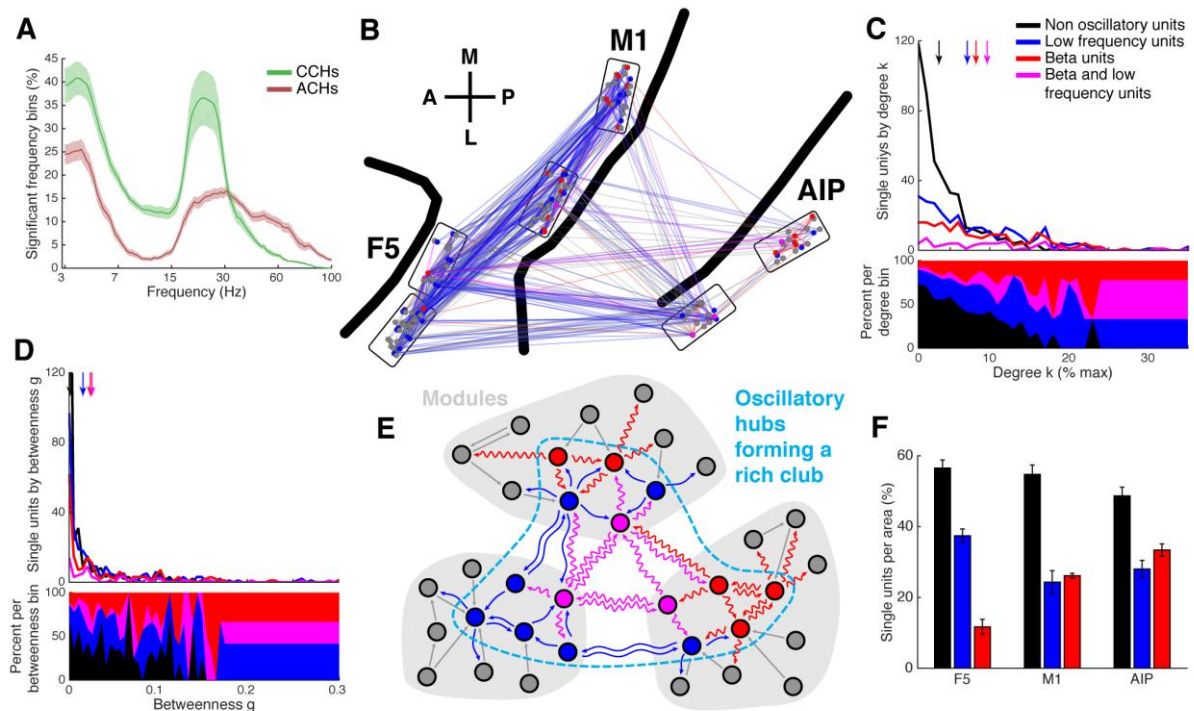


Figure 5. Low frequency and beta oscillators and their network topology. (A) Average number of significant frequency bins of all ACHs and CCHs over all datasets. Frequencies displayed on a logarithmic scale. Line shading bars represent standard error across datasets. (B) Anatomical network representation as in Figure 3D with connections and units color-coded by underlying oscillations (see legend in C). (C) Degree centrality distribution of all datasets separately for beta and low frequency oscillators, non-oscillators, and single units oscillating in both frequency ranges. Upper panel, summed degree centrality distribution of all single units. Median degree is represented by arrows in corresponding color: beta units: 7.5, low frequency units: 6.3, beta and low frequency units: 8.9, and for non-oscillators: 2.7. (D) Same as in C but for the betweenness centrality distribution. Median for beta units: 0.023, low frequency units: 0.016, beta and low frequency units: 0.026, and for non-oscillators: 0.001. (E) Schematic view of the found network topology of oscillators. Oscillators form a rich-club spanning all areas. (F) Distribution of oscillators across areas. The number of single units is normalized to 100% per area. F5 has significantly less beta (red) and significantly more low frequency oscillators (blue) than M1 and AIP. Note that units oscillating in both frequency ranges are counted in both. Non-oscillators (black) still remain the largest group in all areas.

Figure 5C shows the average degree centrality distribution for all networks, as in Figure 4A, but separately for beta and low frequency oscillators, non-oscillators, and units oscillating in both frequencies. There was a clear dominance (high percentage) of oscillators in the high degree range, whereas non-oscillators dominated in the low degree range. The degrees of all three oscillator groups were significantly higher than for non-oscillators (Tukey-Kramer test for rank, $p < 0.001$). Betweenness centrality was also significantly higher for oscillators than for non-oscillators, similar to degree centrality (Figure 5D; Tukey-Kramer test for rank, $p < 0.001$). The number of units oscillating in both frequencies was not higher

2.1. Uniting functional network topology and oscillations

than expected by coincidental overlap of the two frequency bands (permutation test, $p > 0.05$).

Nevertheless, it could be possible that CCHs are more sensitive to oscillatory synchrony than to non-oscillatory synchrony, which would induce a bias when comparing these two groups. At this point, it is important to emphasize that we first tested for significant connectivity independent of oscillatory behavior and only in a second step these connections were tested for their oscillatory behavior as described in the methods section. This ensured that any detected connection is based on a significant amount (or suppression) of coincidental spikes without any selective sensitivity for oscillatory coupling. As an additional test, we simulated pairs of neurons either with an oscillatory or non-oscillatory firing pattern (see Materials and Methods). Since peaks and troughs in CCHs reflect a systematic time lag in spiking between units across trials we simulated different degrees of coupling strengths by systematically varying the trial-wise time offset in spiking for both firing pattern types. Synchronization strength was simply a function of the variation in spike timing offsets between the two neurons and not whether the firing pattern was oscillatory or not (Figure 5 - Figure Supplement 2), confirming that oscillatory coupling is not a priori more detectable than non-oscillatory coupling.

Besides these methodological issues already addressed, it is possible that higher firing rates introduce a bias in the statistical detection of significant frequency bins. To control for this possibility, we applied thresholds for the detection of beta and low frequency oscillations. Thresholds were chosen to give, as closely as possible, the same number of beta and low frequency oscillators as statistical methods. Using this method all three groups had a higher degree and betweenness centrality than non-oscillators, similar to statistical detection (Tukey-Kramer test for rank, $p < 0.001$). To rule out that firing rate dependent detectability of functional connections could cause a spurious inter-dependence of high centrality and detection of oscillatory synchrony, we repeated testing for differences in centrality only with units having a firing rate of 10Hz and above, confirming that oscillators had significantly higher centrality values (Tukey-Kramer test for rank, $p < 0.001$). Similar results were obtained when we tested the data of each monkey individually (Tukey-Kramer test for rank, $p < 0.01$). To our knowledge, the current results represent the first evidence that oscillators have a higher centrality in the single unit network than non-oscillators. Consequently, the rich-club of all networks overlapped significantly with oscillating single

2.1. Uniting functional network topology and oscillations

units (permutation test, $p < 0.05$), highlighting oscillators as the backbone (van den Heuvel et al., 2012) of single unit functional connectivity (Figure 5E).

The number of oscillators did not differ between areas (Tukey-Kramer test for rank, $p < 0.05$), in agreement with the distribution of hubs as well as rich-club units across areas. Closer examination of oscillator types revealed significantly more beta oscillators in AIP and M1 than in F5, and more low frequency oscillators in F5 than in M1 and AIP (Figure 5F; Tukey-Kramer test for rank, $p < 0.05$), reinforcing the notion that different cortical areas operate more strongly in some frequency ranges than others (Brovelli et al., 2004).

A further unresolved question is whether a direct relationship exists between oscillatory synchronization and functional rich-club topology. It is well known that oscillatory synchrony in frontal and motor areas appears in short bursts of only a couple of cycles with variable length and amplitude (Murthy and Fetz, 1996; Lundqvist et al., 2016). We used this property of oscillatory synchrony to split up our data into two equal blocks with high oscillatory and low oscillatory synchrony to investigate the effect on rich-club topology. Since a minimum number of trials are required to properly estimate the functional connectivity for topological analyses, we used the two datasets from monkey M where we recorded more than 900 trials (Supplementary Table 1). The data was split into two blocks with equal number of trials per condition to prevent any biases by different epochs or conditions. Instead of calculating unit-wise ACHs we pooled the activity of all units and estimated single trial population ACHs spectra, reflecting the trial-wise level of oscillatory synchronization. Single trial population ACHs calculations and frequency analyses were performed the same way as for single unit ACHs (see Materials and Methods) and divided by their average power in the beta (18-35 Hz) and low frequency (3-7 Hz) band (Figure 5 - Figure Supplement 3A). After separation into two blocks, the estimation of functional connectivity and network topological analyses were repeated as if they were two separate datasets. For a valid statement about changes in rich-club topology, the network structure and in particular the degree distribution, should not be changed. For both datasets the unit-wise degree as well as the degree distribution were very similar (Figure 5 - Figure Supplement 3B,C), as well as the betweenness centrality distribution (data not shown). However, when comparing the rich-club level there was a striking difference for higher rich-club levels (Figure 5 - Figure Supplement 3D). In both datasets, the high oscillatory state network showed a clear rich-club topology, whereas the low oscillatory state network hardly

2.1. Uniting functional network topology and oscillations

showed any rich-club effect. These results suggest that a rich-club topology is only present when there is a high level of oscillatory synchrony in the network.

Functional network topology and firing rate prediction

Utilizing the identified network topology, the firing rate of individual units can be predicted by the firing rate of input units, providing an estimate on how much of the single unit activity can be explained by functional network connectivity. Each CCH can be understood as a transfer function of spike rates between two units, describing the coincidences per spike at every time point relative to each other. Negative time bins reflect input from the reference unit to the target unit while positive time bins reflect the output. To predict the firing rate of a unit, we convolved the spike trains of all units having a significant connection to the corresponding unit with their respective CCHs (output part). Assuming single units to be simple linear integrators, we summed up the individual convolved spike trains (Figure 6A,B) and correlated these estimated signals with the original spike trains of the target units smoothed with a Gaussian kernel (SD: 3.66 ms), identical to the CCH smoothing. Ninety-nine percent of predicted firing rate curves were positively correlated with the real firing rates of the corresponding target units (Figure 6C).




Figure 6. Prediction of firing rates based on network topology. (A) Average firing rate of one example single unit recorded in F5 in monkey S for the four conditions used in this study during the fixation (Fix), cue (Cue), memory (Mem), and movement period (Mov). The complex tuning patterns for the different task conditions (grip types; free-choice vs. instructed trials) is clearly visible. (B) Predicted firing rate of the same unit as in A based on the population activity of the connected neurons. Curves in (A-B) were smoothed with an additional Gaussian kernel (SD: 40ms). (C) Histogram of correlation coefficients between the true and predicted spike trains of all single units of all datasets. Significant correlations are marked in red. Note that hardly any correlation coefficient were negative. (D) Histogram of correlation coefficients of condition averaged firing rates. Coloring as in C.

However, these correlations could also be due to synchronous up and down states of the brain (Gilbert and Sigman, 2007), which makes proper statistical testing obligatory. Three different permutation tests were applied: shuffling of trials, shuffling of the output

2.1. Uniting functional network topology and oscillations

parts of CCHs, and shuffling of input units. Only if the correlation coefficient significantly exceeded all three permutation distributions ($p < 0.05$) was the correlation considered significant. Remarkably, 45% of the firing rate patterns of our single units could be significantly predicted by their inputs. The differences between grasp types and decision conditions could be significantly predicted in 9% of all cases (Figure 6D; positive correlation: 79%; shuffling of the transfer kernels and input units, $p < 0.05$), even using this simple approach that involved no parameter fitting. The functional network topology presented here allows a surprisingly accurate prediction of temporal firing dynamics, suggesting that the network captured in our recordings, despite being a small subset of the entire network, accurately represents a large portion of the relevant communication in the fronto-parietal grasping network.

Discussion

We analyzed single unit functional network topology across several cortical regions of three monkeys performing a delayed grasping task. The network was structured as a complex network (Bullmore and Sporns, 2009) with a modular SW topology, and highly central hub-units localized in all three areas forming a rich-club. The advantage of such a topology is that it allows for fast and dynamical information processing combined with high robustness against errors (Barabási and Oltvai, 2004; Bassett and Bullmore, 2006; Bullmore and Sporns, 2009; van den Heuvel et al., 2012). More detailed analyses of the kind of synchronization processes within the network revealed that the population of single units could be divided into two groups: oscillatory spiking and synchronized units in the low frequency range or in the beta range, and a group of non-oscillatory spiking units. Importantly, the hubs and therefore the rich-club consisted predominantly of oscillators, while the peripheral neurons were predominantly non-oscillators.

Why is oscillatory synchrony such a central element of functional network topology? More and more evidence supports the hypothesis that information is propagated not only as a simple rate code, but by feed-forward coincidence detection accomplished by oscillatory synchrony (Fries, 2009), meaning that phase-synchronization of neurons with one another is used as a selection mechanism for information transmission. The advantage of this mechanism is not only a reduction of energy cost, but also rhythmic gain modulation. By changing the phase of a synchronous neural population, such as in high-order areas, the input of one group of neurons can be selectively amplified as inputs to another group of

2.1. Uniting functional network topology and oscillations

neurons, allowing for high selectivity and high flexibility, which are exactly the requirements a hub has to fulfill (van den Heuvel and Sporns, 2013a). While feed-forward coincidence detection can theoretically also be accomplished by non-oscillatory processes (Fries, 2009), the coordination of a network spanning different areas requires a larger group of neurons to fire in a coherent manner (Buzsáki and Wang, 2012). A rich-club of oscillating neurons is exactly that, a coherent structure cross-linking functionally segregated modules (van den Heuvel and Sporns, 2013b), suggesting oscillators act as a backbone promoting and coordinating functional communication across different cortical areas (van den Heuvel et al., 2012). This hypothesis is also in accordance with the finding that synchronization over larger distances (>2mm) is almost always oscillatory, whereas synchronization over short distances occurs also in the absence of oscillations (König et al., 1995).

What are the roles of the two different distinct frequency bands present in this network? Parietal and motor areas have been found to communicate via ~20Hz beta synchronization (Pesaran et al., 2002; Brovelli et al., 2004; Pesaran et al., 2008b; Dean et al., 2012) and an increment in beta band activity seems related to the maintenance of the current sensorimotor or cognitive state, in agreement with findings in the basal ganglia (Engel and Fries, 2010). Oscillatory synchrony in the low frequency range (1-4Hz) has been shown to be important for communication within and between the prefrontal and motor areas (Siegel et al., 2009; Nacher et al., 2013) and as a potential population mechanism of movement generation in motor and premotor cortex during reach initiation (Churchland et al., 2012). Therefore, beta seems to be a stabilizing signal, low frequencies a global coordination signal, and both are involved in movement initiation with opposing roles. One possibility is that a function of the rich-club, composed of beta and low frequency oscillators spanning parietal and prefrontal cortex, is coordinating movement generation and initiation. Another possible explanation is that the power of fast oscillations is modulated by the phase of slow oscillations, termed cross-frequency phase-amplitude coupling, which could serve as a neuronal syntax for information transmission (Buzsáki, 2010; Buzsáki and Mizuseki, 2014). Our observation of oscillators in both frequency ranges simultaneously (third row of Figure 2C,D, and Figure 5C,D) support this concept.

Interestingly, we found that beta oscillators were present most frequently in AIP, followed by M1, hardly in F5, and in reverse order for low frequency oscillators (Figure 5F). This is in line with the previous findings that information via beta band is primarily

2.1. Uniting functional network topology and oscillations

transmitted from the parietal to the frontal regions and not vice versa (Brovelli et al., 2004). In areas that are hierarchically lower than the parietal lobe, such as the visual system, beta was identified as a top-down communication frequency (Bastos et al., 2015). Therefore, the parietal lobe might be a center of beta generation. Low frequency oscillatory synchrony during active behavior has been found predominantly in prefrontal areas (Siegel et al., 2009; Nacher et al., 2013). We speculate that the center of low frequency oscillation could be in the prefrontal cortex, suggesting that different anatomical regions generate and communicate with different frequencies. The exact reason for the presence of distinct frequency bands for communication and their detailed interplay needs to be addressed in future studies.

The single unit network topology was highly similar to the regional network of the brain measured by EEG, MEG, DTI or fMRI (Bullmore and Sporns, 2009; Rubinov and Sporns, 2010; van den Heuvel et al., 2012; van den Heuvel and Sporns, 2013a), which strongly suggests that the observed topological properties are scale-invariant (Bullmore and Sporns, 2009). Oscillatory synchrony may therefore act as a global coordination mechanism across the whole cortex.

The modules of the network were primarily composed of the individual areas themselves. Yet, most modules also consisted of a small, but significant, proportion of units from other areas, indicating that the anatomical distance does not necessarily reflect the functional distance. This finding is in line with a recent study showing that the population of neurons within one area can be split up into “choristers,” which are strongly coupled to the rate of the whole population, and “soloists,” which are not (Okun et al., 2015). We speculate that “soloists” could be part of functional circuits centered in other brain areas, in accordance with the present modular topology.

Since we recorded only from a subpopulation of the actual network, it was important to evaluate whether the observed network topology sufficiently represented the fronto-parietal grasping network. We demonstrated that a significant amount of the firing rate of single units could be predicted using only their network inputs, even for complex tuning patterns, suggesting that even a small fraction of the network is enough to characterize a reasonable amount of the spatio-temporal spiking dynamics. Furthermore, we demonstrated on a model that subsampling from a huge network with the same distance-dependent connectivity density as detected in our data did not affect the shape of the

2.1. Uniting functional network topology and oscillations

degree distribution (Figure 4 - Figure Supplement 2). For these reasons, we are confident that our analyzed single unit network constitutes a significant representation of the underlying network dynamics.

One possible point of misinterpretation of the functional network structure could be common drive, resulting in an overestimation of connectivity. Our method to detect functional connectivity corrects for common drive due to stimulus- and movement-locked inputs as well as for trial-wise fluctuations in spiking. Nevertheless, there are two possible additional sources of common drive. The first is the possibility that two neurons receive input from a third neuron while themselves being functionally uncoupled, resulting in a significant peak in the CCHs due to their input similarity. We investigated this possibility using our equal rate model, which included physiologically plausible firing rates and pairwise correlations. Common drive pairs of simulated simple or complex networks were detected as being significant in only around 1% of all cases, suggesting that, irrespective of the underlying topology, our method for detecting functional connectivity is hardly biased by pairwise common drive. The second possibility is that cortical columns or areas could receive common drive input that would cause these neurons to fire in a synchronized fashion even if they were functionally uncoupled. In such a scenario two things would be expected: first, units on the same electrode, as well as units in the same area, should show a similar connectivity pattern. Second, all neurons in the network should show a similar number of functional connections, since they are synchronized by common drive, resulting in a uniform degree centrality distribution. However, we found 43% of all neurons on the same electrode to be not connected, and only sparse connectivity was found in the same area with strongly connected pairs of neurons next to unconnected pairs (e.g., Figure 2 – Figure Supplement 1B,2). Most importantly, the degree distribution of the measured networks was highly heterogeneous and heavy-tailed in contradiction to what would be expected by a strong influence of column- or area-specific common drive. Therefore, it is unlikely that event unrelated common drive can account for a significant amount of the detected functional connections. Further evidence arises from the fact that we found beta, low frequency, and non-oscillatory synchronization with different maximum peak or trough time and phase lags (Figure 2 – Figure Supplement 4), present simultaneously across all areas, also not consistent with a global common drive bias.

2.1. Uniting functional network topology and oscillations

To our knowledge, these results provide the first evidence of oscillatory synchrony as a central coordinating mechanism for the formation of functional network topology at the single neuron level. The combination of communication properties of oscillating single units and their functional topology adds an essential dimension to the understanding of neural circuits. By demonstrating that oscillating neurons form a backbone for functional connectivity, spanning several areas, we provide a unified basis for understanding the neuronal computations coordinating and generating behavior at the network level.

2.1. Uniting functional network topology and oscillations

Materials and Methods

Basic procedures

Neural activity was recorded simultaneously from many channels in two female and one male rhesus macaque monkey (Animals S, Z, and M; body weight 9, 7, and 10 kg, respectively). Detailed experimental procedures have been described previously (Michaels et al., 2015). All procedures and animal care were 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* (Research et al., 2003).

Behavioral Task

Figure 1A illustrates the time course of the behavioral task as described previously (Michaels et al., 2015). Trials started after the monkey placed both hands on the resting positions and fixated a red fixation disk (fixation period). After 600 to 1000ms, cues in the form of disks were shown next to the fixation disk for 300ms to instruct the monkey about the required grip type (power or precision; cue period). During this epoch the grasp target, a handle, was also illuminated. In the instructed task one disk was shown, while in the free-choice task both disks were turned on, indicating that the monkey was free to choose between the two grip types. The monkey then had to memorize the instruction for 1100 to 1500ms (memory period). The switching off of the fixation light cued the monkey to reach and grasp the target (movement period) in order to receive a liquid reward. Importantly, during free choice trials the reward was iteratively reduced every time the monkey repeatedly chose the same grip type. All trials were randomly interleaved and executed in darkness. The behavioral task also contained delayed instructed trials, which were not analyzed in this study.

Chronic electrode implantation

Surgical procedures have been described previously (Michaels et al., 2015). In short, each animal was implanted with two floating microelectrode arrays per area (FMAs; Microprobes for Life Sciences; 32 electrodes; spacing between electrodes: 400 μ m; length: 1.5 to 7.1 mm monotonically increasing to target grey matter along the sulcus). Animal S and Z were implanted with four FMAs in area AIP and F5 in the left and the right hemisphere, respectively. Animal M was implanted with a total of six FMAs in the same cortical areas and two additional arrays in area M1, in the left hemisphere (Figure 1B).

2.1. Uniting functional network topology and oscillations

Neural recordings and spike sorting

Neural 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 band-pass hardware filter; for monkey S and Z) or a 256 channel Tucker-Davis system (TDT RZ2; sampling rate 24.414 kS/s; 0.6-10000Hz band-pass hardware filter; monkey M).

For spike detection, data were first low-pass filtered with 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 common noise-sources principal component (PC) artifact cancellation was applied for all electrodes of each array as described previously (Musial et al., 2002). To ensure that no individual channels were eliminated, PCs with any coefficient greater than 0.36 (conservatively chosen and with respect to normalized data) were retained. Spike waveforms were detected 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 PCs, 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 different average waveforms (templates) was done to detect overlaid waveforms (Gozani and Miller, 1994). To achieve this, 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 discriminate analysis for classification of waveforms. After spike identification, the target template was subtracted from the filtered signal of the corresponding channel to reduce artifacts during the detection of the next template. This procedure allowed us to detect spikes with a temporal overlap up to 0.2 ms. Unit isolation was evaluated again, based on the five criteria mentioned above, to determine the final

2.1. Uniting functional network topology and oscillations

classification of all units into single or non-single units. Stationarity of firing rate was checked for all units and in case it was not stable over the entire recording session (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 single units). Only single units fulfilling all of these criteria, and no multi-units, were further used in this study.

Functional connectivity analysis

After sorting, spike events were binned in non-overlapping 1-ms windows to produce a continuous firing rate signal (1 kHz) and aligned to cue and movement onset. Two time windows were chosen for further analysis (Cue onset: -700 to 1500ms; Movement onset: -300 to 500ms), since neuronal activity was locked to both events, with a variable memory period between them. Note that all three monkeys had very consistent movement times (mean SD across datasets = 39ms).

The functional network topology of single-unit populations was derived from analyses of pairwise correlations (Yu et al., 2008). We calculated cross-correlation histograms (CCHs; time lags: -500 ms to 500 ms) between all pairs of single units of each dataset (Bair et al., 2001):

$$CCH_{n_1, n_2}(\tau) = \frac{1}{M} \sum_{i=1}^M \sum_{t=1}^N \frac{x_{n_1}^i(t) x_{n_2}^i(t + \tau)}{(N - |\tau|) \sqrt{\lambda_1 \lambda_2}} \quad (1)$$

where M is the number of trials, t is time, N is the number of time bins in the trial, $x_{n_1}^i$ and $x_{n_2}^i$ are the spike trains of single units n_1 and n_2 on trial i , τ is the time lag, and λ_1 and λ_2 are the mean firing rates of the two single units across the entire time interval M . The denominator is normalizing for the degree of overlap ($N - |\tau|$) in the CCH and the geometric mean spike rate $\sqrt{\lambda_1 \lambda_2}$, which is the most common normalization used for CCHs (Bair et al., 2001; Smith and Kohn, 2008). The normalized CCHs were then averaged across all time periods and task conditions (e.g., see Figure 2 – Figure Supplement 1A).

Subsequently, all CCHs were corrected for correlations induced by common stimulus drive or global state changes, such as arm and hand movements, as well as for trial-wise fluctuation in spiking, by simulating and subsequently subtracting surrogate CCHs. Surrogate CCHs contain the same stimulus locked correlation, but no pairwise temporal correlation. To this end, peri-stimulus time histograms (PSTH) were calculated for the same two time

2.1. Uniting functional network topology and oscillations

windows and alignments (Cue and Movement onset) as mentioned above, separately for each single unit and task condition (smoothed with a Gaussian kernel, SD: 3.66 ms). Artificial spike trains were generated from an inhomogeneous Poisson process using the PSTHs as the rate function (Ramalingam et al., 2013). These artificial spike trains preserved the number of trials and the number of spikes per trial, but varied in the timing of individual spikes (surrogate data; e.g., Figure 2 – Figure Supplement 1A). Since the number of spikes per trial was preserved for all units recorded simultaneously, any trial-wise common drive is equally present and therefore accounted for in the surrogate data (Smith and Kohn, 2008). From these surrogate data, surrogate CCHs were calculated by replacing x_n^i with the trials of the artificial spike trains for the corresponding single unit (surrogate CCHs). This procedure was repeated 1000 times. The resulting surrogate CCHs reflected the level of correlation when both units are statistically independent. Finally, average surrogate CCHs were subtracted from the CCHs to yield the corrected CCHs.

Auto-correlation histograms (ACHs) were generated by setting $x_{n_1}^i = x_{n_2}^i$ in Eq. 1 for all i , and corrected by generating artificial spike trains and substituting them for $x_{n_1}^i$ and $x_{n_2}^i$ in Eq. 1 for the calculation of surrogate ACHs.

Cluster-based surrogate test

For statistical purposes, all surrogate CCHs were corrected by their own average to achieve an equally processed set compared to the corrected CCHs, containing just the chance level of correlation (corrected surrogate CCHs). These 1000 corrected surrogate CCHs were then used to run a nonparametric cluster-based surrogate test, a variation of the cluster-based permutation test (Maris and Oostenveld, 2007), to deal with the multiple comparison problem of testing all time lags. Cluster-based tests are tests for dependent variables, which consider contiguous values fulfilling a certain criterion as a cluster. Instead of calculating a test statistic for individual values, the accumulated values of clusters are tested against a null distribution of accumulated cluster values by chance. In our case, adjacent time lags are not independent, since functional coupling of neurons does not follow millisecond precision. We checked significance for a time window of -200ms to 200ms. Calculation of this test statistic involved the following steps:

1. For every time bin the standard deviation of corrected surrogate CCHs was calculated. Subsequently, the corrected CCH and the corrected surrogate CCHs were

2.1. Uniting functional network topology and oscillations

normalized by these standard deviations (z transformation of the data).

2. A z-score of 2 corresponds to a p-value of ~ 0.05 . So we marked all time lags exceeding a z-score of 2 or -2. Please note that the statistical inference is not directly based on this z-score criterion, but rather on the subsequent non-parametric test.
3. As already mentioned, in CCHs neighboring time lags are not independent. Clusters of marked bins were selected on the basis of temporal adjacency.
4. From each corrected surrogate CCH, the largest cluster was selected (independent of the sign) based on its accumulated z-score, creating a distribution of 1000 largest clusters. Since we used each unit as $x_{n_1}^i$ and as $x_{n_2}^i$, we obtained two CCHs per pair of units. These two CCHs are identical, except for being inverted in time. We merged their distributions to a final distribution of the 2000 largest chance clusters.
5. In a final step, cluster-level statistics were calculated. The accumulated z-score of each real cluster was tested against the distribution of biggest clusters occurring by chance. The obtained p-value of each cluster was saved for further corrections.

This procedure was repeated for every CCH. A critical alpha-level of 0.05 was selected. Nevertheless, at this processing step we still have a total alpha-error equal to our set criterion times the number of single unit pairs tested. For complete multiple comparison correction, false discovery rate correction was applied on all found clusters across all compared pairs of single units (Benjamini and Hochberg, 1995) to yield

$$P_{(k)} \leq \frac{k}{m} q \quad (2)$$

where q is our set criterion of 0.05 false positives, m the total number of clusters, $k = 1, \dots, m$, and $P_{(k)}$ are the p-values of all clusters in increasing order. All clusters whose p-values did not fulfill Eq. 2 were rejected. By doing so we achieved a total alpha-level of 0.05 for each dataset.

Network analysis

For every pair of neurons it was evaluated if there were significant troughs or peaks in their CCHs. If there was only a trough or peak with negative (or positive) time lags, this pair was denoted as having a connection from the input to the target (or the target to the input) unit (Figure 2E). In case there were several clusters on both sides of the zero time lag, or a cluster

2.1. Uniting functional network topology and oscillations

straddling the zero time lag, we checked the unsigned maximum peak of the corresponding CCH. If the maximum peak was shifted more than 2 ms to either side, the connection was considered unidirectional, as described before. Otherwise, the connection between the two single units was considered functional bidirectional (Figure 2E), since the units are driven by the circuit at the same time. We systematically varied the maximum peak shift (0-5 ms) for bidirectional classification with little to no change to the results. Repeating this procedure for all pairs of single units led to a binary directed connectivity matrix (Figure 3A).

To characterize brain networks on every scale, network measures from the multidisciplinary field of graph theory were utilized (Rubinov and Sporns, 2010). A network is defined by the nodes (N) and connections between pairs of nodes. In our network nodes represented single units. For all following network measures, n is the number of nodes and l the number of connections. a_{ij} is the connection between nodes i and j : $a_{ij} = 1$ if the link (i, j) exists and $a_{ij} = 0$ otherwise ($a_{ii} = 0$ for all i). Furthermore, we define:

Degree centrality, k_i , is the number of connections to a node i .

$$k_i = \sum_{j \in N} a_{ij} \quad (3)$$

Shortest path length, d_{ij} , is the minimum number of nodes connecting nodes i and j .

$$d_{ij} = \sum_{a_{uv} \in g^{i \leftrightarrow j}} a_{uv} \quad (4)$$

where $g^{i \leftrightarrow j}$ is the shortest path between i and j .

Characteristic path length, L , is the average shortest path length between all pairs of nodes of the network.

$$L = \frac{1}{n(n-1)} \sum_{\substack{i, j \in N \\ i \neq j}} d_{ij} \quad (5)$$

Betweenness centrality, g_i , is the average fraction of shortest paths that pass through node i .

2.1. Uniting functional network topology and oscillations

$$g_i = \frac{1}{(n-1)(n-2)} \sum_{\substack{h,j \in N \\ h \neq j, h \neq i, j \neq i}} \frac{\rho_{hj}^{(i)}}{\rho_{hj}} \quad (6)$$

where ρ_{hj} is the number of shortest paths between h and j , and $\rho_{hj}^{(i)}$ is the number of shortest paths between h and j that pass through i .

Clustering coefficient of the network, C , is the average fraction of existing to maximal possible interconnections between all directly connected nodes to node i .

$$C = \frac{1}{n} \sum_{i \in N} \frac{2t_i}{k_i(k_i - 1)} \quad (7)$$

Where k_i are all connected neighbors to node i and t_i is the number of links between them.

Small-worldness, SW , is the ratio of C and L each normalized by the same measurements for a size matched random network.

$$SW = \frac{C/C_{rand}}{L/L_{rand}} \quad (8)$$

Small-world networks are formally defined as networks that are significantly more clustered than random networks, yet have approximately the same characteristic path length as random networks (Watts and Strogatz, 1998).

Modularity, Q , is the proportion of all links within modules M with links between modules, when the network is fully subdivided into non-overlapping modules in a way that maximizes the number of within-group connections and minimizes the number of between-group connections.

$$Q = \sum_{u \in M} \left[e_{uu} - \left(\sum_{v \in M} e_{uv} \right)^2 \right] \quad (9)$$

where e_{uv} is the fraction of all links that connect nodes in module u with nodes in module v .

Rich-club coefficient, R , at degree k is the fraction of connections between all nodes of degree k or higher, with respect to the maximum possible number of such connections.

$$R(k) = \frac{2E_{>k}}{N_{>k}(N_{>k} - 1)} \quad (10)$$

2.1. Uniting functional network topology and oscillations

where $E_{>k}$ is the number of connections among the $N_{>k}$ nodes having degree of k or higher (Colizza et al., 2006). To reduce inaccuracy for large degrees we calculated the rich-club coefficient only in degree bins containing at least 5 single units ($N_k \geq 5$).

Statistics for network measures

For statistical purposes we created two types of surrogate network sets per dataset (1000 partitions each). All surrogate networks were created by shuffling the connectivity matrix. Since connectivity is a function of distance (Smith and Kohn, 2008; Gerhard et al., 2011), distance dependency was reflected in our surrogate data. During shuffling, the number of connections for single units on the same electrode, the same array, the same cortical area, and the different inter-area connections were always held constant (Figure 3B). For all surrogate networks, the total number of single units, number of connections, and the distance-dependent ratio of bi- and uni-directional connections were kept as similar as possible to the original connectivity matrix with only the required network parameter shuffled. We used these sets of surrogate networks to test the small-world coefficient, the degree centrality distribution, and the betweenness centrality distribution. Statistical testing of the rich-club coefficient and conservative testing of modularity requires surrogate networks with a matched degree centrality distribution. To this end, we generated a second set of surrogate networks with the degree distribution preserved. One issue that could arise due to shuffling is that the connectivity matrix of some units or groups of units could become disconnected from the main part of the network, since the calculation of most network measures requires a fully connected, not segregated, network. For this purpose, each surrogate network was tested for segregation into different components. If a network was segregated, it was discarded and the process repeated until 1000 non-segregated networks were generated.

To determine if the degree, the betweenness centrality distribution, or the rich-club level were significantly different to surrogate networks, we used a nonparametric cluster-based permutation test (Maris and Oostenveld, 2007). Briefly, this test evaluates the t-statistic (independent samples) between centrality or rich-club distributions and their surrogate distributions over all data points exceeding a critical alpha-level set to 0.05. In a second step, adjacent degree, betweenness values, or rich-club coefficients exceeding the set alpha-level are considered as clusters, extracted, and their t-value summed. A test

2.1. Uniting functional network topology and oscillations

distribution was generated by randomly permuting the centrality or rich-club distributions across recording days and monkeys with the corresponding surrogate distributions by randomly reassigning them to one of the two groups while maintaining the group size. For each partition (1000 partitions) the t-statistics and clustering was repeated. From every partition the largest cluster-level statistic was used to generate a largest chance cluster distribution. For each real cluster-level statistic a nonparametric statistical test was performed by calculating a p-value under the largest chance cluster distribution. Thus, the multiple comparisons for each sample are replaced by a single comparison, replacing the need to make multiple comparisons.

Since some electrode pairs between F5 and M1 are closer than some other pairs within M1 for monkey M, we repeated statistics for network measures for all datasets from monkey M with physical distance dependent shuffling instead of the above mentioned categories such as “same electrode”, “same array,” and “same area”. To this end, we calculated the pairwise physical distance between all pairs of electrodes based on an anatomical diagram (Figure 1B) and defined distance groups with a stepsize of 3.6mm including 0mm as one group. The physical distance between AIP and the two other areas is misleading, since the neuronal axons have to pass the central sulcus. Therefore, we set all distances between AIP and the two other areas as a separate maximum distance group. Note that we had to define groups to be able to shuffle connections. Nevertheless, the categorical distance dependent shuffling was subdivided into 8 groups, which is more conservative than the 6 groups defined in the original analysis. All statistics for network measures gave nearly identical results, with no case where a measure was significant when it was not for categorical distance dependent shuffling, and vice versa for non-significant measures. In addition, the normalized rich-club coefficient, which depends on the surrogate networks, was highly correlated ($r = 0.98$) between the two different ways of distance dependent shuffling.

Equal rate model

For validation of the estimates of directed functional connectivity, as well as to check for a possible bias in the detected network topology obtained using CCHs, we modeled artificial directed neuronal networks with the same firing rate distribution as the recorded single units. Two sets of networks were generated, one simple network (SN) set with normally distributed connectivity and one complex network (CN) set with heterogeneously distributed

2.1. Uniting functional network topology and oscillations

connectivity, and in agreement with previous studies both with weak connection strength between neuronal pairs (Cohen and Kohn, 2011).

For each simulated neuron, artificial spike trains were generated with Poisson distributed firing and an average rate randomly drawn from the real firing rate distribution. For the SN set, the number of connections from each neuron to other neurons was drawn randomly from a Gaussian distribution (mean: 5.22, SD: 3.214), mirroring the average degree centrality distribution of surrogate networks. For the complex network set (CN), the number of connections followed precisely the EXPTPL model for the average degree centrality distribution of the measured networks (Figure 4A), with a weak rich-club and small-world topology. In case one neuron was connected to another, spikes were added in a probabilistic manner for a certain amount of time, starting with time point $t + 1$ in ms relative to the spike event, reflecting the axonal delay. The network was updated every millisecond, allowing for multiple interactions. Gamma functions were used as temporal transfer kernels, given by

$$f(t|a, b) = \frac{1}{b^a \Gamma(a)} t^{a-1} e^{-\frac{t}{b}} \quad (11)$$

where f is the probability of an additional spike appearing, t is time in ms, a is a constant set to 5 and b is randomly varied between 0 and 3 (Figure 2 – Figure Supplement 3A). The integral of each gamma kernel was set to 0.02, reflecting the connection strength. Since we added spikes to the network, which increases the average firing rates, we lowered the starting rates by a factor and repeated the process until the average rate resembled the rate before adding the connections. As a criterion for similarity we correlated the randomly drawn rates with the network rates and stopped when the residual error was below 0.005. For the results in Figure 2 – Figure Supplement 3 and Figure 4 – Figure Supplement 1 we did not vary the connection strength in order to avoid interaction effects between connection strength and firing rate. However, we varied connection strength randomly between 0.005 and 0.035 with no detectable change to the results. Alternatively, we used a Boxcar kernel (20 ms, integral: 0.02) instead of gamma functions as transfer kernel, which did not degrade the results of this model.

For both sets of networks (SN and CN), ten artificial networks with 100 neurons were calculated and processed identically to the real data. Signal detection theory was used to evaluate detectability of connections based on significant CCH peaks or troughs with the

2.1. Uniting functional network topology and oscillations

originally modeled networks as a reference. Each pairing was classified into one of four categories: “*Hit*”, if a connection was correctly detected, “*Miss*”, if a connection was not detected, “*Correct rejection*” (CR), if a non-existing connection was detected as no connection, and “*False Alarm*” (FA), if a non-existing connection was detected as a connection.

Subsampling model

We generated an artificial neuronal plane with random (Poisson distributed), distance-dependent connectivity density based on our empirically collected data (Figure 3B). We modeled 2 cortical areas, each divided into 5 sub-regions coverable by an array, each sub-region covered with 160 electrode positions, and 20 single units per electrode, giving a total of 32,000 neurons. Figure 4 – Figure Supplement 2A shows the degree centrality distribution of the full network with an average degree of 3000 and a standard deviation of 70.

Next, we randomly selected 12 subsamples from the neuronal plane with exactly the number of neurons detected as in the real datasets. Subsampling was done with the restriction that always both areas were chosen, with 2 array sub-regions per area and 32 electrode positions per sub-region, reflecting the real recording configuration in most of the datasets. Subsampled networks were then analyzed with the same complex network measures as the real data.

To address the problem that subsampling could artificially cause a heavy tailed degree centrality distribution, even if the underlying connectivity is random, as described in (Han et al., 2005), we had a closer look at the parameters mentioned in this study. The average degree of their analyzed networks was 2.19 (SD = 0.45, min = 1.84, max = 2.98), in contrast to our average (non-normalized) degree of 8.28 (SD = 5.73, min = 3.87, max = 25.59). Note that the highest average degree of their analyzed networks was smaller than the lowest average degree of our analyzed networks. More importantly, the underlying networks of their study were strongly fragmented into components (min = 70, max = 591 components), while we excluded all single units which were not part of the largest component, resulting in one component for analysis, while their largest average component size was 20.2. Our network analysis was done on average on 70 single units (min 30, max 148 single units). Based on these different network parameters we concluded that the detected topology, in particular falsely detected power law degree distribution, could be due to the fragmentation into different components. To evaluate this, we created neuronal

2.1. Uniting functional network topology and oscillations

planes with distance dependent connection density of 1/5, 1/4, 1/3, 1/2, 1, 2, 3, 4, and 5 times of the empirically collected data. After subsampling, we estimated the goodness of fit for the power law model to the degree centrality distribution, the size of the largest component relative to the whole network, and the level of compartmentalization, described by

$$\text{Compartmentalization} = \frac{P - 1}{N - 1} \quad (12)$$

where N is the number of neurons in the network and P the number of separate components (Figure 4 – Figure Supplement 2C).

Frequency analyses

We estimated the oscillatory behavior of significant connections of single units (according to CCHs) and the spiking of single units themselves (Bair et al., 1994; Mureşan et al., 2008) (according to ACHs). Since different oscillation frequencies could be present, we computed power spectra of all corrected CCHs and ACHs (Mureşan et al., 2008). The power spectrum gives the magnitude of a signal as a function of frequency. To avoid distortions by sharp peaks with small delays that are occasionally present in CCHs (Fujisawa et al., 2008), which cause a broad band increase in power due to their impulse like properties, we cut out the time range from -5ms to 5ms and interpolated the segment linearly. Importantly, sharp peaks were only removed for spectral analyses and not for functional connectivity analyses. Frequency spectra were computed using a discrete Fourier transform algorithm (Siegel et al., 2009) (100 logarithmically scaled frequencies from 3 to 100 Hz). Note that computing power spectra of CCHs and ACHs instead of raw spike trains reduced the influence of firing rate on the power spectrum as well as the problem of frequency leakage due to the binary properties of the spike train (Bair et al., 1994). In analyzing such a large range of frequencies we had to take the specific characteristics of CCHs into account. Underlying oscillation frequencies in physiology are not phase stable, which leads to a limited number of side lobes in the CCH or ACH. The number of side lobes are also strongly frequency dependent, which makes the ideal window length for Fourier transformation around the 0 time lag frequency dependent. We used Hanning windows of four times the frequency of interest period (with a maximum of 1000ms and a minimum of 150ms) aligned on the 0 time bin of the CCHs (Figure 5 – Figure Supplement 1A), resulting in approximately 1/frequency and half octave

2.1. Uniting functional network topology and oscillations

spectro-temporal bandwidth. Each frequency bin was divided by its window length for correct scaling of all frequency bins. To determine significance, we repeated spectral analysis on the corrected surrogate CCHs and ACHs, subtracted their mean spectra from the corresponding spectra of real data and used a cluster-based surrogate test as described before to evaluate the significance of the underlying frequencies in the CCHs.

Spectral analysis of the ACHs differed in one point. Hanning windows covering only one half of the ACHs (with a maximum of 500ms and a minimum of 75ms) aligned on the 0 time lag were used (Figure 5 – Figure Supplement 1B). By doing so, an accurate measure of the full frequency range with little distortion of refractory effects present in ACHs (Mureşan et al., 2008) was obtained.

Oscillatory vs non-oscillatory synchronization model

We generated pairs of neurons with 600 trials and a trial length of 3.1 seconds, similar to our recorded data. Spike trains of neurons were generated as a probabilistic process. In case of oscillatory firing neurons, the probability function was a 20Hz sinusoid. For non-oscillating neurons, we first randomized the 20Hz sinusoid, in a second step filtered it with a non-causal 50 Hz low-pass filter (Butterworth filter, 4th order) in order to produce a similar decay in spiking probability, and in a last step the filtered probability vector was variance matched with the 20Hz sinusoid to have a maximum degree matching between the two kinds of probability functions. For each trial the same probability function was used for both neurons with a spiking probability of 0.05 per ms to stay in a physiological range. Independent Poisson distributed noise was added to both neurons representing background stochastic firing, resulting in an average rate of around 5 Hz per neuron. Varying the different parameters within physiological ranges did not alter the results. To simulate different degrees of coupling strengths we systematically varied the trial-wise time offset in spiking of the pair of neurons to each other from completely synchronized to a jitter of a complete cycle (50ms) in steps of 1 ms.

2.1. Uniting functional network topology and oscillations

References

- Bair W, Koch C, Newsome W, Britten K (1994) Power spectrum analysis of bursting cells in area MT in the behaving monkey. *J Neurosci* 14:2870–2892.
- Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci* 21:1676–1697.
- Barabási A-L, Oltvai ZN (2004) Network biology: understanding the cell's functional organization. *Nat Rev Genet* 5:101–113. doi:10.1038/nrg1272.
- Bassett DS, Bullmore E (2006) Small-World Brain Networks. *The Neuroscientist* 12:512–523. doi:10.1177/1073858406293182.
- Bastos AM, Schoffelen J-M (2016) A Tutorial Review of Functional Connectivity Analysis Methods and Their Interpretational Pitfalls. *Front Syst Neurosci* 9:413. doi:10.1126/science.1139597.
- Bastos AM, Vezoli J, Bosman CA, Schoffelen J-M, Oostenveld R, Dowdall JR, De Weerd P, Kennedy H, Fries P (2015) Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron* 85:390–401. doi:10.1016/j.neuron.2014.12.018.
- Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B (Methodological)*:289–300.
- Bonifazi P, Goldin M, Picardo MA, Jorquera I, Cattani A, Bianconi G, Represa A, Ben-Ari Y, Cossart R (2009) GABAergic hub neurons orchestrate synchrony in developing hippocampal networks. *Science* 326:1419–1424. doi:10.1126/science.1175509.
- Bosman CA, Schoffelen J-M, Brunet N, Oostenveld R, Bastos AM, Womelsdorf T, Rubehn B, Stieglitz T, De Weerd P, Fries P (2012) Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* 75:875–888. doi:10.1016/j.neuron.2012.06.037.
- Brovelli AA, Ding MM, Ledberg AA, Chen YY, Nakamura RR, Bressler SLS (2004) Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by Granger causality. *Proc Natl Acad Sci USA* 101:9849–9854. doi:10.1073/pnas.0308538101.
- Bullmore E, Sporns O (2009) Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* 10:186–198. doi:10.1038/nrn2575.
- Buzsáki G (2010) Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron* 68:362–385. doi:10.1016/j.neuron.2010.09.023.
- Buzsáki G, Mizuseki K (2014) The log-dynamic brain: howskewed distributions affect networkoperations. *Nat Rev Neurosci* 15:264–278. doi:10.1038/nrn3687.

2.1. Uniting functional network topology and oscillations

- Buzsáki G, Wang X-J (2012) Mechanisms of Gamma Oscillations. *Annu Rev Neurosci* 35:203–225. doi:10.1146/annurev-neuro-062111-150444.
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. *Nature* 487:51–56. doi:10.1038/nature11129.
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nature Publishing Group* 14:811–819. doi:10.1038/nn.2842.
- Colizza V, Flammini A, Serrano MA, Vespignani A (2006) Detecting rich-club ordering in complex networks. *Nat Phys* 2:110–115. doi:10.1038/nphys209.
- Dean HL, Hagan MA, Pesaran B (2012) Only coherent spiking in posterior parietal cortex coordinates looking and reaching. *Neuron* 73:829–841. doi:10.1016/j.neuron.2011.12.035.
- Engel AK, Fries P (2010) Beta-band oscillations--signalling the status quo? *Curr Opin Neurobiol* 20:156–165. doi:10.1016/j.conb.2010.02.015.
- Fries P (2009) Neuronal Gamma-Band Synchronization as a Fundamental Process in Cortical Computation. *Annu Rev Neurosci* 32:209–224. doi:10.1146/annurev.neuro.051508.135603.
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11:823–833. doi:10.1038/nn.2134.
- Gerhard F, Pipa G, Lima B, Neuenschwander S, Gerstner W (2011) Extraction of Network Topology From Multi-Electrode Recordings: Is there a Small-World Effect? *Front Comput Neurosci* 5 doi:10.3389/fncom.2011.00004.
- Gilbert CD, Sigman M (2007) Brain states: top-down influences in sensory processing. *Neuron* 54:677–696. doi:10.1016/j.neuron.2007.05.019.
- Gollo LL, Mirasso C, Sporns O, Breakspear M (2014) Mechanisms of Zero-Lag Synchronization in Cortical Motifs Gutkin BS, ed. *PLoS Comput Biol* 10:e1003548. doi:10.1371/journal.pcbi.1003548.s011.
- Gozani SN, Miller JP (1994) Optimal discrimination and classification of neuronal action potential waveforms from multiunit, multichannel recordings using software-based linear filters. *Biomedical Engineering, IEEE Transactions on* 41:358–372.
- Gregoriou GG, Gotts SJ, Desimone R (2012) Cell-Type-Specific Synchronization of Neural Activity in FEF with V4 during Attention. *Neuron* 73:581–594. doi:10.1016/j.neuron.2011.12.019.
- Gregoriou GG, Gotts SJ, Zhou H, Desimone R (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* 324:1207–1210.
- Haegens S, Nacher V, Hernandez A, Luna R, Jensen O, Romo R (2011) Beta oscillations in the

2.1. Uniting functional network topology and oscillations

- monkey sensorimotor network reflect somatosensory decision making. *Proc Natl Acad Sci USA* 108:10708–10713. doi:10.1073/pnas.1107297108.
- Han J-DJ, Dupuy D, Bertin N, Cusick ME, Vidal M (2005) Effect of sampling on topology predictions of protein-protein interaction networks. *Nat Biotechnol* 23:839–844. doi:10.1038/nbt1116.
- Hilgetag CC, Burns GA, O'Neill MA, Scannell JW, Young MP (2000) Anatomical connectivity defines the organization of clusters of cortical areas in the macaque monkey and the cat. *Philos Trans R Soc Lond, B, Biol Sci* 355:91–110. doi:10.1098/rstb.2000.0551.
- König P, Engel AK, Singer W (1995) Relation between oscillatory activity and long-range synchronization in cat visual cortex. *Proc Natl Acad Sci USA*.
- 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. doi:10.1016/j.neuron.2009.12.010.
- Lundqvist M, Rose J, Herman P, Brincat SL, Buschman TJ, Miller EK (2016) Gamma and Beta Bursts Underlie Working Memory. *Neuron*:1–14. doi:10.1016/j.neuron.2016.02.028.
- 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.
- Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods* 164:177–190. doi:10.1016/j.jneumeth.2007.03.024.
- Maris E, Schoffelen J-M, Fries P (2007) Nonparametric statistical testing of coherence differences. *Journal of Neuroscience Methods* 163:161–175. doi:10.1016/j.jneumeth.2007.02.011.
- 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. doi:10.1523/JNEUROSCI.1714-15.2015.
- Moore GP, Segundo JP, Perkel DH, Levitan H (1970) Statistical signs of synaptic interaction in neurons. *Biophys J* 10:876–900. doi:10.1016/S0006-3495(70)86341-X.
- Mureşan RC, Jurjuţ OF, Moca VV, Singer W, Nikolic D (2008) The oscillation score: an efficient method for estimating oscillation strength in neuronal activity. *J Neurophysiol* 99:1333–1353. doi:10.1152/jn.00772.2007.
- Murthy VN, Fetz EE (1996) Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior.
- Musial PG, Baker SN, Gerstein GL, King EA, Keating JG (2002) Signal-to-noise ratio improvement in multiple electrode recording. *Journal of Neuroscience Methods* 115:29–43.
- Nácher V, Ledberg A, Deco G, Romo R (2013) Coherent delta-band oscillations between

2.1. Uniting functional network topology and oscillations

cortical areas correlate with decision making. *Proc Natl Acad Sci USA* 110:15085–15090. doi:10.1073/pnas.1314681110.

Nigam S, Shimono M, Ito S, Yeh FC, Timme N, Myroshnychenko M, Lapish CC, Tosi Z, Hottowy P, Smith WC, Masmanidis SC, Litke AM, Sporns O, Beggs JM (2016) Rich-Club Organization in Effective Connectivity among Cortical Neurons. *J Neurosci* 36:670–684. doi:10.1523/JNEUROSCI.2177-15.2016.

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. doi:10.1038/nature14273.

Pesaran B, Nelson MJ, Andersen RA (2008a) Free choice activates a decision circuit between frontal and parietal cortex. *Nature* 453:406–409. doi:10.1038/nature06849.

Pesaran B, Nelson MJ, Andersen RA (2008b) Free choice activates a decision circuit between frontal and parietal cortex. *Nature* 453:406–409. doi:10.1038/nature06849.

Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5:805–811. doi:10.1038/nn890.

Quiroga RQ, Nadasdy Z, Ben-Shaul Y (2004) Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput* 16:1661–1687. doi:10.1162/089976604774201631.

Ramalingam N, McManus JNJ, Li W, Gilbert CD (2013) Top-down modulation of lateral interactions in visual cortex. *J Neurosci* 33:1773–1789. doi:10.1523/JNEUROSCI.3825-12.2013.

Research COGFTUOAINAB, National Research Council, Studies DOEAL, Institute for Laboratory Animal Research (2003) Guidelines for the Care and Use of Mammals in Neuroscience and Behavioral Research. National Academies Press.

Roelfsema PR, Engel AK, König P, Singer W (1997) Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385:157–161. doi:10.1038/385157a0.

Rubinov M, Sporns O (2010) Complex network measures of brain connectivity: Uses and interpretations. *Neuroimage* 52:1059–1069. doi:10.1016/j.neuroimage.2009.10.003.

Salazar RF, Dotson NM, Bressler SL, Gray CM (2012) Content-specific fronto-parietal synchronization during visual working memory. *Science* 338:1097–1100. doi:10.1126/science.1224000.

Schaffelhofer S, Scherberger H (2016) Object vision to hand action in macaque parietal, premotor, and motor cortices. *Elife* 5. doi:10.7554/eLife.15278.

Schroeter MS, Charlesworth P, Kitzbichler MG, Paulsen O, Bullmore ET (2015) Emergence of Rich-Club Topology and Coordinated Dynamics in Development of Hippocampal

2.1. Uniting functional network topology and oscillations

- Functional Networks In Vitro. *J Neurosci* 35:5459–5470. doi:10.1523/JNEUROSCI.4259-14.2015.
- Shimono M, Beggs JM (2014) Functional Clusters, Hubs, and Communities in the Cortical Microconnectome. *Cerebral Cortex* doi:10.1093/cercor/bhu252.
- Siegel M, Warden MR, Miller EK (2009) Phase-dependent neuronal coding of objects in short-term memory. *Proc Natl Acad Sci USA* 106:21341–21346.
- Smith MA, Kohn A (2008) Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex. *J Neurosci* 28:12591–12603. doi:10.1523/JNEUROSCI.2929-08.2008.
- Stephan KE, Hilgetag CC, Burns GA, O'Neill MA, Young MP, Kötter R (2000) Computational analysis of functional connectivity between areas of primate cerebral cortex. *Philos Trans R Soc Lond, B, Biol Sci* 355:111–126. doi:10.1098/rstb.2000.0552.
- van den Heuvel MP, Kahn RS, Goñi J, Sporns O (2012) High-cost, high-capacity backbone for global brain communication. *Proc Natl Acad Sci USA* 109:11372–11377. doi:10.1073/pnas.1203593109.
- van den Heuvel MP, Sporns O (2013a) Network hubs in the human brain. *Trends in Cognitive Sciences* 17:683–696. doi:10.1016/j.tics.2013.09.012.
- van den Heuvel MP, Sporns O (2013b) An anatomical substrate for integration among functional networks in human cortex. *J Neurosci* 33:14489–14500. doi:10.1523/JNEUROSCI.2128-13.2013.
- Van Eck NJ, Waltman L (2007) VOS: A new method for visualizing similarities between objects.
- Van Eck NJ, Waltman L, Dekker R (2010) A comparison of two techniques for bibliometric mapping: Multidimensional scaling and VOS. *Journal of the ...*
- Vicente R, Gollo LL, Mirasso CR, Fischer I, Pipa G (2008) Dynamical relaying can yield zero time lag neuronal synchrony despite long conduction delays. *Proc Natl Acad Sci USA* 105:17157–17162. doi:10.1073/pnas.0809353105.
- Watts DJ, Strogatz SH (1998) Collective dynamics of “small-world” networks. *Nature* 393:440–442. doi:10.1038/30918.
- Womelsdorf T, Valiante TA, Sahin NT, Miller KJ, Tiesinga P (2014) Dynamic circuit motifs underlying rhythmic gain control, gating and integration. *Nature Publishing Group* 17:1031–1039. doi:10.1038/nn.3764.
- Yu S, Huang D, Singer W, Nikolic D (2008) A Small World of Neuronal Synchrony. *Cerebral Cortex* 18:2891–2901. doi:10.1093/cercor/bhn047.

Supplementary Figures and Tables

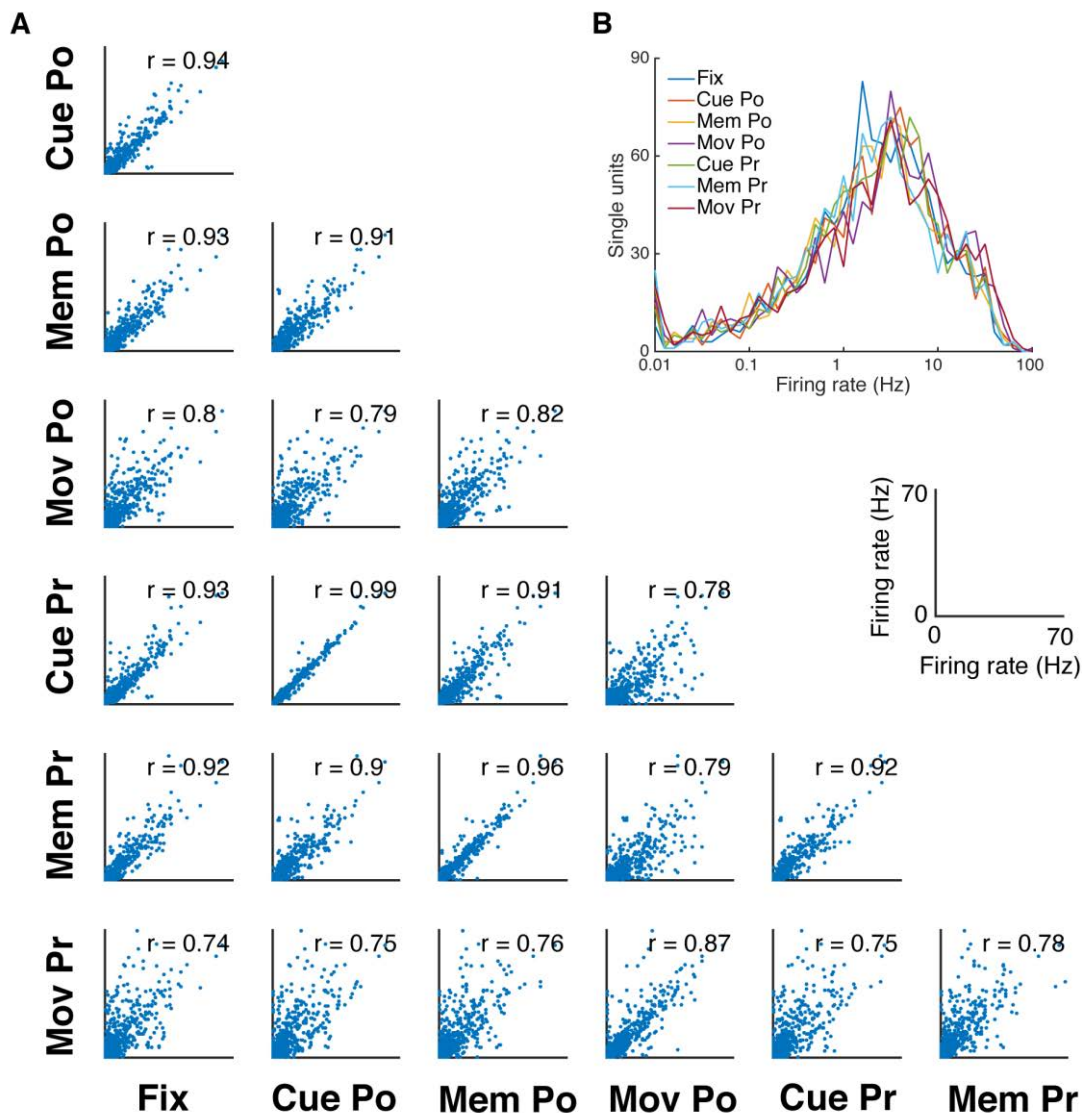


Figure 1 – Figure Supplement 1. Firing rate distribution and stability across task epochs and conditions. (A) Scatter plots of all pairs of condition- and epoch-wise average firing rates of all recorded single units of all datasets (fixation (Fix), cue power (Cue Po), memory power (Mem Po), movement power (Mov Po), cue precision (Cue Pr), memory precision (Mem Pr) and movement precision (Mov Pr)). Due to the high degree of similarity, free-choice and instructed trials were collapsed. In each panel the corresponding correlation coefficient is displayed (mean $r = 0.85$, $SD = 0.08$; for all: $p < 0.001$). (B) Firing rate distribution averaged as in A, displayed on a logarithmic x-axis. The firing rate distribution is very similar for all conditions and epochs and close to log-normal.

2.1. Uniting functional network topology and oscillations

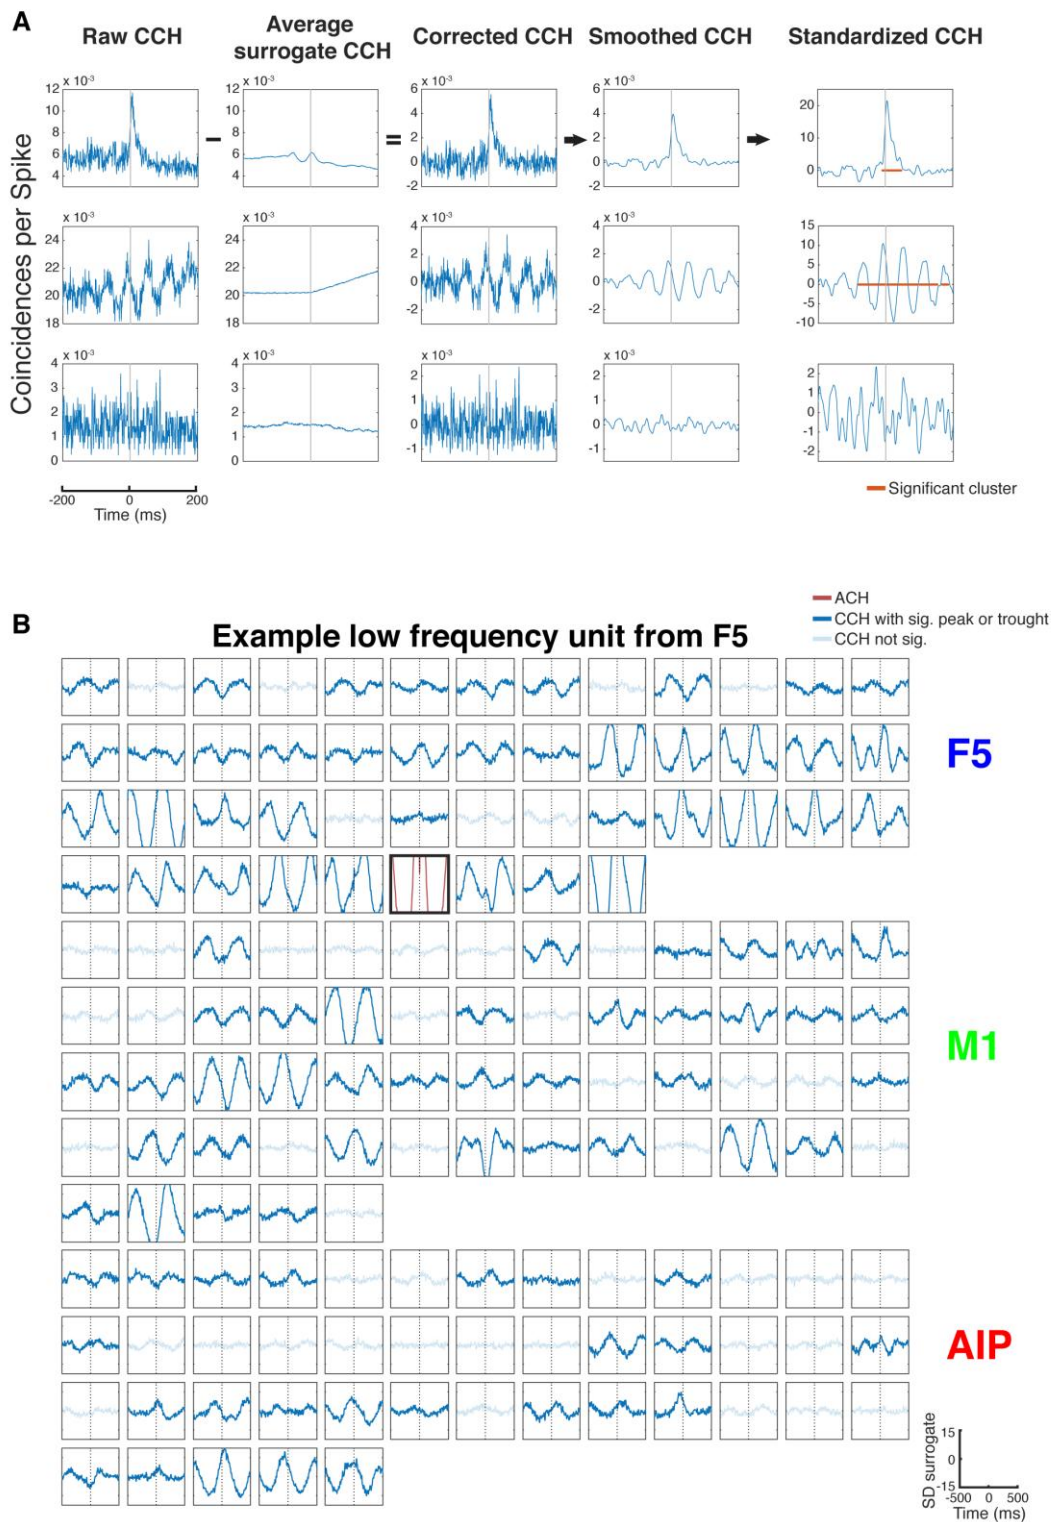


Figure 2 – Figure Supplement 1. CCH processing and statistics, and all connections of an example unit oscillatory synchronized in the low frequency range. (A) Processing steps of three example CCHs. From left to right: illustration of the processing steps involving surrogate subtraction, smoothing, and cluster statistics to evaluate if a peak or trough in a CCHs was significant. From top to bottom: A CCH with one significant peak, a CCH with multiple significant peaks and troughs having an underlying frequency in the beta range, and a CCH with no significant peak or trough. (B) An examples of all CCHs (small panels) and the ACH of one unit with all other units of one dataset of a unit communicating and oscillating in the low frequency range. The ACH is displayed in red, significant connections are indicated by dark lines in CCHs and not significant connections as transparent lines. Directionality information, which is also derived from the CCHs, is not represented.

2.1. Uniting functional network topology and oscillations

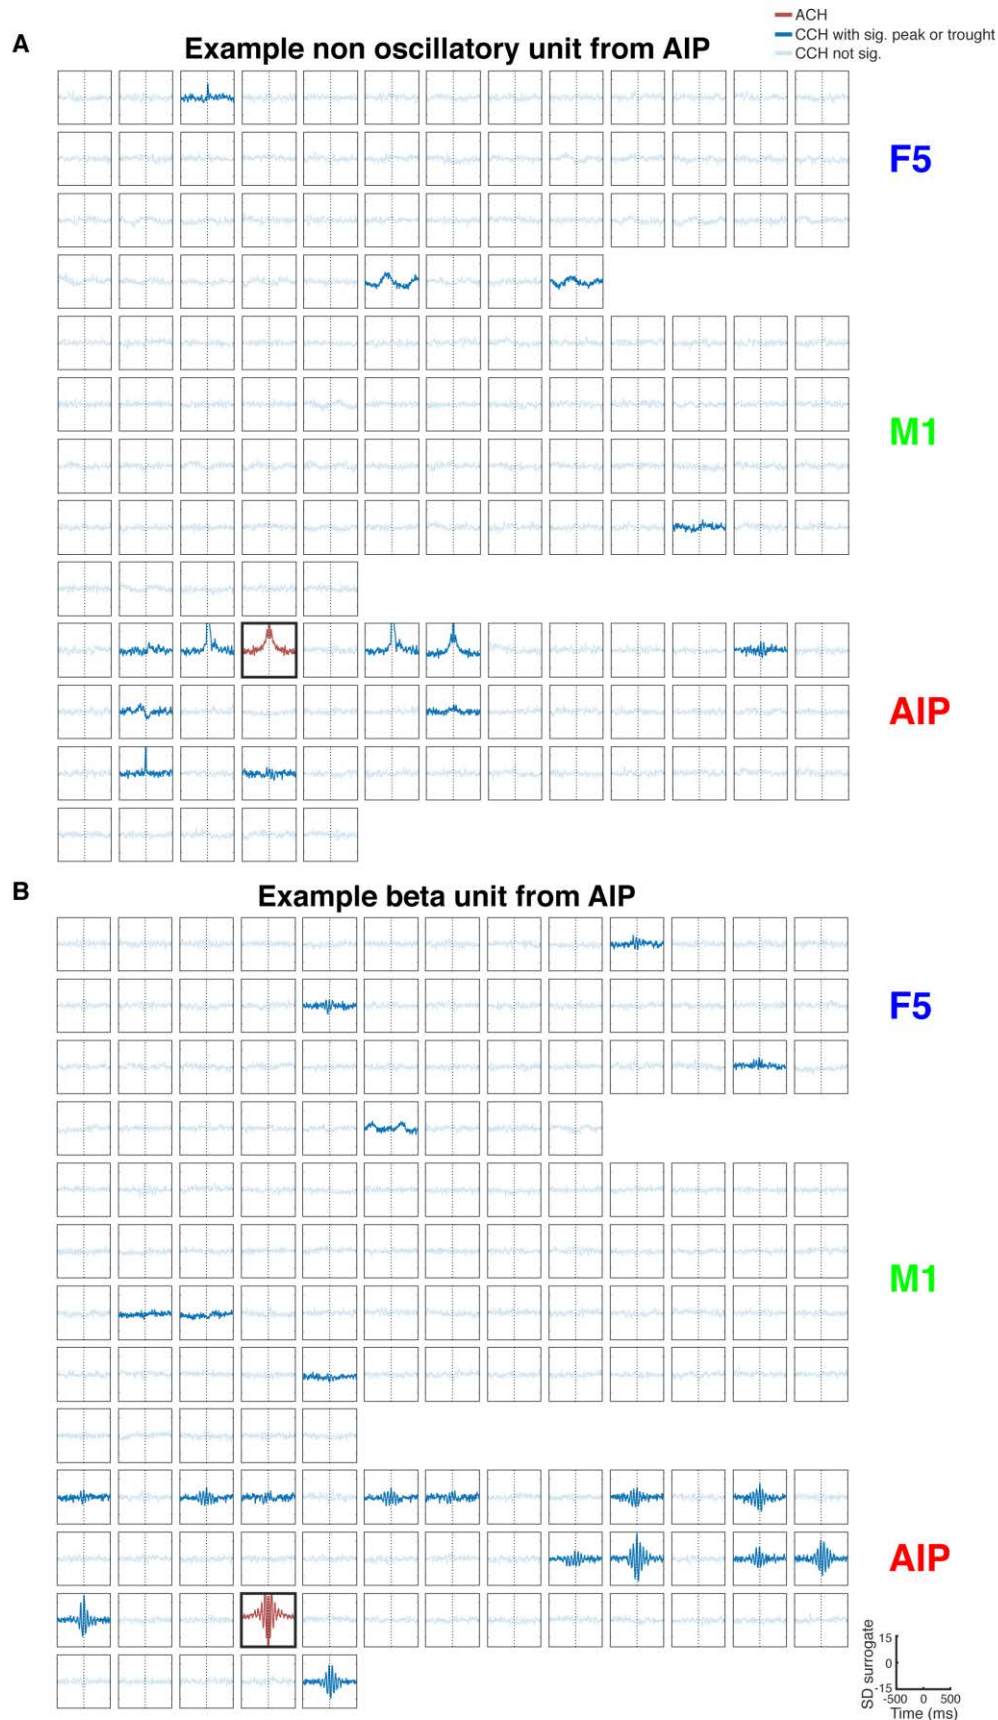


Figure 2 – Figure Supplement 2. All connections of two example units, one non-oscillatory synchronized and one oscillatory synchronized in the beta range. (A) Same as in Figure 2 – Supplementary Figure 1B, but for a non-oscillatory synchronized unit. (B) Same as in Figure 2 – Supplementary Figure 1B, but for a unit communicating and oscillating in the beta range.

2.1. Uniting functional network topology and oscillations

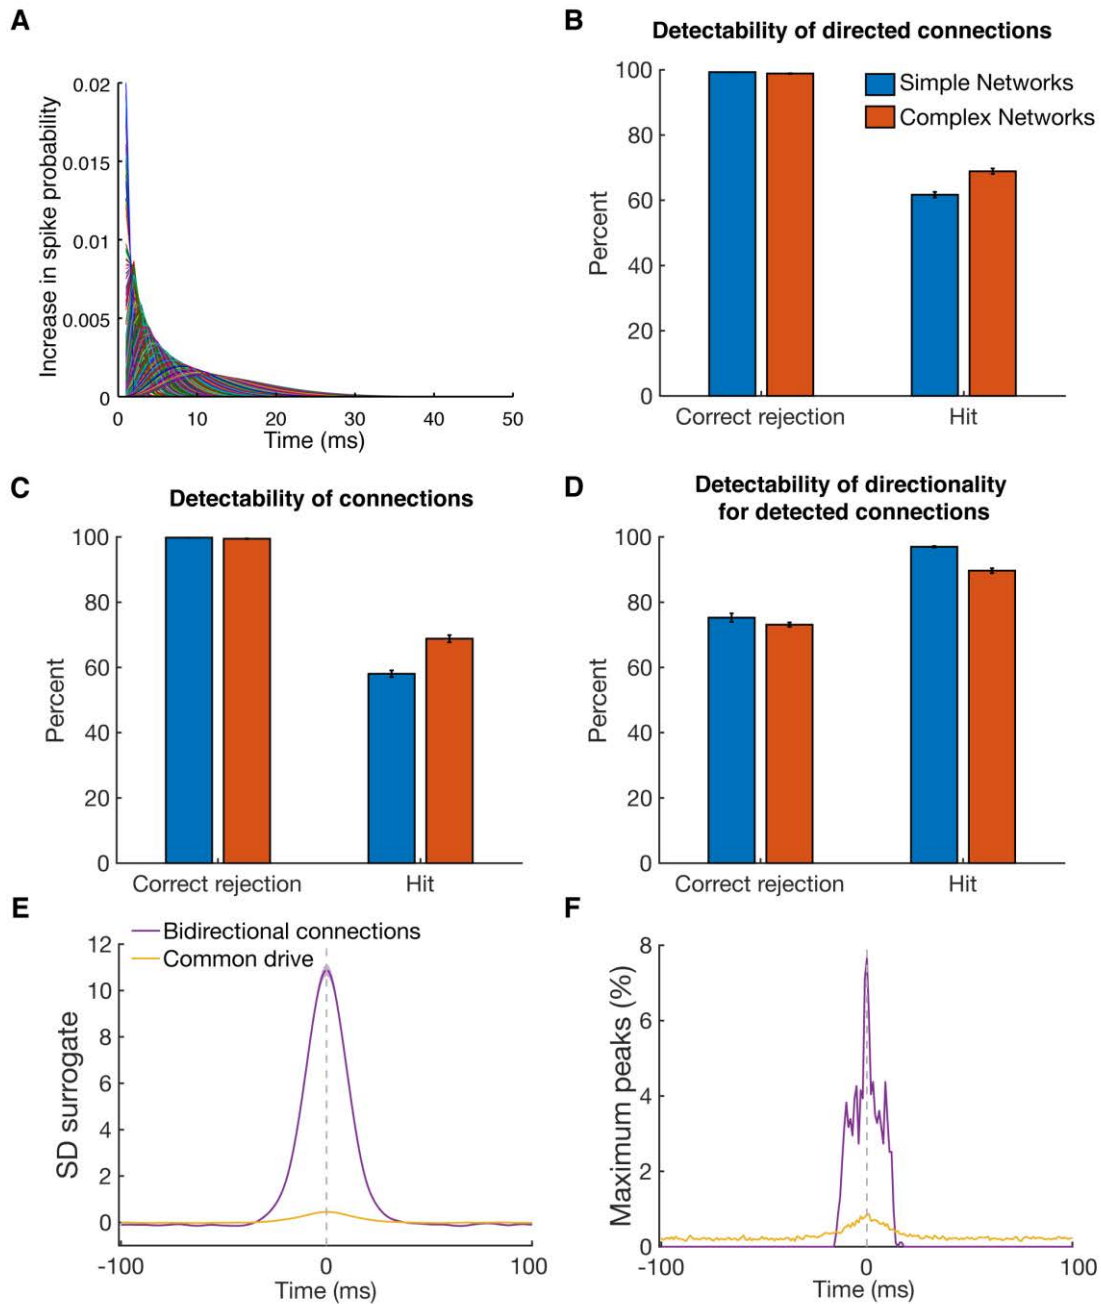


Figure 2 – Figure Supplement 3. Detectability of directed functional connections using equal rate model simulations. (A) Transfer kernels of one modeled dataset. Gamma functions with different maxima and lengths were used as temporal transfer kernels. The area under the curve was always normalized to 0.02. (B) Histogram of detectability of directed connections. Average number of correct rejections and hits are shown for 10 simulated simple networks (SN) and 10 simulated complex networks. Error bars show the standard error across simulated networks. (C) Same as in B, but for detectability of connections. Any directional information was ignored and it was just estimated if a connection between two units was detected or not. (D) Same as in B, but for detectability of directionality for detected connections. The percent of correct rejections and hits is only for the correctly detected connections as displayed in B, thus a pure evaluation of directionality detectability unbiased by connection detectability. (E) Average CCHs for bidirectional connections and common drive pairs of all 20 simulations. The data was pooled, since no considerable difference between the two types of simulations was found. All simulated pairs of both groups are included irrespective of whether they were detected as significant. Error bars show the standard error across CCHs. Note that even though the average peak is at the zero time lag, many pairs had peaks on either side of the zero time lag. (F) Maximum peak count of bidirectional and common drive pairs (for each ms bin) displayed in E. In case CCHs had two peaks or just showed noise fluctuations, only the time lag of the maximum value was considered in order to avoid preselection biases.

2.1. Uniting functional network topology and oscillations

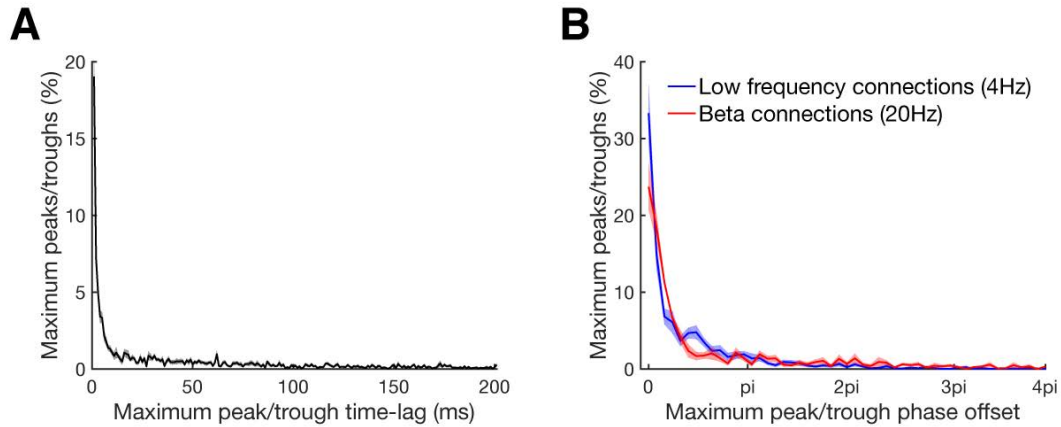


Figure 2 – Figure Supplement 4. Maximum peak or trough time and phase lag distributions. (A) Maximum peak or trough time lag distribution of all significant connections relative to the zero time lag. In case that more than one significant cluster was detected, only the cluster with the highest absolute value was considered. For bidirectional connections time lags were considered for both directions. Line shadings show standard error across datasets. (B) Maximum peak or trough phase relative to the zero time lag for all connections with significant underlying oscillation classified by a significant peak in their corresponding frequency spectra. Results are shown separately for beta at 20Hz (red) and low frequency at 4Hz (blue) oscillations. Note that 4π (two cycles) corresponds to 100ms for beta and to 500ms for low frequency oscillations. Line shadings show standard error across datasets.

2.1. Uniting functional network topology and oscillations

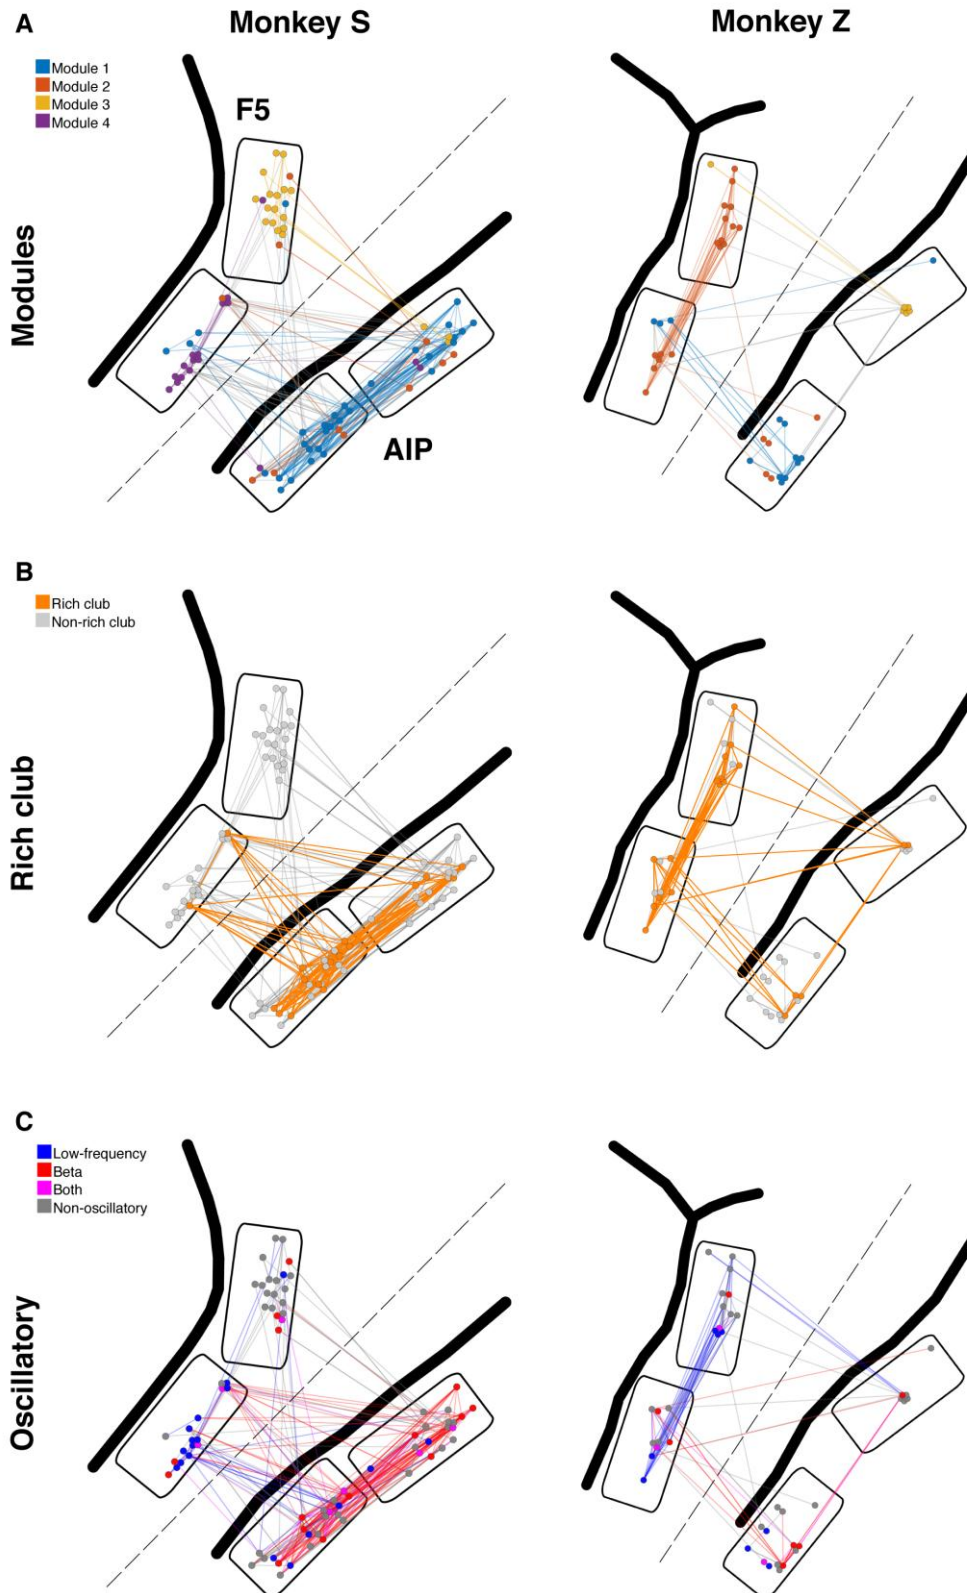


Figure 3 – Figure Supplement 1. Example anatomical networks from Monkey S and Z. Since no data were recorded from area M1 for these monkeys, the F5 and AIP arrays are presented closer together than in reality for better illustration (dashed line marks anatomical discontinuity). (A) Each node colored based on the module, as in Figure 3C. (B) Nodes and connections colored based on rich-clubness, as in Figure 4E. (C) Nodes and connections colored based on oscillatory components in the ACHs and CCHs, respectively, as in Figure 5B.

2.1. Uniting functional network topology and oscillations

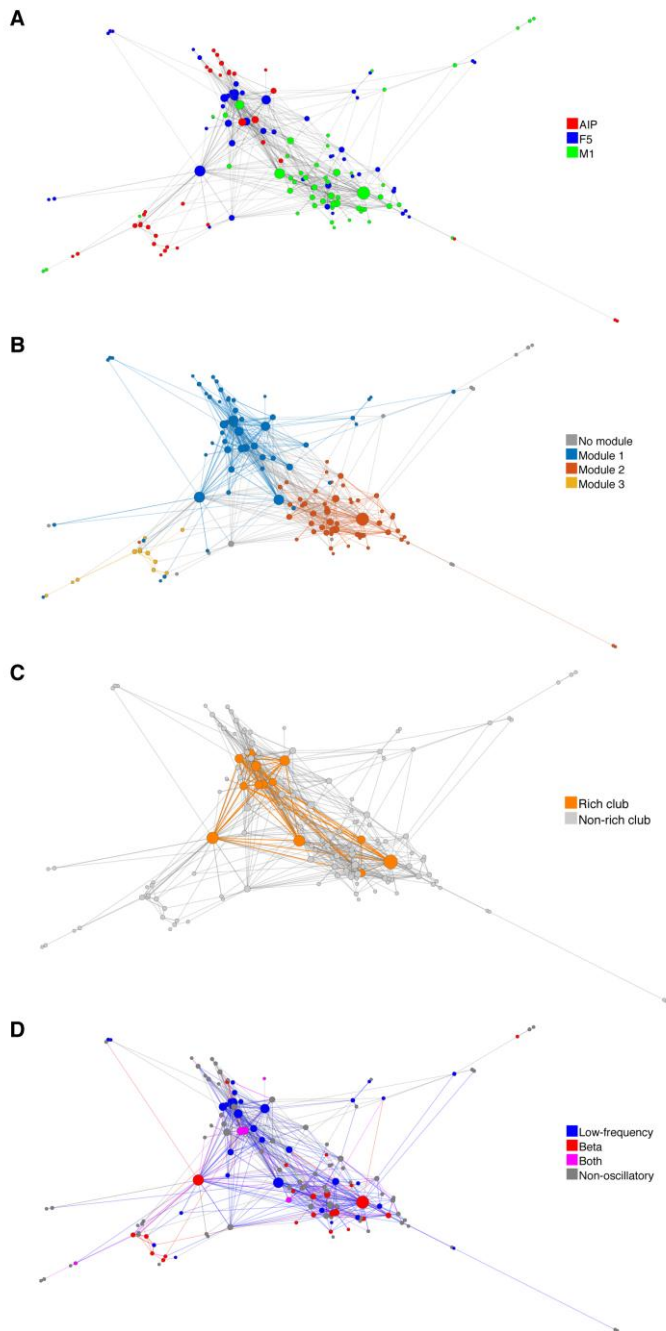


Figure 3 – Figure Supplement 2. Functional network connectivity of an exemplar data set displayed as a web where the locations of all neurons were determined using the visualization of similarities (VOS) approach (Van Eck and Waltman, 2007). (A) Each node is colored based on the area it was recorded. (B) Each node colored based on its module. (C) Nodes and connections colored based on oscillatory components in the ACHs and CCHs, respectively. (D) Nodes and connections colored based on rich-clubness. Each circle represents a single neuron and is scaled based on the degree of connectivity. VOS aims to find locations in a low-dimensional space (in this case 2D) in such a way that the distance between each node reflects the similarity between these nodes. Similarity is typically found by calculating the association strength (also known as proximity index) on the co-occurrence matrix of items, which is in this case the weighted network connectivity matrix. Association strength is simply the co-occurrence of two items divided by the product of the number of occurrences of each item. The location of each node is then found by minimizing the sum of the squared distance between all nodes, weighted by the computed similarity between each node. To avoid trivial solutions in which all nodes are assigned the same location, there is an additional constraint that the average distance between all pairs of items must be equal to one. Mathematically, VOS bares much similarity to the method of multi-dimensional scaling (Van Eck et al., 2010). All implementations of VOS were performed using the freely available software, Pajek (<http://mrvar.fdv.uni-lj.si/pajek/>), and then plotted in Matlab.

2.1. Uniting functional network topology and oscillations

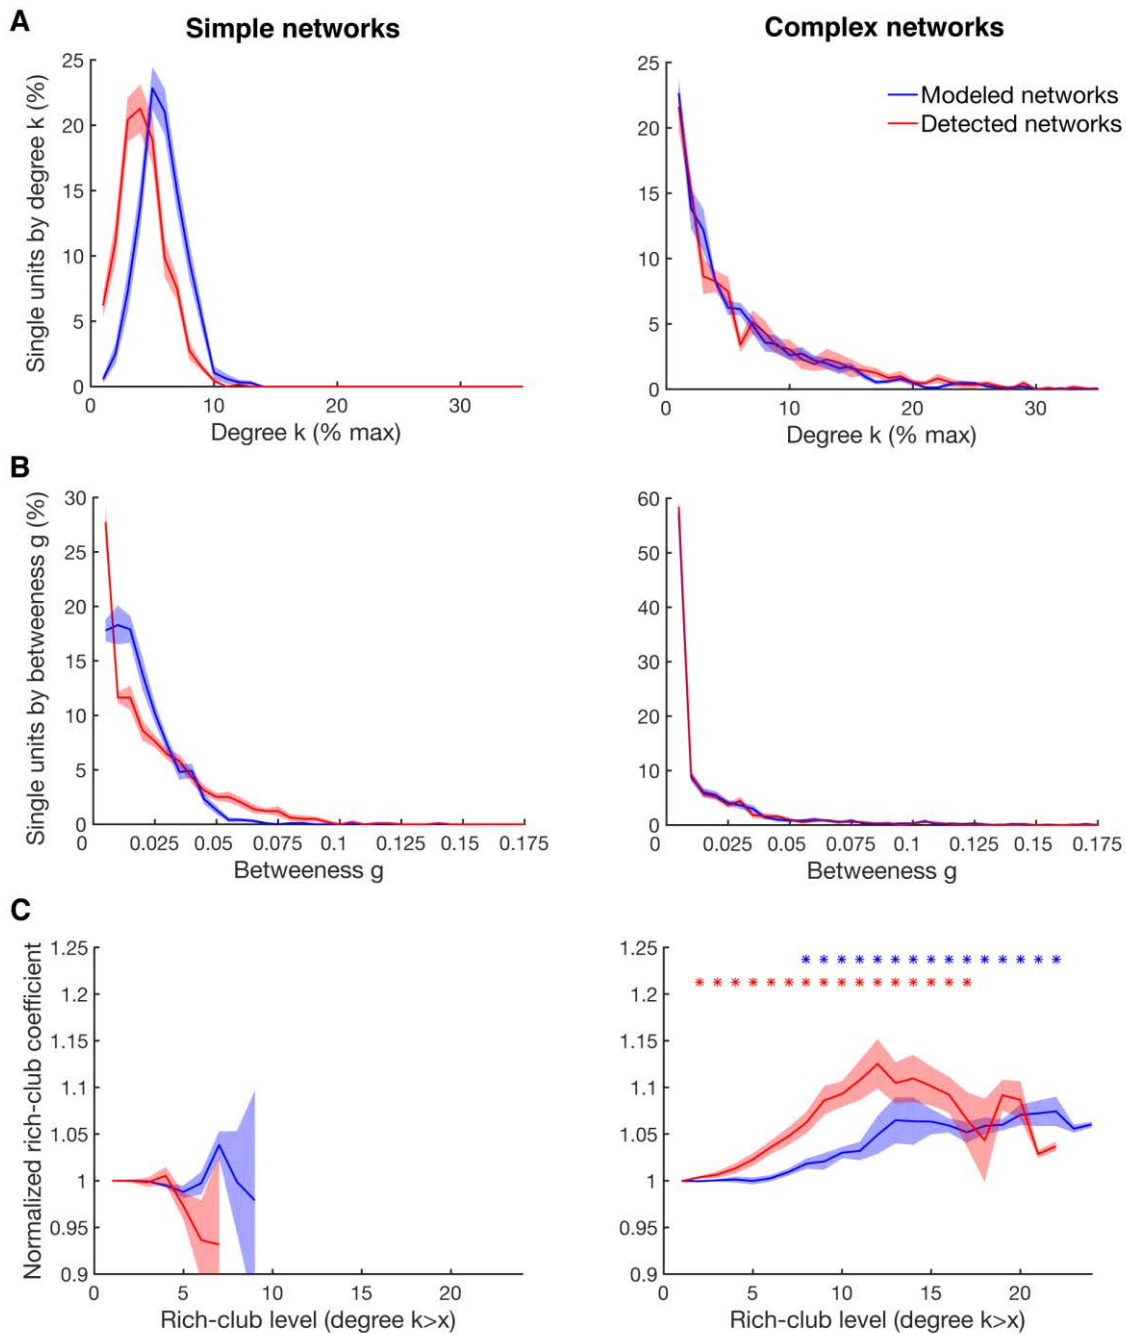


Figure 4 – Figure Supplement 1. Detectability of the underlying network topology using equal rate model simulations. (A) Average degree centrality distribution of all networks simulated with the equal rate model (blue) and the corresponding detected networks with the described method for detecting directed functional connectivity (red). Results are shown for the same 10 simulated simple networks and 10 simulated complex networks as in Figure 2 – Figure Supplement 3. Error bars show the standard error across simulated networks. (B) Same as in A, but for the betweenness centrality distributions. (C) Same as in A, but for the rich-club level relative to surrogate datasets. Asterisks represent significant difference of rich-club level to surrogate networks. Two different sets of surrogate networks were calculated per dataset, one for the simulated network and one for the detected network.

2.1. Uniting functional network topology and oscillations

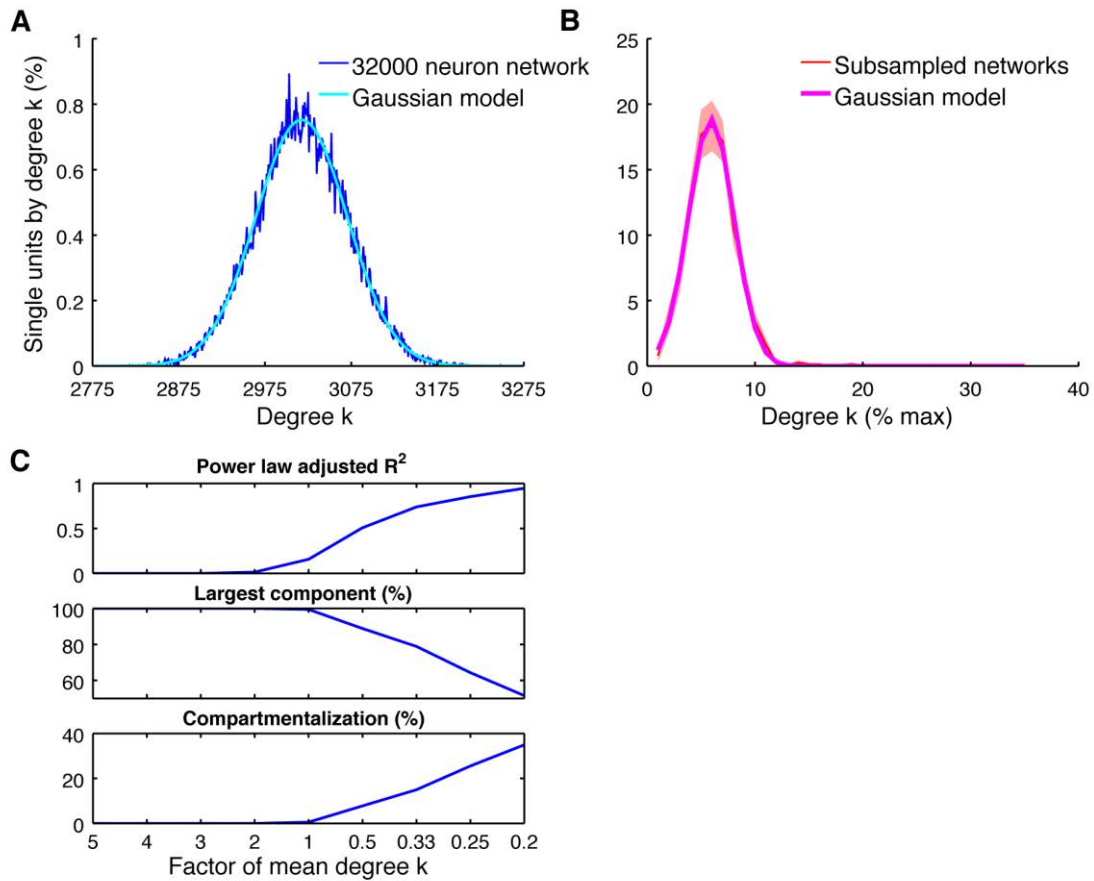


Figure 4 – Figure Supplement 2. Subsampling model. (A) Average degree centrality distribution of the modeled neuronal plane (32000 neurons, 2 areas, each divided into 5 subregions coverable by an array, 160 possible electrode position, and a maximum of 20 single units per electrode) with distant dependent random connectivity (Figure 3B). The distribution could be best described by a Gaussian model (adjusted $R^2 = 0.98$). (B) Average degree centrality distribution of 12 different subsamplings of the modeled neuronal plane with exactly the same number of neurons as in the real datasets. Line shadings show standard error across subsamplings. Datasets were processed as in Figure 4A. Average degree distribution could be best described by a Gaussian model (adjusted $R^2 = 1$) and only poorly by a power law model (adjusted $R^2 = 0.17$). (C) Dependency of goodness of power law fit, the size of the largest component relative to the whole network, and the level of compartmentalization on average degree k. Different average degrees were generated by varying the distance-dependent connectivity density of the empirically gained data (Figure 3B) by factors of 1/5, 1/4, 1/3, 1/2, 1, 2, 3, 4, and 5 times to create a neuronal plane. Goodness of power law fit was highly correlated with the size of the largest component (adjusted $R^2 = 0.93$) and the compartmentalization (adjusted $R^2 = 0.93$).

2.1. Uniting functional network topology and oscillations

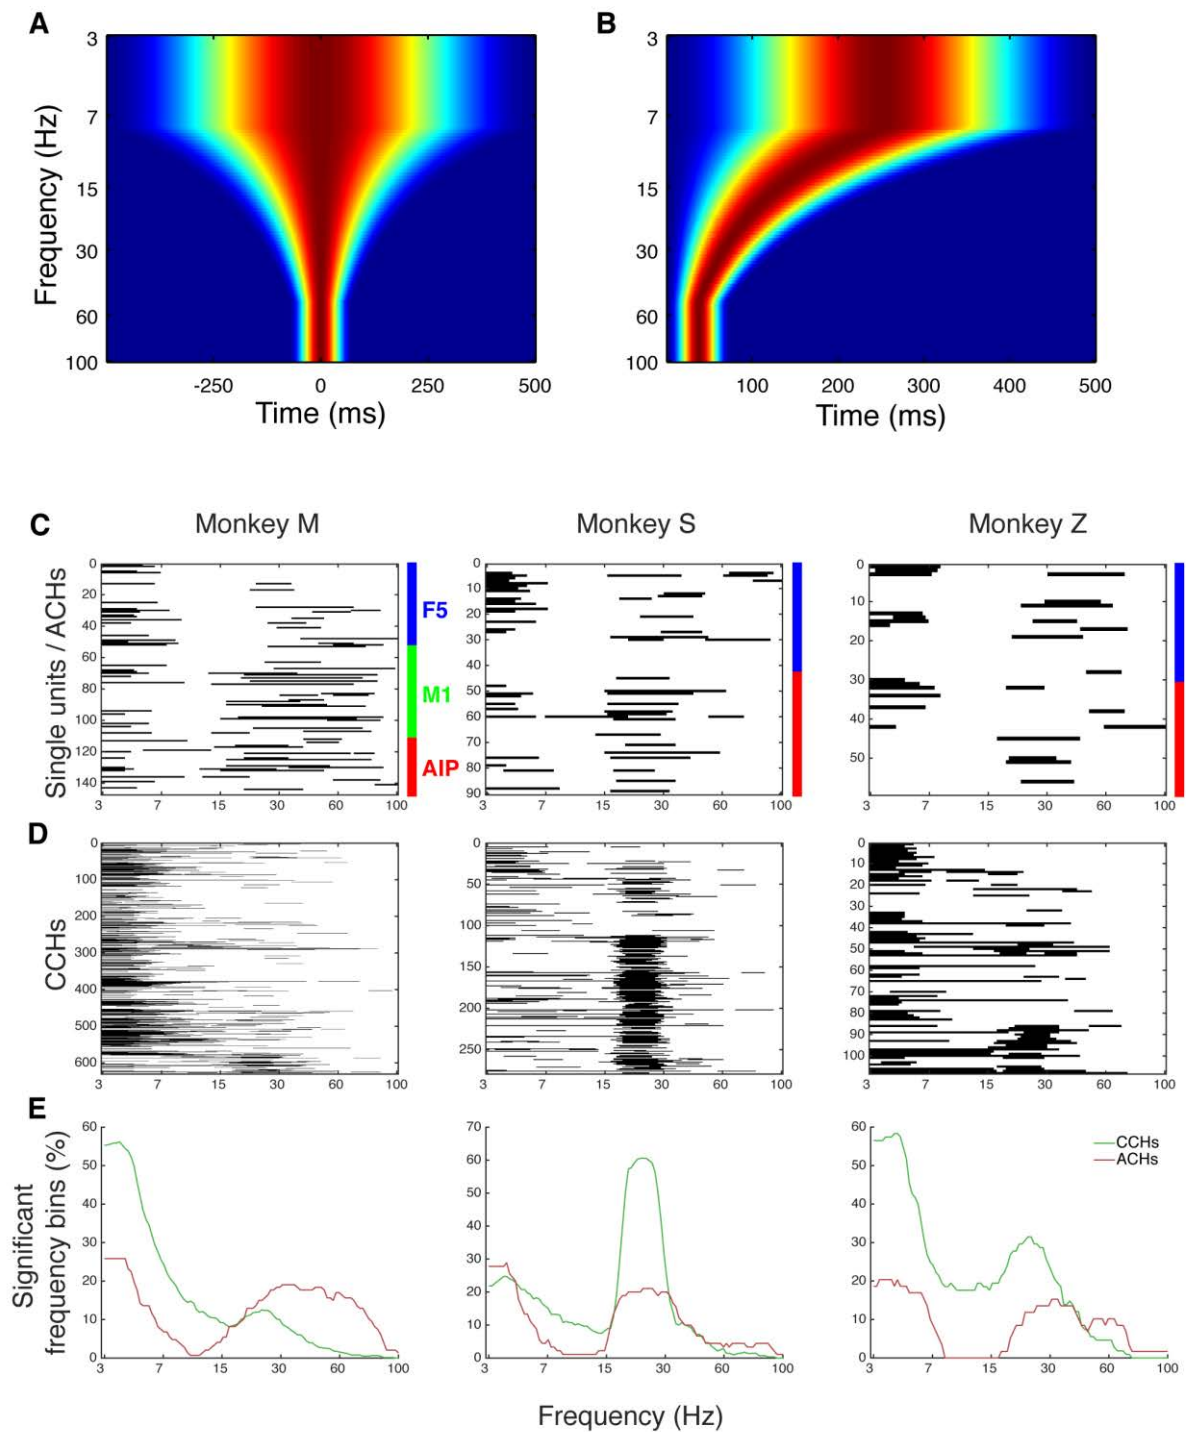


Figure 5 – Figure Supplement 1. Frequency dependent Hanning windows used for discrete Fourier transform. (A) Hanning windows used for discrete Fourier transform of all CCHs. All windows were aligned to the zero bin and span four times the frequency of interest period (with a maximum of 1000ms and a minimum of 150ms). Frequencies of interest were scaled logarithmically (100 frequencies from 3 to 100 Hz). (B) Hanning windows used for discrete Fourier transform of all ACHs. All windows were aligned to the zero bin and span two times the frequency of interest period (with a maximum of 500ms and a minimum of 75ms). (C) Significant frequency bins of power spectra of all ACHs of one example dataset per monkey. Frequencies were calculated and displayed on a logarithmic scale. (D) Significant frequency bins of power spectra of all CCHs of the same example datasets as in C. (E) Average number of significant frequency bins of all ACHs and CCHs of the same example datasets as in C and D.

2.1. Uniting functional network topology and oscillations

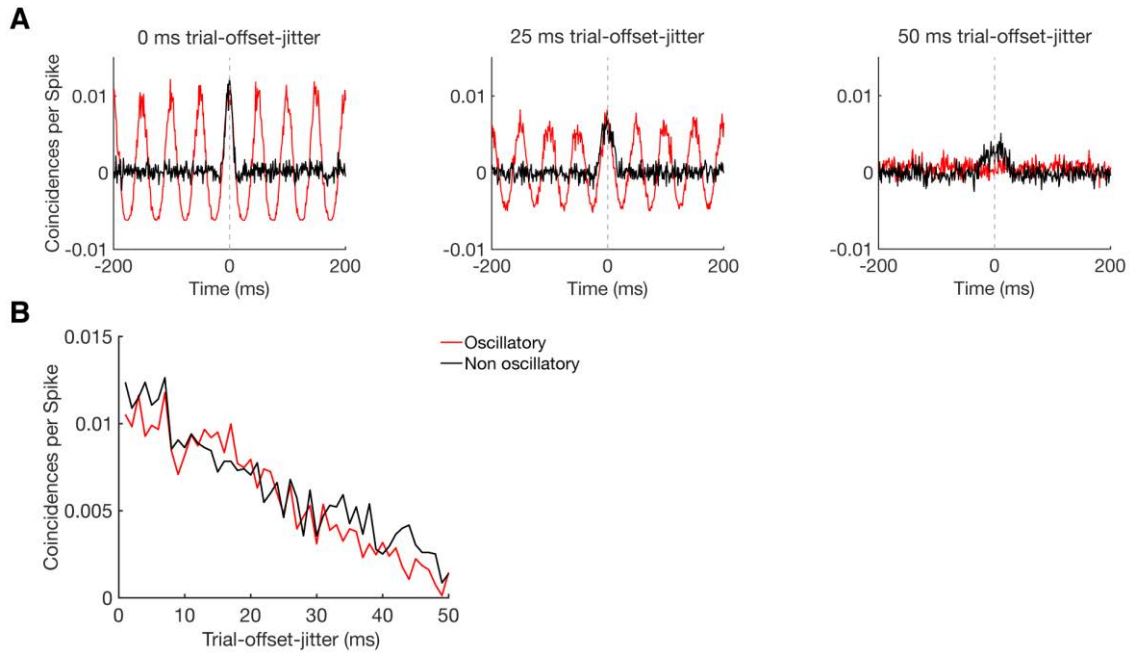


Figure 5 – Figure Supplement 2. Sensitivity of CCHs in detecting oscillatory synchrony and non-oscillatory synchrony. (A) CCHs for pairs of simulated neurons with an average firing rate around 5Hz, either firing in an oscillatory (20Hz, red curve) or non-oscillatory manner (black curve). By jittering their trial-wise temporal offset in firing, we simulated different levels of coupling strength, without disturbing the firing pattern of the individual neurons nor the similarity in firing between the two neurons. Results are shown for a trial-wise jitter of 0ms (perfect synchronization), 25ms, and 50ms (hardly synchronized). (B) Maximum CCH peak heights of oscillatory and non-oscillatory neurons with a systematical trial-offset-jitter from 0 to 50ms.

2.1. Uniting functional network topology and oscillations

Datasets	Trials	Single units total	F5	M1	AIP	Single units used	F5	M1	AIP
M 1	958	149	48	57	44	148	48	57	43
M 2 *	900	147	52	58	37	137	50	52	35
M 3	621	107	49	32	26	79	41	20	18
S 1	503	86	46	-	40	57	28	-	29
S 2	565	76	39	-	37	64	30	-	34
S 3	460	76	35	-	41	64	28	-	36
S 4	460	82	35	-	47	64	26	-	38
S 5 *	557	90	42	-	48	78	37	-	41
S 6	374	83	42	-	41	47	25	-	22
Z 1	400	52	29	-	23	33	21	-	12
Z 2	436	48	24	-	24	30	17	-	13
Z 3 *	608	59	30	-	29	41	21	-	20
Average	570.2	87.9	39.3	49	36.4	70.2	31	43	28.4
SD	177.4	31.2	8.5	12.0	8.5	35.8	10.3	16.4	10.5

Table 1. Trial and single unit counts for all datasets. Marked datasets correspond to the displayed example networks in Figure 3-5 and Figure 3 – Figure Supplements 1 and 2. Columns 3-6 show the total and area specific number of units recorded. Columns 7-10 show total and area specific number of units of the largest component of the network, which is the basis for all topological analysis.

2.1. Uniting functional network topology and oscillations

Datasets	Oscillators total	Non-Oscillators	Beta Oscillators	Low Frequency oscillators	Oscillators in both frequency ranges
M 1	83	65	37	60	14
M 2 *	60	77	28	37	5
M 3	34	45	12	25	3
S 1	31	26	14	26	9
S 2	32	32	14	22	4
S 3	31	33	15	20	4
S 4	26	38	14	19	7
S 5 *	40	38	22	25	7
S 6	21	26	14	10	3
Z 1	13	20	5	10	2
Z 2	13	17	6	9	2
Z 3 *	18	23	10	11	3
Average	33.5	36.7	15.9	22.8	5.3
SD	19.4	17.4	8.7	13.8	3.4

Table 2. Number of oscillators in all networks analyzed. Marked datasets correspond to the displayed example networks in Figure 5 and Figure 3 – Figure Supplements 1 and 2.

Three information subspaces explain the category-free population dynamics in the macaque fronto-parietal network

Benjamin Dann^{1,*}, Jonathan A Michaels^{1,2,3,*}, Andres Agudelo-Toro¹, Hansjörg Scherberger^{1,4}

¹German Primate Center, Kellnerweg 4, 37077 Göttingen, Germany

²Electrical Engineering Department, Stanford University, Stanford, CA 94305, USA

³Howard Hughes Medical Institute, Stanford University, Stanford, CA 94305, USA

⁴Faculty of Biology and Psychology, University of Göttingen, 37073 Göttingen, Germany

*Equal contribution

Acknowledgements: We thank R. Lbik, N. Bobb, and S. Borchert for animal support, and M. Dörge for technical support. This work was supported by the Deutsche Forschungs Gemeinschaft (SCHE 1575/1-1 & 3-1 and SFB889, C9), the Bundesministerium für Bildung und Forschung (BCCN-II, 01GQ1005C) and the European Union (FP7-611687, NEBIAS).

2.2. Three information subspaces explain the category-free population dynamics

Abstract

The fronto-parietal network is known to be essential for sensory to movement transformations, including decision-related processes. However, how information is encoded in this neural circuit is still unclear. Here, we recorded many neurons in parallel in the fronto-parietal grasping network of two macaque monkeys who were visually instructed, or freely choose, to grasp a handle in one of two different ways. Despite high numbers of tuned neurons, individual neuron response patterns were often complex, multiphasic, and heterogeneous across the neuronal population. In contrast, when we considered neural population activity as one dynamical process, all task-specific single trial activity could be explained by an evolution through subspaces representing visual, preparatory, and movement information, into which all neurons contributed uniformly. A recurrent neuronal network model with a decision-making process and generating muscle patterns reproduced the recorded neuronal dynamics. These results suggest that sensorimotor transformation can be well explained as a dynamic transformation between information-subspaces according to the behavioral demands.

Introduction

Deciding to act on the environment involves the flexible preparation and execution of movements based on external and internal drives. To this end, different kinds of sensory information and internal motivations have to be transformed into the same movement plans¹⁻⁴. Yet, the internal and external circumstances can change which requires the system to dynamically adapt to the new situation potentially resulting in a changed movement plan⁵. Furthermore, once the moment has come to move the movement plan needs to be transformed into movement related activity^{4,6-8}. The fronto-parietal network has been identified to be strongly involved in the flexible transformation of visual information into movement plans and in turn into movement related activity, with its neurons being modulated for visual features^{1,9-11}, movement preparation^{6,12-14}, movement execution^{6,7,15}, and decision making processes^{9,14,16-18}. However, despite the high number of conducted studies it is still unclear how information is encoded and transformed in the network.

Based on the representational framework the firing rate of each neuron is described as tuning to various parameters^{19,20}. The transformations taking place in the fronto-parietal network have been described as interactions between different categories of neurons, specifically visual, visuomotor, and motor neurons¹¹. An increase in firing rate of the motor-related neurons passing a threshold was thought to cause movement initiation²¹. However, individual neuron tuning analyses only explain a fraction of the rich heterogeneity of neural population response, and the assumed tuning functions only roughly match the individual neurons²²⁻²⁴, leaving a lot of neural variance unexplained.

In contrast, from the dynamical system perspective neural circuit function is assumed to arise from the activation of the whole network of neurons, which cannot be understood by studying one neuron at a time²⁵. In this perspective, neural population activity evolves through a lower-dimensional space where the current state causes the next state²⁶. This framework explains crucial aspects of the preparatory activity and the transition to movement activity in premotor and motor cortex^{6,7,27-29}, as well as decision process in prefrontal cortex (PFC)^{3,30}. Furthermore, neural contributions to low-dimensional subspaces during preparatory activity have been found to be randomly distributed across neurons, eschewing specific neuron types or categories^{4,31}. Interestingly, there is increasing evidence that population activity transitions through a limited number of independent subspaces, which was found for PFC activity during a perceptual decision task³, and was described for

2.2. Three information subspaces explain the category-free population dynamics

the transition from preparatory to movement activity in premotor and motor cortex⁸. However, finding the underlying dimensionality of the population activity is difficult since time-shifts between neurons^{23,32}, as well as uncontrolled moment-by-moment fluctuations, as present during decision processes^{5,33}, can result in an artificial high number of estimated dimensions. For this reason, a proper estimate of the dimensionality explaining all task-specific single trial variance during the encoding and transformation from sensory to movement activity is still missing.

Here, we analyzed how information is encoded and transformed in the fronto-parietal grasping network, while two monkeys performed different combinations of visually-instructed or freely-chosen delayed grasping movements. Tuning analyses revealed all neurons to be significantly modulated by task parameters. However, single-neuron response patterns were often complex, multiphasic, strongly changing over time, and heterogeneously across the neuronal population with no sign of categories. Exploratory population analyses revealed a clear temporal and conditional structure with periods where only visual, preparatory, or movement information was present. We found that the population response for the three types of information explored orthogonal subspaces explaining nearly all task specific single trial variance, with neurons across areas contributing randomly, without evidence of categories, to the three information subspaces. Neurons from a regularized recurrent neuronal network (RNN) trained to generate muscle patterns for completing the same task strikingly resembled the recorded activity on the single-neuron and population level. These findings suggest that the encoding and transformation of information in the fronto-parietal network can be well understood as a dynamical evolution through subspaces allowing for an independent moment-by-moment readout of parallel-encoded information.

Results

Task and Behavior

Two monkeys (S and Z) were trained to perform three variants of a delayed grasping task where they had to grasp a handle with either a power and precision grip (Fig. 1a) as described previously^{29,34}. In the instructed-task monkeys were visual cued by one of two different discs displayed on a monitor to perform the associated grip type. In the free-

2.2. Three information subspaces explain the category-free population dynamics

choice-task both discs were displayed, but monkeys were ensured to switch grip types (mean power choice: $38.5 \pm 4.2\%$ and $53.3 \pm 5.2\%$ for S and Z respectively) by iteratively reducing the reward every time the monkey repeatedly chose the same grip type. In the delayed-instructed-task first both discs were displayed, but after a variable period of time one of the two discs were displayed again giving a clear instruction identical to the instructed-task. Equal number of trials of the three tasks were presented in in random order to the monkeys. For all following analyses 6 datasets from monkey S and 3 from monkey Z were used. Both monkeys learned to perform the task with high accuracy and high trial counts (mean successful trials: 95% and 96%, successful trials: 730 ± 106 and 722 ± 167 for S and Z respectively) and stable reaction and movement times across task types (Supplementary Fig. 1).

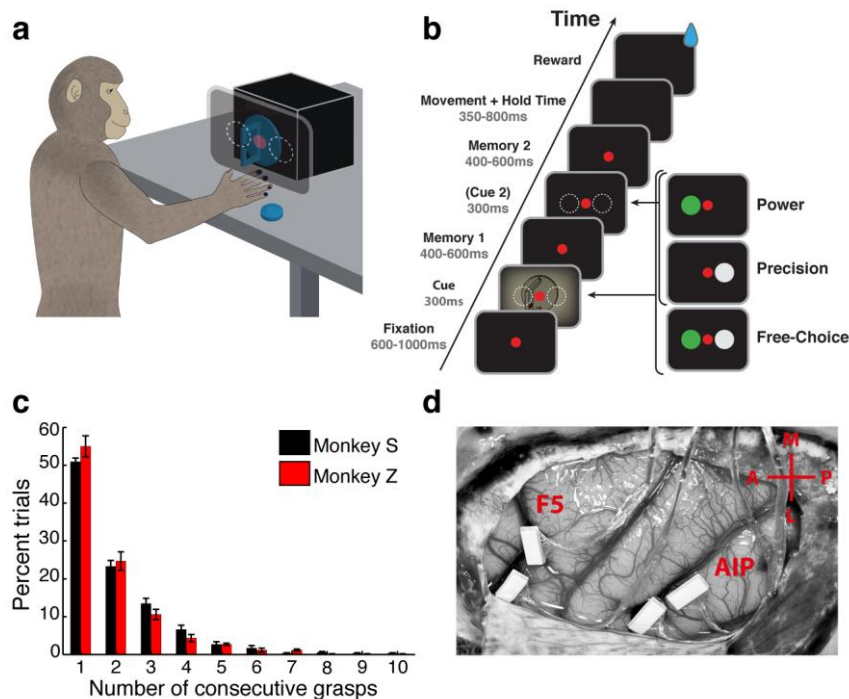


Figure 1 (a) Setup: Monkeys were visually cued to grasp a target (handle) with one of two different grip types displayed on a monitor appearing superimposed on the handle. Task: Monkeys had to fixate a red disk for 600-1000ms (Fixation), followed by a cue period of 300ms (Cue). Then, either a clear instruction was given by a disk to the left or right of the fixation disc indicating a power or precision grip respectively, or a free-choice cue was given by turning on both disks. After the cue, a memory period followed (duration: 1100-1500ms) before the fixation dot was turned off (go-signal) indicating the monkey to execute the grasp movement

(maximum duration:1000ms). In 50% of all cases where a free-choice cue was presented an instruction cue was presented for 300ms after 400-600ms in the middle of the memory period. **(b)** Distribution of number of consecutive trials for the free-choice-task. Note, that trials of the free-choice-task were randomly interleaved by trials from the other two tasks. The distribution was estimated separate per dataset and displayed is the average per monkey with standard error across datasets. **(c)** Electrode array implantation of monkey S with 4 floating microelectrode arrays in areas AIP and F5. Arrays were implanted at the lateral end of the Intraparietal Sulcus in AIP, and in the posterior bank of the Arcuate Sulcus in area F5.

Different types of decision making have been described with different implications for the underlying neuronal mechanism^{14,18,35,36}, which makes it essential to analyze the choice behavior of the free-choice-task. Intuitively one might expect the monkeys to switch their performed grasp type every trial (rule-based-decision), since this would be the ideal

2.2. Three information subspaces explain the category-free population dynamics

solution to maximize the reward. However, both monkeys showed a graded decay in number of consecutive grasps (Fig. 1b), which still leaves two other possible types of choice behavior: reward-ratio-based-decisions or random-decisions. Both decision types would result in a graded, predictable decay of consecutive grasps with a faster decay for reward ratio based decisions (see Online Methods). The switching distribution for all datasets of both monkeys was highly correlated with both decision type probability distributions (reward ratio based decision: $r = 0.970 \pm 0.023$, random decision: $r = 0.989 \pm 0.009$, which was expected due to their high similarity. Nevertheless, the random-decision distribution was significantly higher correlated (Wilcoxon signed rank test, $p = 0.027$), indicating random choice behavior.

Neuronal recordings and single unit tuning

We recorded from the grasping motor network, including part of the ventral premotor (F5) and anterior intraparietal (AIP). In each area, recordings were obtained from two floating microelectrode arrays (FMAs), for a total of 64 channels (32 per microarray) per area (Figure 1c). For all analyses only well isolated single units with an average rate above 1Hz were used, resulting in an average number of single units per dataset of 31.0 ± 3.2 and 21.6 ± 3.2 for area F5 and 32.3 ± 3.3 and 16.0 ± 1.0 for area AIP for monkey S and Z respectively. Note that no other preselection criterion for units was applied.

100% of single units of both areas and all datasets showed significant modulation from baseline (cluster-based surrogate t-test, $p < 0.05$, see Online Methods) and $98 \pm 3\%$ of F5 units, and $89 \pm 7\%$ of AIP units were significantly tuned for either grip-type, task-type or grip-task-interaction (Cluster-based permutation 2-way anova, $p < 0.05$, see Online Methods) with no significant difference between datasets of monkey S and Z (Wilcoxon ranksum test, $p = 1$ for F5 and $p = 0.9$ for AIP). The high number of task modulated as well as tuned neurons clearly proved that our arrays were implanted at grasp and task relevant spatial location of both areas.

Based on the representational framework one might expect either visual units tuned to visual parameters, motor units tuned for movement related parameters, or visuomotor units tuning in a meaningful way to both^{1,10} and indeed some units in both areas showed the expected tuning characteristic (Fig. 2a). However, a large proportion of units was not classifiably in this way and showed complex, multiphasic tuning patterns (Fig. 2b) as previously described for reach movements^{6,7} in premotor and motor cortex as well as

2.2. Three information subspaces explain the category-free population dynamics

prefrontal and parietal cortex^{3,4}. Nevertheless, when we pooled all significant time points across units we found a high increasing number of grip type tuned units starting from cue onset with an maximum during movement (Fig. 2c), well in line with previous literature^{12,13}. Task-type and grip-task interaction tuning was highest shortly after cue onset and dropped nearly to zero at the end of the memory epoch. The number of significant tuned neurons across time was highly similar for AIP and F5 ($r = 0.90$, $r = 0.80$ for significant tuning, tuning amplitude, respectively) with slightly less grip-type tuned neurons in AIP. In line with grip-type tuning the number of significant task dependent modulated units showed a similar temporal sequence for all conditions and was also similar for both areas (Fig. 2d; $r = 0.90$, $r = 0.78$, significant modulation, and modulation amplitude, respectively). Note that all following results were pooled across datasets from both monkeys due to their high degree of temporal similarity (AIP: $r = 0.94$, $r = 0.89$, $r = 0.93$, $r = 0.88$; F5: $r = 0.96$, $r = 0.94$, $r = 0.96$, $r = 0.87$ for significant tuning, significant modulation, tuning amplitude, and modulation amplitude, respectively). The relative high and stable number of significant tuned units could be easily misinterpreted as a clear sign for tuning stability well in line with the representational framework and the increasing number of tuning units with a maximum during movement even as sign for a movement threshold³⁵, however, tuning preference was not stable over time for many units (Fig. 2a).

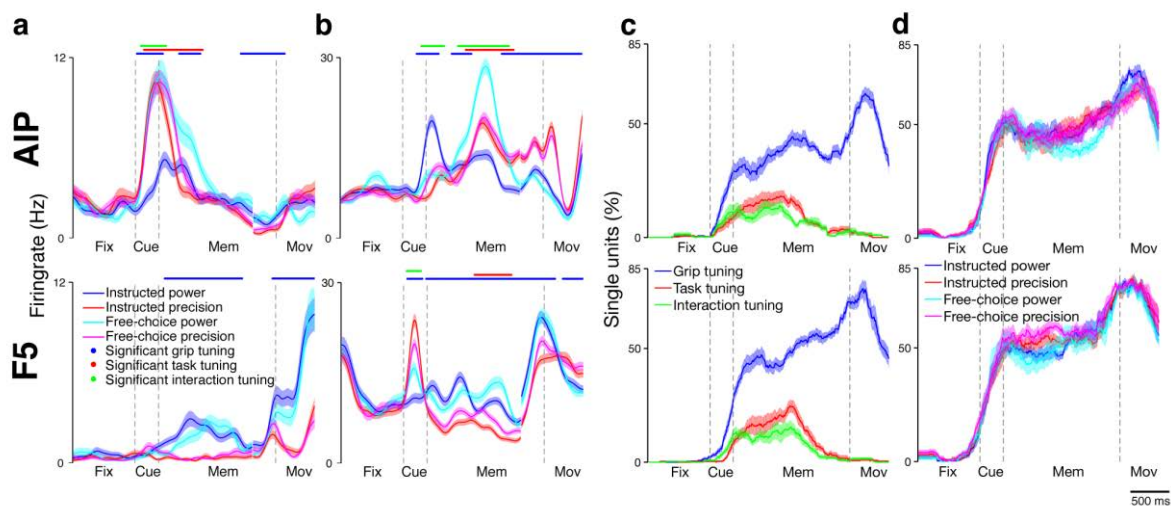


Figure 2 Average firing rates per condition and significant temporal tuning characteristics of the population of recorded single units in AIP and F5. (a) Average firing rates across time and conditions of an example putative AIP visual unit tuned for the precision cue and an example putative F5 motor unit only modulated for grip-type differences. Line shadings represent standard error across trials. The lines on top represent significant tuning for grip-, task-, and grip-task-interaction. **(b)** Average firing rates across time and conditions of an AIP and a F5 example complex tuned units. Standard error and significances are displayed as in **a**. **(c)** Average number of significant tuned units for AIP and F5 units across time. Line shadings represent standard error across datasets. Note that the set of tuned units across time does not have to be the same. **(d)** Same as in **c**, but for significant modulation from baseline (fixation period).

2.2. Three information subspaces explain the category-free population dynamics

Tuning stability, rate distributions and single unit tuning similarity

In order to quantify tuning preference stability over time for the entire neuronal population we created tuning similarity matrices starting from cue onset (Fig. 3a, b). Each value within each matrix represent the tuning similarity between the two corresponding time points where a value of 1 reflects stable significant tuning preference over time, a value of 0 independent tuning preference and a value of -1 an inversion of tuning preference. Across all datasets and for both areas tuning preference was highly dynamic over time for grip-type tuning (Fig. 3a) as well as significant modulation from baseline (Fig. 3b). Between epochs or between farer apart time points tuning similarity was even close to 0, indicating independent tuning of the neuronal population. Thus, despite of a stable number of tuned units with a maximum during movement these findings are not in line with a representational view where all epochs of the task should share similar tuning to at least a certain extend. Yet, under the dynamical systems view there is little reason why tuning should be similar for the initial and subsequent states of the system⁶.

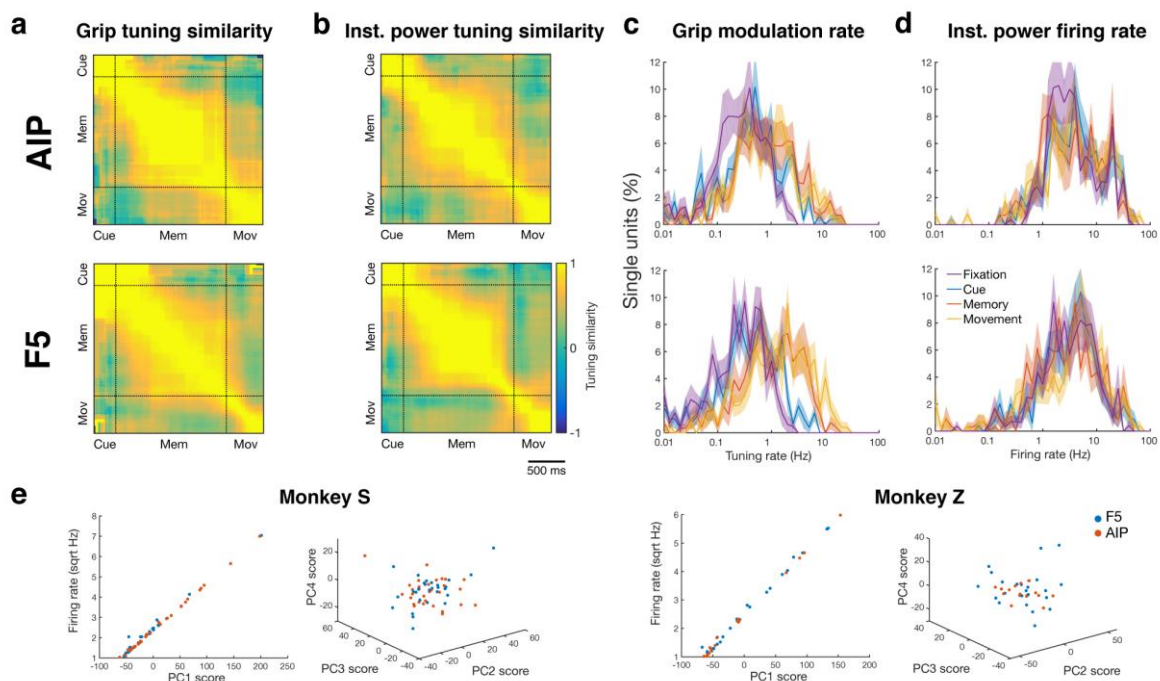


Figure 3 Tuning characteristics across the neuronal population and time. (a) Tuning stability of significant grip tuning across time separate for units from AIP and F5. Each value within each matrix represent the tuning similarity between the two corresponding time points where a value of 1 reflects stable significant tuning preference over time, a value of 0 independent tuning preference and a value of -1 an inversion of tuning preference. (b) Same as in a, but for significant modulation from baseline. Only tuning similarity for the instructed-power condition is shown as an example, since tuning similarity was highly similar for instructed-precision, free-choice-power and free-choice-precision condition. (c) Tuning rate distribution of all four task periods separate for AIP and F5. Cue onset -300ms, +300ms, +1100ms and movement onset +150 were selected for fixation, cue, memory, and movement period respectively. The distribution from each dataset was

2.2. Three information subspaces explain the category-free population dynamics

first normalized by the number of single units and then averaged across datasets. Line shadings represent standard error across datasets. **(d)** Same as in c, but for raw firing rates of the instructed-power condition. The raw distributions of the other instructed and the two free-choice conditions were highly similar. **(e)** Projection of all units from one example dataset per monkey onto the first 4 PCs. PCA was applied on the trial averaged responses of all single units recorded in parallel per dataset with conditions x time as variables and units as observations. The square root of average firing rates was taken before calculating PCA that is also why the first PC is plotted against the square root of the average firing rate.

Another important property assumed from a representational framework is that first the distribution of firing rate difference between different condition should be multi- or at least bimodal distributed, since not all units are assumed to participate in the task, while from a dynamical system perspective modulation rate can be even randomly distributed across the entire population resulting an unimodal distribution of rate differences⁴.

Displayed in Fig. 3c is the tuning rate distribution at four randomly picked time points during the four epochs of the task for both areas. Interestingly, tuning rate was unimodal distributed at any time point for all tuning types (Cluster-based surrogate Hartigans Dip test, $p < 0.05$, see Online Methods). Also, the firing rate was unimodal distributed for all time points of all conditions and for both areas (see for example Fig. 3d; Cluster-based surrogate Hartigans Dip test, $p < 0.05$). Tuning rate as well as firing rate distributions turned out to be best displayed on a log-scale which is well in line with previous findings showing that neuronal firing rates are log normal distributed across cortex^{34,37}. Note that the Hartigans Dip test is a nonparametric test for multimodality, which makes the finding of unimodal rate distribution independent from the choice of scale. The unimodal distribution of tuning and firing rate critically questions the usage of tuned unit counts for AIP and F5, since in this case the number of tuned units is simply a function of recorded trials and spikes, which results in a misleading interpretation of information representation in the recorded areas.

However, despite the lag of tuning stability and side lobes in the rate distributions across time, it is still possible that the temporal modulation in firing rate is similar for distinct categorical groups of single units. Dimensionality reduction using principal component analyses (PCA) as commonly used for spike-sorting is a suitable approach to detect any categorical structure in the neuronal space if present were each resulting dimension represents a linear combination of all time points across conditions. To exclude a possible bias of the results by different firing rate per area we tested for differences per dataset with no significant effect (Wilcoxon ranksum test, $p < 0.05$). Interestingly, the first PC already explained $74 \pm 7\%$ of all variance and nearly perfectly resembled the average firing rate per neuron (Fig. 3e, Supplementary Fig. 2). In the space of PC 2-4 explaining together with PC 1

2.2. Three information subspaces explain the category-free population dynamics

91±2% of all variance no signature of any neuronal clustering became apparent neither within or between areas for all datasets (Fig. 3e, Supplementary Fig. 2) rather showing a heterogeneous distribution of neuronal response patterns.

In summary, the lack of tuning stability across time, the unimodal rate distributions at all time points, and the heterogeneity of single unit activity with no sign of suggests that the activity of the neuronal population of area AIP and F5 can not be properly described by a representational framework.

Structure within the neuronal state-space

By contrast, the presented findings are in line with the dynamical system perspective^{4,7,22} where network function is assumed to arise from the neuronal population activity evolving through different states by which the current state causes the next state^{25,26}. One common way to analyze the neuronal population activity is the state-space framework, in which the firing rate of each neuron is a dimension and the firing rates over time form a trajectory through this space^{4,6,7,27}.

In order to explore the evolution of condition wise neuronal trajectories we estimated the Euclidean distance between all time points of trial-averaged activity within and across conditions. The distance structure of the different datasets was similar (F5: $r = 0.88 \pm 0.07$, AIP $r = 0.71 \pm 0.14$) and pooled for all further analyses. Displayed in Fig. 4a are the condition wise distance matrices where each value represents the distance between the two corresponding time points. Remarkably, the temporal structure of the condition wise neuronal trajectories was highly similar between conditions (AIP: 0.98 ± 0.01 , F5: 0.98 ± 0.01) and even across areas ($r = 0.92$). The trajectories appeared to resemble the task structure with an expected temporal delay, which we tested for by applying an assumption free cluster algorithm (see Online Methods) on the full distance matrix including all time points of all conditions. Intriguingly, population activity within the four task periods were clustered across conditions (shown in light gray) in a similar way for both areas.

To analyze condition dependent differences on top of the found temporal structure we compared the distance between conditional trajectories at the same time point of the task (Fig. 4b). For the instructed power and instructed precision conditions a significant split (Cluster-based permutation ttest, $p < 0.05$, see Online Methods) was present shortly after cue onset which remained stable through the task and peaked during movement. The

2.2. Three information subspaces explain the category-free population dynamics

distance between free power and free precision condition was quite similar to the instructed trajectories. However, the trajectories split significant later in AIP compared to the instructed task and showed a slower rise in distance in both areas, suggesting that the task specific information in this early period is mainly visual, since this is only present in the instructed task. Interestingly, there was no significant split between free and instructed trajectories shortly before movement onset. The lack of task specific information at this time point where the monkeys still remained silence and no go cue was given yet, let us assume the neuronal activity only represented movement preparation free of any visual or task rule specific activity. Also, later on during the movement there was no task specific information present however a clear increase in distance for the two grip types during the movement was present.

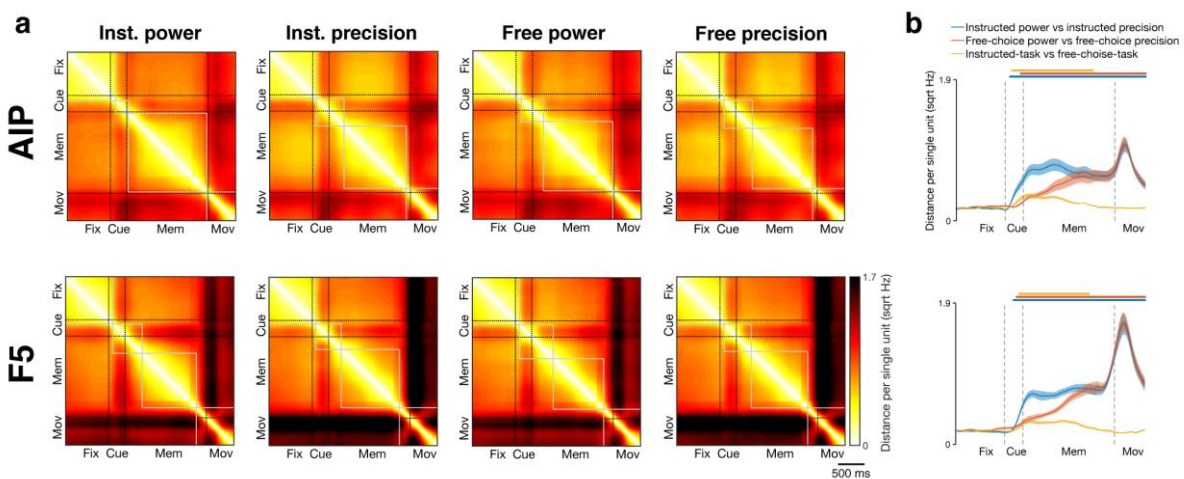


Figure 4 Euclidian distance structure of the full neuronal space of trial averaged population activity per area. (a) Condition wise distance matrices where each value represents the Euclidean distance between the two corresponding time points across the task. Distance matrices of the different datasets were collapsed due to their high degree of similarity. Low values correspond to the population activity at the two time points to be similar and high values represent correspond to the population activity two to be far apart from each other. The square root of average firing rates was taken before calculating the Euclidian distance and in order to compare days and areas the Euclidian distance was normalized to distance per single unit (see Online Methods). The gray boxes show the four identified clusters across conditions per area. (b) Dataset averaged Euclidian distance of the trial averaged population activity at the same time point of the task. Since only the same time points are compared global modulations across conditions are not visible. For the comparison of the instructed- and free-choice population response unit wise activity for the two grip types was collapsed. The lines on top represent significant distance for the corresponding distance comparison with the same colour. Line shadings represent standard error across datasets.

Crucially, the three specific time periods for visual, preparatory and movement information were always within the found clusters for cue, memory and movement period and never between the clusters (Fig. 4a). One possible explanation for this distinct structure could be that the whole neuronal activity transitions through a low dimensional state-space

2.2. Three information subspaces explain the category-free population dynamics

with subspaces for visual, preparatory and movement information. The high degree of similarity between the two areas even suggests that this is happening across the areas.

Single trial neuronal trajectories

For a proper evaluation of the temporal and conditional structure population activity structure moment-by-moment readouts of the population state on single-trials is required^{5,33}. This is especially important out of two reasons, first since neuronal activity can be a function of internal cognitive processes as it is in particular the case for decision-making process, and secondly since single-trial activity reflects the ground truth of the ongoing processes in the neuronal population. Dimensionality reduction methods were shown to be well-suited for revealing low dimensional representations within high dimensional data with the additional advantage to be applicable on single trial data^{5,29,33}.

We used linear discriminant analysis (LDA) to project the single-trial neuronal population activity of both areas into a one-dimensional space best separating the task specific variance separate for each dataset (see Online Methods), which has been proven to be a robust method for single trial activity^{4,5,33}. This was done in steps of 10ms through the time course to the task to always capture the optimal separating projections, which together corresponds to the whole task specific neuronal variance. For this purpose, we just used successful instructed task trials, since we know a priori that the population response is either related to the power or precision condition. The degree of task specific information quantified by the degree of overlap of the population response for the two conditions was high at all time points and for all datasets after cue onset ($94.4 \pm 5.8\%$) confirming the relevance of the neuronal population for the executed task.

To assess if the neuronal population evolves through the same or different subspaces over time and task-types we projected the activity of all other time points of all task-types into all optimal instructed-task spaces. In case of a stable representation across time and task-types the degree of task specific information should be high at all time points for all projections across task-types, while in case of a dynamic population response transitioning through different subspaces the task specific information of different projection should be time or task specific and up to independent for different projections. Task specific information turned out to be dynamic, with some stability during the memory period and an increasing degree of similarity between the conditions over time (Fig. 5a). The optimal

2.2. Three information subspaces explain the category-free population dynamics

projections during the cue period were highly specific to this time period and not yielding any task specific information at any other time point of the instructed-task nor for any time point of the free-choice-task. However, when the same cues as for the instructed-task were shown at a later time point for the delayed-instructed-task selective information was present in exactly that period, indicating these projections to be specific for visual information. Task specific information was more sustained during the memory period for the instructed-task, with an increasing degree in information for the free-choice-task and once the determining cue was given also for the delayed-instructed-task.

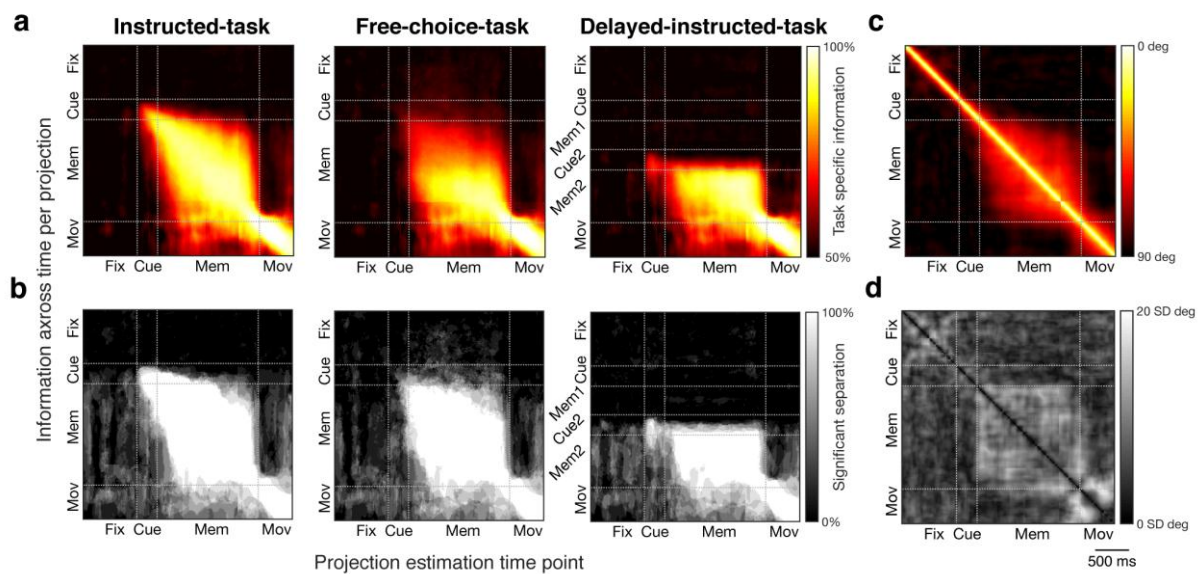


Figure 5 Optimal separating projections of the single trial population response for power and precision condition per time point of the instructed-task. (a) The degree of task specific information for all time points for all optimal instructed-task projections. In case of the instructed-task the estimation of the optimal projection was leave-one-out cross-validation to prevent overfitting (see Online Methods). Projections were estimated based on the square root of smoothed single trial firing rates. The degree of task specific information of the different datasets was collapsed due to their high degree of similarity. **(b)** The percentage of datasets with significant task specific information above chance level per time point, task-type and projection (see Online Methods). Each data point of the three matrices corresponded one to one to the matrices in **a**. **(c)** Average angle in degree between all optimal single trial population projections of the instructed task across datasets. 0 degree correspond to a complete overlap of the projections, while 90 degrees correspond to complete orthogonality. **(d)** Standard deviation across datasets of the angle between all optimal projections in **c**.

This was to be expected since every specific information present before the second cue (Fig. 1a) would simply cancel out. Interestingly, the degree of task specific information for the same projections was task-type unspecific at later time points (Cue onset +1970 - +2500ms, Movement onset -300 - +500ms, Cluster-based permutation 1-way anova, $p < 0.05$, see Online Methods). The projections estimated during late memory period were not giving any specific information during the cue and later movement period, suggesting them

2.2. Three information subspaces explain the category-free population dynamics

to represent preparatory information. In contrast, projections from the movement period were highly selective just representing information during the movement period. As an additional control, we also estimated optimal separating projections based on the free-choice-task. The degree of task specific information for the late memory and movement period was indistinguishable from the instructed-task results (Cluster-based permutation 1-way anova, $p < 0.05$) confirming these periods to be task-type unspecific. The described pattern of task specific information was similar and significant above chance level for all datasets (Fig. 5b; permutation test, $p < 0.01$, see Online Methods).

Intriguingly, when we checked the projections from the three identified specific periods for their independence they turned out to be nearly perfectly orthogonal to each other for all datasets (Fig. 5c-d). Taken together, these results suggest that the population response of AIP and F5 evolves through three independent information-subspaces for visual, preparatory and movement activity.

In order to evaluate the temporal evaluation of the single trial trajectories through the three information-subspaces and to compare the different task-types we took one projection from each orthogonal period (Fig. 6) in accordance with the 3 specific periods from the Euclidean distance analyses (Fig. 4) to ensure independence and selectivity based on both analyses. The task specific information for the three subspaces clearly overlapped before returning to baseline level (Fig. 6a), which is considered an important requirement for information transfer between them^{3,8}. Note that preparatory- and movement-subspace were task-type unspecific as described earlier is now clearly visible by the nearly perfect overlap of task-specific-information at later time points. Displayed in Fig. 6b are 150 randomly selected single trial trajectories per task-type from one dataset of monkey Z (see Supplementary Fig. 3 as example for monkey S). The specific separation for the three subspaces is clearly visible as well as the overlapped between the different subspaces. Note that the trajectories in the preparatory-subspace for the delayed-instructed-task and the free-choice-task showed the same heterogeneity up to cue2 as previously expected. In the same period for both tasks preparatory-trajectories showed a lot of vacillation in agreement with previous findings⁵. The forced switch of movement plans for the delayed-instructed-task in case cue2 was not consistent with the pre-existing movement plan was nicely visible by a lot of crossings shortly after onset of cue2.

2.2. Three information subspaces explain the category-free population dynamics

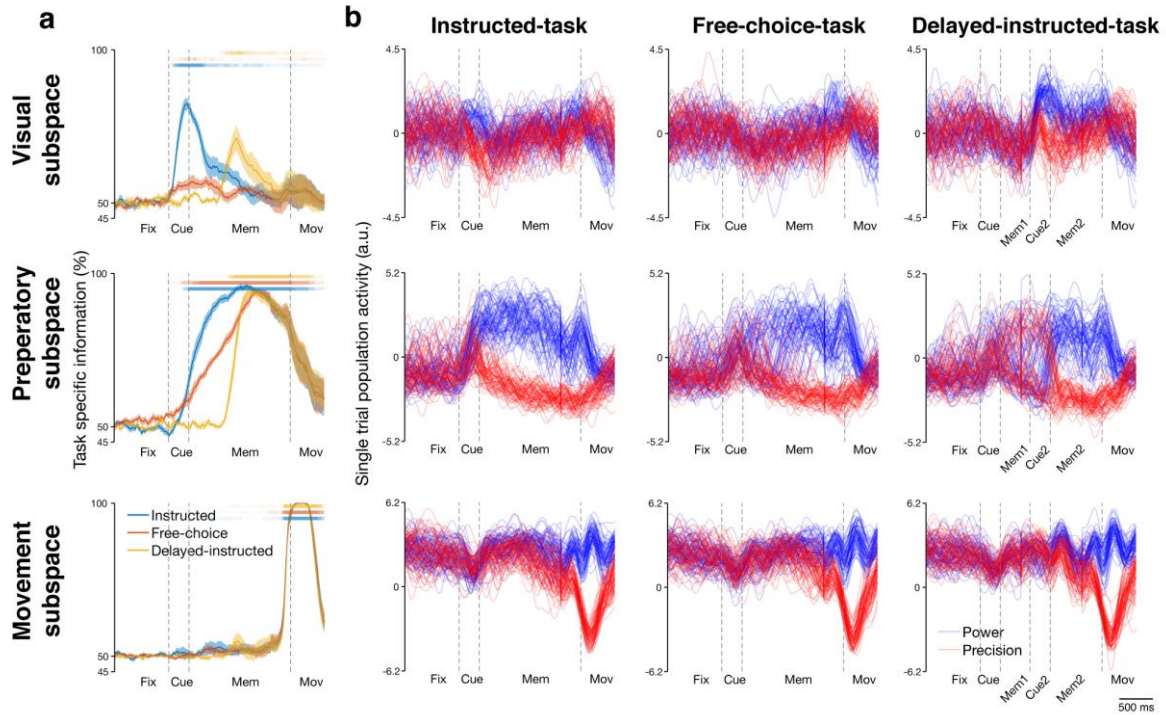


Figure 6 Projections of the population response into the three subspaces for visual, preparatory, and movement information. (a) The degree of task specific information for the three selected subspace projections for all time points and task-types (see Online Methods). The lines on top represent the percentage of datasets with significant task specific information above chance level per time point and task-type as in Fig. 5b, but just for the subspace projections. The different colours correspond to the different task-types. **(b)** Single trial trajectories for the projections into the three subspaces of one example dataset from monkey Z. 150 trials (75 trials per grip-type) were randomly chosen per task-type. Not all single trial trajectories were displayed for a better overview. The range per subspace projection was fixed for better comparison between the different task-types.

Given the confirmed specificity and independence of the visual-, preparatory- and movement-subspaces one key question is whether these three subspaces are enough to explain the whole task specific variance over time within the AIP-F5 single unit network or do we miss an essential part of it. A possible alternative could be a rapid series of state transitions which are not even orthogonal to each other but yet contain specific information not captured by the three orthogonal spaces as it was previously suggested for prefrontal cortex populations³⁰. To this end, we estimated two sets of time point wise optimal projections for each task-type separately to be sure to catch all task-type specific variance if present. One set was estimated from the whole neuronal population as described before, while the other one was estimated just based on the three-dimensional trajectories from the visual-, preparatory- and movement-subspaces, which are as already mentioned just based on the same three projections for all task-types. The degree of task-specific information over time as well as the single trial trajectories of the two sets of optimal projection nearly

2.2. Three information subspaces explain the category-free population dynamics

perfectly matched (Fig. 7a,b) and were highly correlated (degree of task-specific information: $r = 0.97 \pm 0.02$, average trajectories per condition: $r = 0.98 \pm 0.01$, and single trial trajectories: $r = 0.86 \pm 0.04$ with standard deviation over datasets). As a consequence of this analyses, nearly all task-specific variance of the whole neuronal population across time, task-types, and areas can be explained by just three orthogonal subspaces. Especially, the simplicity of just three subspace capturing nearly all task-specific variance on the network level speaks in favour of the dynamical system perspective, which would be much more complex if not impossible using a representational framework.

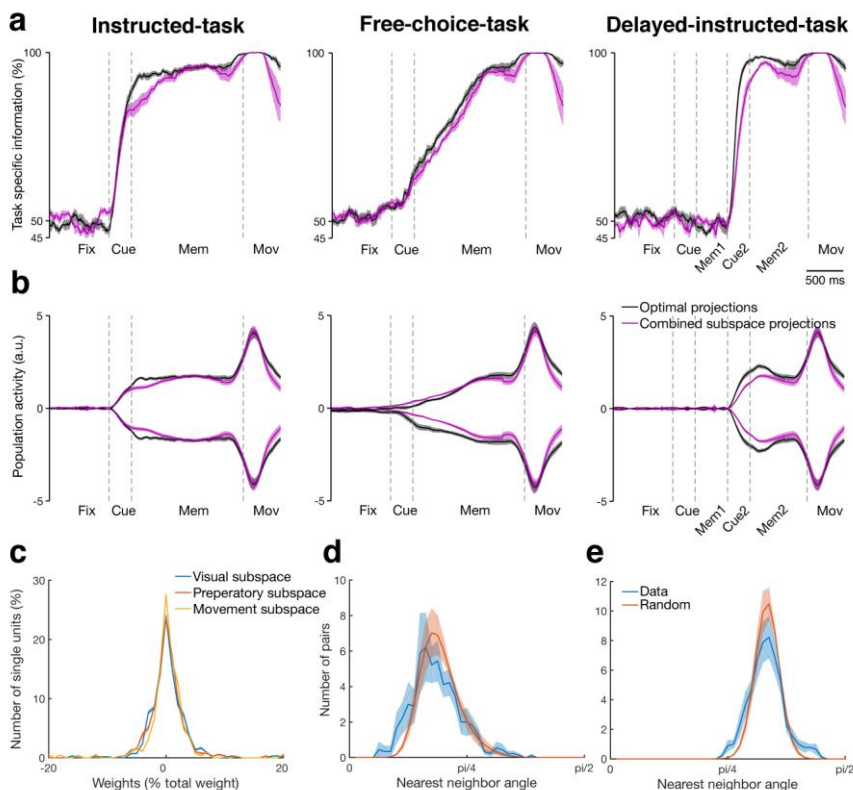


Figure 7 Amount of task-type specific variance explained by the three subspaces and the distribution of neuronal contributions to the subspaces. **(a)** The maximum degree of information estimated by the optimal projection at each time point of each task-type (3 x 310 projections) in comparison to the combined degree of information from the three subspace projections (see Online Methods). Black lines represent the results from the optimal projections estimated from the whole neuronal population and purple lines results just based on the three-dimensional subspace trajectories. The

average across all datasets is displayed with line shadings representing standard error across datasets. **(b)** Same as in **a**, but for condition wise average single trial trajectories. **(c)** Average neuronal weight distributions for the three subspace projections across datasets. For better comparison, the neuronal weights of each subspace projection were first normalized by the total absolute weight per projection. Line shadings represent standard error across datasets. **(d)** Average histogram of the distribution of angles between each single unit and its nearest neighbours across datasets. The blue line represents the distribution of nearest-neighbour angles of the neuronal weights for the three subspace projections and the red line represents the distribution of nearest-neighbour angles of random three-dimensional vectors. Line shadings represent standard error across datasets. **(e)** Same as in **d**, but for the first 15 PCs of trial averaged data.

The weights of the three subspaces allowed for more assumption free estimation of area wise encoding of information, since they were independent of the length of time a certain information is present in the network, which in turn strongly depends on the task-design. Visual information was stronger encoded in AIP, preparatory information was equally present in both areas and movement information was stronger in F5 (Supplementary Fig. 4).

2.2. Three information subspaces explain the category-free population dynamics

However, the neuronal population of both areas contributed to all three subspaces and none was presented in just one area, which is well in line with current findings^{8,9}. Both areas had the strongest encoding of movement information, followed by preparatory information, with the weakest representation for visual information, suggesting both areas to be stronger involved in movement related processing.

Although we haven't found any signature of neuronal categories yet (Fig. 3, Supplementary Fig. 2) it is possible that the projections for the individual subspaces are just based on distinct groups of units, which could not be detectable without unmixing their response. In this case, the weight distributions reflecting the individual contributions of all units for the three subspaces should be bi- or multimodal, which was not the case for any of the datasets or subspaces (Fig. 7c; Hartigans Dip test, $p < 0.05$). However, this still leaves the possibility that certain combinations of contributions across subspaces appear more often than by chance, which would be evidence for categories across subspaces. To test for this, we compared the weight vector of each unit for the three subspaces with its nearest neighbors. If there were any categories across as well as within subspaces there should be pairs of units which should have closer neighbors than expected by chance, which can be tested for by a new developed statistic called "PAIRS"⁴ (see Online Methods). The PAIRS test did not indicate any evidence of categories across the three subspaces and datasets (Fig. 7d; $p = 0.45$). Also, when we used PCA as an unsupervised method for dimensionality reduction (see Online Methods), there was no sign of categories (Fig. 7e; $p = 0.60$). The finding of randomly distributed neuronal contributions for the three subspaces were, which in turn capture nearly all task-specific variance across both areas gives even more evidence for this network to be a dynamical system.

Comparison of a recurrent neuronal network model to data

A crucial indication for the network of neurons to be a dynamical system is basically generating a model where we know that it is a dynamical system by definition. To be a valid model it should get the same inputs, create the same outputs and resembling the neuronal responses as close as possible. This has been recently done for the transition of preparatory activity to movement activity in the PMd-M1 reaching network by optimizing a regularized recurrent neuronal network (RNN)³⁸ to produce muscle patterns. The dynamics of the RNN resembled the dynamics seen in the recorded neuronal population at both the single-neuron

2.2. Three information subspaces explain the category-free population dynamics

and population levels. However, it is unclear how well such a model can resemble fronto-parietal network activity for different task-types as in our case visual instructions as well as for a random free-choice task. We trained a similar network of 200 neurons to generate muscle patterns for the whole reach to grasp movement estimated from a musco-skeletal model based on recorded kinematics for the two grasps (Fig. 8a; see Online Methods). The network received randomly weighted independent inputs for the two cues resembling our visual stimuli with an identical temporal profile to our trained task (Fig. 1a), which means in case if a free-choice cue both inputs were given at the same time. Additionally, noise was added to the network. All time periods resembled the real task, which means it had to memorize the instruction or generate a movement from noise for the free-choice-task, as well being able to flip movement plans if necessary for the delayed-instructed-task. Importantly, we trained the network not just to generate the required muscle patterns but also to encourage extremely simple solutions by heavily regularizing it³⁸ (see Online Methods).

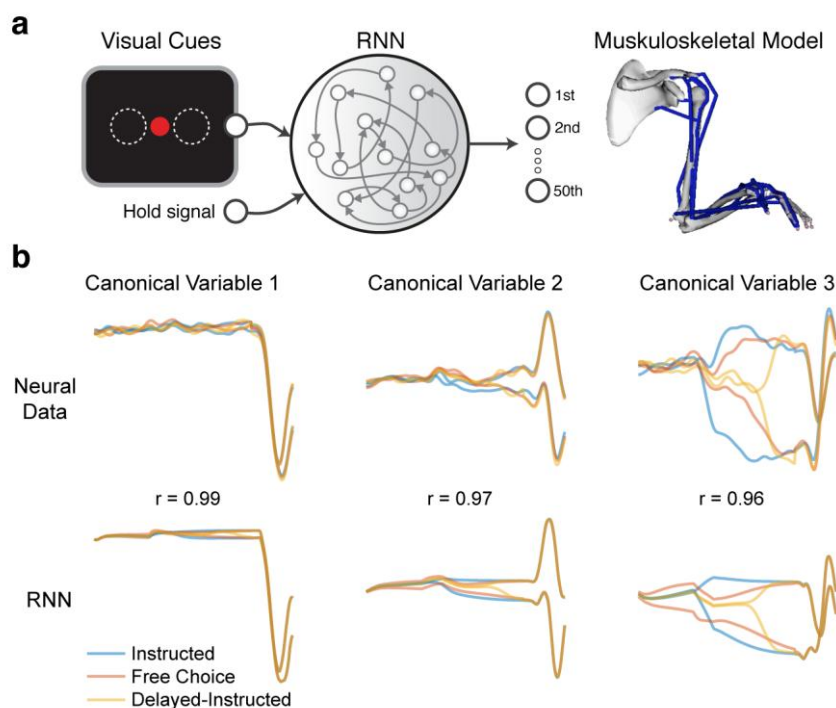


Figure 8 Recurrent neural network (RNN) for decision making. (a) The RNN containing 300 internal neurons receives four input signals. Three of the input signals represent the visual signal on individual trials, i.e. when the precision cue, power cue, and fixation cue were presented. The fourth input represents a hold signal that is released 200 ms before the onset of movement. The output of

the network is a 50-dimensional signal representing the condition-averaged velocity of each muscle during either power or precision grip. (b) The trial-averaged signal of both the recorded neural activity and the simulated activity of the RNN were compared using Principal Component Analysis followed by Canonical Correlation Analysis (CCA). CCA finds low-dimensional projections of the data that are most highly correlated and where each dimension is orthogonal.

We were able to successfully train the network to generate the two grasp

2.2. Three information subspaces explain the category-free population dynamics

movements for all three task-types. Intriguingly, the population activity as well as the single-unit responses well matched the dynamics of the recorded data (Fig. 8b) with a high average canonical correlation for the first 8 dimensions ($r = 0.73$ for monkey Z, and $r = 0.67$ for monkey S). To our knowledge this is the first model including random decisions which also resembles the population and single-unit activity. Even if the results from this model are not a causal proof they strongly suggest that the encoding and transformations of information within the fronto-parietal grasping network can be well understood from a dynamical system perspective.

Discussion

In this study, we addressed the question how information is presented in fronto-parietal grasping network for different tasks and grip types. By analyzing the population of neurons from the classical representational perspective describing the activity of neurons as a function of tuning for various parameters^{19,20} we found high numbers of tuned neurons in both areas. However, crucial aspects of neuronal activity could not be explained by this classical perspective. The tuning characteristic of many neurons were complex and the tuning of the population of single units changed dynamically and were even independent over time especially between different periods of the task. Also, neuronal responses were heterogeneously distributed with no evidence of distinct categorical groups of neurons tuned for any task parameters.

In contrast, when we considered the population of neurons as one interconnected ensemble^{25,33} with its activity as a whole evolving through state-space in the framework of dynamical systems^{3-7,26,30,38-41} a clear temporal and conditional structure became apparent. Based on this data driven exploratory approach³³, we could identify three orthogonal subspaces for visual, preparatory and movement information explaining nearly all task specific single trial variance across time, tasks-, grip-types, and even areas in form of a dynamical trajectory through this informational subspaces. Contributions to all subspaces were randomly distributed across all neurons of both areas tested with a new highly sensitive test PAIRS⁴, which showed an involvement of all neurons in the encoding and transformation of all three information types.

Intriguingly, a regularized recurrent neuronal network³⁸ trained to produce muscle

2.2. Three information subspaces explain the category-free population dynamics

activity for the two grasps, which is by definition a dynamical system, could well reproduce the neuronal dynamics on the single unit and population level.

It is important to state that the representational view and the dynamical system view are not completely at odds^{6,26}. Since the contributions to the three informational spaces are randomly distributed across the neuronal population (Fig. 7b) individual neurons could be indeed described to be mixed selective³¹ for visual, preparatory and movement selective^{11,42} to a certain extent. However this would always result in subselection on neurons not explaining an essential part of the neuronal variance^{22,23} and how information is transformed up to the generation of muscle movements^{26,38}.

The finding of the whole fronto-parietal population response evolving through three independent subspaces raises the question, what is the encoding advantage of such an encoding and transformation structure? An interesting concept is the idea that activity of different orthogonal subspaces cancel out at the level the population readout, which was shown for the communication of the motor cortex to the spinal cord and arm muscles to avoid causing involuntary movements^{40,43}. From this point of view different subspaces would allow for multiplexed information encoding in the same network, which could be flexibly transformed on the bases of the behavioral demands^{3,4,31}. This observed activity in the preparatory subspace flexibly driven by visual activity or presumably by noise in the system as well as the behavioral dependent interaction of visual and preparatory activity for the delayed-instructed-task (Fig. 6b, Supplementary Fig. 3) supports this idea.

Interestingly, as assumed by many studies the encoding and transformation of information was not happening between areas^{10,18,40} but between subspaces spanning the two areas. This findings are in line with recent findings for preparatory and movement subspaces spanning premotor and motor cortex⁸ as well as a study showing that sensory as well as preparatory information were encoded across the entire visuomotor pathway, albeit with different strength⁹. A possible explanation is given by the distributed consensus model, in which decisions occur on multiple levels in parallel, and the final decision is achieved through a distributed consensus³⁶. While the model itself is rather abstract about the exact encoding and transformation of information, a neuronal population activity evolving through subspaces for different information spanning different areas could be solution to it.

A rather global than local encoding and transformation of information raises the question how these processes are coordinated. For limited network with a limited number

2.2. Three information subspaces explain the category-free population dynamics

of conditions a dynamical system, modeled by a trained regularized RNNs receiving the same inputs and generating muscle patterns, offer a remarkably solution how information is coordinated³⁸. In particular, the high degree of similarity of population activity as well as the single-unit responses with the recorded data, makes this solution very convincing. However, the coordination of information across many brain areas for a nearly endless number of different behaviors requires presumably a more complex coordination mechanism. Oscillatory synchrony has been proposed as a key mechanism for global coordination of information^{44,45}. Especially synchronization in the beta-band was shown to be involved in behavioral relevant coordination of information in the fronto-parietal network^{16,46}. Intriguingly, we found that the same neuronal population analyzed in this study was coordinated by an area spanning strongly interconnected group of oscillatory synchronized neurons in the beta- and low-frequency range³⁴, which let us assume that oscillatory synchrony could play a crucial role in the coordination of cross area information encoding and the transition of information between different subspaces.

Despite many aspects of the fronto-parietal encoding and transformation of information are still to be understood, the simplicity that nearly all single trial neuronal variance of all neurons across areas can be understood as a dynamical process through just three information subspaces offers a new perspective to analyze activity from this network. The independence of the subspaces if visual, preparatory, and movement information even allows to disentangle and analyze them separately, which could be useful for analytical studies as well as possible decoding approaches.

2.2. Three information subspaces explain the category-free population dynamics

Materials and Methods

Basic procedures

Neural activity was recorded simultaneously from many channels in two female rhesus macaque monkey (Animals S and Z; body weight 9 and 7kg, respectively). Detailed experimental procedures have been described previously^{29,34}. All procedures and animal care were 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*⁴⁷.

Behavioral Task

Figure 1A illustrates the time course of the behavioral task as described previously^{29,34}. Trials started after the monkey placed both hands on the resting positions and fixated a red fixation disk (fixation period). After 600 to 1000ms, cues in the form of disks were shown next to the fixation disk for 300ms to instruct the monkey about the required grip type (power or precision; cue period). During this epoch the grasp target, a handle, was also illuminated. In the instructed-task (33% of all cases) one disk was shown, while in the free-choice-task and the delayed-instructed-task (33% of all cases each) both disks were turned on, indicating that the monkey was free to choose between the two grip types. The monkey then had to memorize the cued condition for 1100 to 1500ms (memory period). In the delayed instructed-task a second cue identical to the instructed-task was given in the middle of the memory period after 400 to 600ms keeping the total memory period the same as for the other tasks. The switching off of the fixation light cued the monkey to reach and grasp the target (movement period) in order to receive a liquid reward. Importantly, during free choice trials the reward was iteratively reduced every time the monkey repeatedly chose the same grip type. All trials were randomly interleaved and executed in darkness.

Chronic electrode implantation

Surgical procedures have been described previously^{29,34}. In short, each animal was implanted with two floating microelectrode arrays per area (FMAs; Microprobes for Life Sciences; 32 electrodes; spacing between electrodes: 400 μ m; length: 1.5 to 7.1 mm monotonically increasing to target grey matter along the sulcus). Animal S and Z were implanted with four FMAs in area AIP and F5 in the left and the right hemisphere, respectively. Animal M was

2.2. Three information subspaces explain the category-free population dynamics

implanted with a total of six FMAs in the same cortical areas and two additional arrays in area M1, in the left hemisphere (Figure 1B).

Choice behavior

In order to evaluate which choice behavior was present we compared the distribution of consecutive grasps (Fig. 1B) to the probability distribution of consecutive grasps of rule-based-decision, reward-ratio-based-decisions or random-decisions separate for each dataset. The probability distribution of rule based decision is simply 1 for switching after one grasp and 0 for all other values. Since the histogram of consecutive grasps showed a graded decay ruled based decision behavior could be excluded right away. The probability distribution of reward-ratio-based-decision is given by:

$$P(n) = \frac{1}{(2^n - 1) + 1} \quad (1)$$

In contrast, the probability distribution of random-based-decision is given by:

$$P(n) = \frac{1}{2^n} \quad (2)$$

where n of the number of consecutive grasps. The distribution of consecutive grasps was correlated with both probability distributions per dataset giving a direct estimate how good the choice behavior can be described by both probability distributions.

Neural recordings and spike sorting

Neural 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 band-pass hardware filter; for monkey S and Z) or a 256 channel Tucker-Davis system (TDT RZ2; sampling rate 24.414 kS/s; 0.6-10000Hz band-pass hardware filter; monkey M).

For spike detection, data were first low-pass filtered with a median filter (window length 3ms) and the result subtracted from the raw signal, corresponding to a nonlinear high-pass filter. The signal was then low-pass filtered with a non-causal Butterworth filter (5000 Hz; 4th order). To eliminate common noise-sources, principal component (PC) artifact cancellation was applied for all electrodes of each array, as described previously⁴⁸. To ensure

2.2. Three information subspaces explain the category-free population dynamics

that no individual channels were eliminated, PCs with any normalized coefficient greater than 0.36 (conservatively chosen) were retained. Spike waveforms were detected and semi-automatically sorted using a modified version of the offline spike sorter *Wave_clus*^{43,49}. 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 PCs, 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 different average waveforms (templates) was done to detect overlaid waveforms⁵⁰. To achieve this, 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 discriminate analysis for classification of waveforms. After spike identification, the target template was subtracted from the filtered signal of the corresponding channel to reduce artifacts during the detection of the next template. This procedure allowed us to detect spikes with a temporal overlap up to 0.2 ms. Unit isolation was evaluated again, based on the five criteria mentioned above, to determine the final classification of all units into single or non-single units. Stationary of firing rate was checked for all units and in case it was not stable over the entire recording session (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 single units). Only single units fulfilling all of these criteria, and no multi-units, were further used in this study.

Data preprocessing

After sorting, single neuron spike events were binned in non-overlapping 1-ms windows to produce a continuous firing rate signal (1 kHz) and smoothed with a Gaussian window ($\sigma = 60$ ms). Data were aligned to cue and movement onset for the instructed- and free-choice-task (cue onset: -800 to 1500 ms, and movement onset: -300 to 500 ms) and additionally for the second cue for the delayed-instructed-task (first cue onset: -800 to 1500 ms; second cue onset -150 to 750ms, and movement onset: -300 to 500 ms), since neuronal activity was

2.2. Three information subspaces explain the category-free population dynamics

locked to these events, with variable memory periods between them. The time range of the alignments was chosen in order to have as small as possible overlap and an as smooth as possible transition between them. The different alignments were combined to produce a continuous signal. Average firing rates were then calculated by averaging over all trials and alignments per condition and single unit.

Statistics

For most statistics tests standard functions from MATLAB Statistics and Machine Learning Toolbox were used if not stated otherwise below. In order to test for unimodality we used the nonparametric Hartigan's Dip Test⁵¹ implemented by F. Mechler and freely downloadable at: <http://www.nicprice.net/diptest/>. In case of time series where multiple time points were tested we used three versions of a non-parametric cluster-based permutation test⁵² (cluster-based permutation *t*-test, 1-way ANOVA and 2-way ANOVA) and two versions of a non-parametric cluster-based surrogate test³⁴ (cluster-based surrogate *t*-test, and cluster-based surrogate Hartigan's Dip Test) to deal with the multiple comparison problem. Cluster-based permutation and surrogate tests are based on clustering of adjacent time-samples exceeding a set threshold. The four tests used by us only differ in the statistic used for selecting the threshold and whether they were used to compare different conditions with each other (permutation tests) or for testing against a surrogate condition (surrogate tests). For the cluster-based surrogate *t*-test, which was used to test for significant modulation from baseline (fixation period) for each condition, an equal number of trials was generated from a homogeneous Poisson process using the baseline firing rate distribution. In case of the cluster-based Hartigan's Dip Test an equal sized set of random values, corresponding to the number of neurons and time points tested, was used as a surrogate condition. The clustering and testing for all four tests was carried out as follows: First, the statistic for selecting the threshold was estimated for all time points. Next, all values (*t*-values, *F*-values, and dip-values, respectively) exceeding a threshold corresponding to an alpha-level of 0.05 were selected. In the next step, adjacent values exceeding the set alpha-level were considered as clusters, extracted, and their values summed. A test distribution was generated by randomly permuting trials of the different conditions by randomly reassigning them to the different groups while maintaining the group size. For each partition (1000 partitions) the thresholding procedure and clustering was repeated.

2.2. Three information subspaces explain the category-free population dynamics

From every partition the largest cluster-level statistic was used to generate a largest chance cluster distribution.

The final statistical test was performed by calculating a p-value for each real cluster under the largest chance cluster distribution and comparing them with a critical alpha-level.

Thus, the multiple comparisons for each sample are replaced by a single comparison, replacing the need to make multiple comparisons.

Euclidean distance analyses

The Euclidean distance was calculated between all time points of trial-averaged activity within and across conditions of all single units across both areas in steps of 10ms. The square root of average firing rates was taken before calculating the Euclidean distance to ensure that results were not dominated by a few high-rate neurons^{6,7}. The square-root-transform was used although firing rates at all time were found to be approximately log normal distributed (Fig. 3d), since the log-transform distorts values close to 0 and is not defined for 0, while the square-root-transform is robust in the range between 0 and 1 and commonly used in literature for the same purpose^{28,29,53}. In order to make the Euclidean distance comparable between areas and datasets, for which different number of neurons were recorded, we normalized by the square root of number of neurons to obtain Euclidean distance per neuron. As a control, we also calculated the Euclidean distance of raw average rates as well as log average rates with minimum log-rates set to -3 and giving a similar temporal pattern as in Fig. 4b (data not shown).

For cluster analyses we used a well-known community structure analyses from Newman⁵⁴ that iteratively finds non-overlapping groups of conditions that minimizes the within-group distance and maximizes the between-group distance implemented by M. Rubinov Mechler and freely downloadable at: <https://sites.google.com/site/bctnet/>. Importantly, no assumption on the number of clusters is required.

Linear discriminant analyses (LDA)

LDA was used to estimate projections of single trial activity of all parallel recorded single units into a one-dimensional space best separating the power and precision condition. We used the standard function from MATLAB Statistics and Machine Learning Toolbox to calculate linear discriminant analysis (LDA). In all cases LDA were weighted for number of

2.2. Three information subspaces explain the category-free population dynamics

trials of the power and precision condition and to prevent overfitting we applied leave-one-out cross-validation. In agreement with the Euclidean distance analyses, LDA were calculated based on the square root of smoothed single trial firing rates. Projections were estimated in steps of 10ms separately for the three task-types. For most analyses only the projections from the instructed-task were used if not stated otherwise.

The degree of task specific information was estimated as the percentage of correctly separated single trial trajectories for grip-type per task-type and dataset. Since only two grip-types were performed the chance level was 50%. In order to test if the degree of task specific information was significant different than chance level we applied a permutation test estimating the degree of task specific information for 10000 random separations. To this end, single trial trajectories for the two grip-types were randomly permuted into two equal sized sets of trajectories and the degree of task specific information was estimated again. Finally, the p-value was calculated under the distribution of all random degree of task specific information estimations separate for each time point of each projection and compared to an alpha-level of 0.01.

The nearly orthogonal projections estimated from the instructed-task at the time points cue onset + 180ms, cue onset + 1170ms, and movement onset + 150ms were selected as visual, preparatory, and movement subspace projections, respectively. For the evaluation of how much task-specific information is captured by the three information subspaces compared to the optimal projection at each time point, another set of projections was estimated as described above, but using only the activity of the three subspaces as input, as opposed to all neurons.

Projection angle index of response similarity (PAIRS)

To test for clustering in the individual contribution of neurons to the three subspaces, we used the PAIRS analyses as described previously⁴ using freely available code at:

<http://repository.cshl.edu/30912/>. The three projection vectors for the three subspaces

were used as input for the PAIRS analyses separately per dataset. To ensure that they

captured completely independent variance, although they were nearly orthogonal, we orthogonalized them using the Gram-Schmidt algorithm. Briefly, the PAIRS analysis

computes the distribution of the average angle of each units n -dimensional weight vector to its k nearest neighbours. In our case n was three weights from the subspace projections and

2.2. Three information subspaces explain the category-free population dynamics

k was set to 10. We set k to a fixed value to be in a comparable range for all datasets. Setting k to any value between 6 to 15 had a negligible effect on the results (data not shown). The median of the nearest neighbour angle distribution was compared to median values of 10,000 simulated datasets with the same number of neurons and dimensions randomly drawn from a Gaussian distribution, which were processed as described before. The percentile of instances where the simulated distribution values exceeded or undercut (two-sided test) the empirical value corresponds to the p -value. One test across all datasets was performed by subtracting each real median value from the 10,000 simulated median value per dataset, pooling them, and testing against 0. As a control, we also tested the datasets of both monkeys separately and found the same results (data not shown).

To test if there was any potential clustered neuronal variance missed by the three subspace projections, we repeated the PAIRS analysis by using unsupervised principle component analyses (PCA) for dimensionality reduction. We used the weight vectors of the first 15 PCs of each dataset explaining more than 99% of the trial-averaged variance and the number of k nearest neighbours was set to 3. As a control, we also used the weight vectors of the first 8, 10, and 12 PCs, with no effect on the results (data not shown).

Recurrent neural network

In order to model the planning and execution of a grasping task with a decision making component on a single trial basis, we implemented the dynamical system, $\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}, \mathbf{u})$, using a standard continuous RNN equation of the form

$$\tau \dot{x}_i(t) = -x_i + \sum_{k=1}^N J_{ik} r_k(t) + \sum_{k=1}^I B_{ik} u_k(t) + b_i^x + \varepsilon_i(t) \quad (3)$$

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 . Each unit has an offset bias, b_i^x , and each unit receives normally distributed noise, ε_i , with standard deviation 0.01, at every time point. In order to allow for the

emergence of “decisions” on individual trials, the noise injected to the system on each trial

2.2. Three information subspaces explain the category-free population dynamics

was unique to that trial and remained fixed over training. 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 200 ms before movement onset, and two inputs representing the visual presentation of the grip type, either power or precision grip. 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 1.5⁵⁵. We used a fixed time constant of 50 ms for τ , with Euler integration every 10 ms.

In a separate recording session, the kinematics of multiple repetitions of power and precision grip were recorded using a tracking glove⁵⁶ to produce 27 degrees of freedom in joint angles. These kinematics were further transformed into a set of 50 muscle length measurements using a musculoskeletal model⁵⁷. The network was required to generate the average muscle velocities in 50 dimensions during power or precision grip over the first 300 ms of movement, where movement onset was determined by a threshold crossing in elbow position. In order to account for neural conduction delays and muscle activation times, the desired kinematics were shifted 50 ms backward relative to the corresponding neural signal. The output of the network was defined as a linear readout of the internal network

$$z_i(c, t) = \sum_{k=1}^N \mathbb{E}_{ik} r_k(c, t) + b_i^z \quad (4)$$

where z represents the two kinematic 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, and b_i^z , which is a bias term for each output dimension.

The input weights, B , internal connectivity, J , output weights, W , and all biases, were trained using Hessian Free Optimization⁵⁸ (freely available code: <https://github.com/sussillo/hfopt-matlab>) also utilized in Sussillo et al.³⁸ and Michaels et al.²³. The error function used to optimize the network considered the difference between the output of the linear readout and the desired muscle velocity profiles, v ,

$$E_i(c, t) = z_i(c, t) - v_i(c, t) \quad (5)$$

at each time point, t , each output dimensions, i , and each individual trial, c . We report normalized error, which is the sum of the squared error from Eq 3 over all times,

2.2. Three information subspaces explain the category-free population dynamics

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.³⁸). 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 contraction⁵⁹). The hyper-parameters used for these regularizations were $3e-2$, $1e-4$, and $1e-4$, respectively.

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.

Canonical Correlation Analysis

Canonical Correlation Analysis (CCA) was undertaken to compare the simulated activity within the neural network to the recorded neural population data. While the simulation was carried out for individual trials, a single CCA analysis was carried out on trial-averaged data aligned to both grip cue onset and movement onset that was concatenated to form a single trajectory. Before CCA, all units in both the neural data and the simulated data were reduced to 8 principal components, where the data was of the form $ct \times n$, where c is the number of conditions, t is the amount of time per trial, and n is the number of units. CCA produces new dimensions that are linear combinations of the principal components of each data set (neural or simulated) that are highly correlated between data sets and orthogonal to all other canonical variables.

2.2. Three information subspaces explain the category-free population dynamics

References

1. Gail, A. Neural Dynamics in Monkey Parietal Reach Region Reflect Context-Specific Sensorimotor Transformations. *J Neurosci* **26**, 9376–9384 (2006).
2. Rishel, C. A., Huang, G. & Freedman, D. J. Independent Category and Spatial Encoding in Parietal Cortex. *Neuron* **77**, 969–979 (2013).
3. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. 1–19 (2013). doi:10.1038/nature12742
4. Raposo, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nature Publishing Group* **17**, 1784–1792 (2014).
5. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *Elife* **4**, e04677 (2015).
6. Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* **68**, 387–400 (2010).
7. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
8. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications* 1–15 (2016). doi:10.1038/ncomms13239
9. Siegel, M., Buschman, T. J. & Miller, E. K. Cortical information flow during flexible sensorimotor decisions. *Science* **348**, 1352–1355 (2015).
10. Schaffelhofer, S. & Scherberger, H. Object vision to hand action in macaque parietal, premotor, and motor cortices. *Elife* (2016). doi:10.7554/eLife.15278.001
11. Murata, A., Gallese, V., Luppino, G., Kaseda, M. & Sakata, H. Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP. *J. Neurophysiol.* **83**, 2580–2601 (2000).
12. Baumann, M. A., Fluet, M. C. & Scherberger, H. Context-Specific Grasp Movement Representation in the Macaque Anterior Intraparietal Area. *J. Neurosci.* **29**, 6436–6448 (2009).
13. Fluet, M. C., Baumann, M. A. & Scherberger, H. Context-Specific Grasp Movement Representation in Macaque Ventral Premotor Cortex. *J. Neurosci.* **30**, 15175–15184 (2010).
14. Andersen, R. A. & Cui, H. Intention, action planning, and decision making in parietal-frontal circuits. *Neuron* **63**, 568–583 (2009).
15. Menz, V. K., Schaffelhofer, S. & Scherberger, H. Representation of continuous hand and arm movements in macaque areas M1, F5, and AIP: a comparative decoding study. *J. Neural Eng.* **12**, 1–19 (2015).
16. Pesaran, B., Nelson, M. J. & Andersen, R. A. Free choice activates a decision circuit between frontal and parietal cortex. *Nature* **453**, 406–409 (2008).
17. Klaes, C., Westendorff, S., Chakrabarti, S. & Gail, A. Choosing Goals, Not Rules: Deciding among Rule-Based Action Plans. *Neuron* **70**, 13–13 (2011).
18. Freedman, D. J. & Assad, J. A. Neuronal Mechanisms of Visual Categorization: An Abstract View on Decision Making. *Annu. Rev. Neurosci.* **39**, 129–147 (2016).
19. Todorov, E. Direct cortical control of muscle activation in voluntary arm movements: a

2.2. Three information subspaces explain the category-free population dynamics

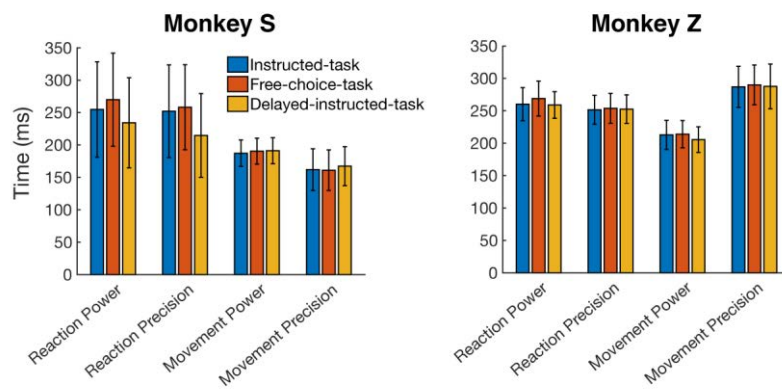
- model. *Nat. Neurosci.* **3**, 391–398 (2000).
20. Georgopoulos, A. P. & Carpenter, A. F. Coding of movements in the motor cortex. *Curr. Opin. Neurobiol.* **33**, 34–39 (2015).
 21. Cisek, P. Integrated Neural Processes for Defining Potential Actions and Deciding between Them: A Computational Model. *J Neurosci* **26**, 9761–9770 (2006).
 22. Churchland, M. M. & Shenoy, K. V. Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *J. Neurophysiol.* **97**, 4235–4257 (2007).
 23. Michaels, J. A., Dann, B. & Scherberger, H. Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLoS Comput Biol* **12**, e1005175–22 (2016).
 24. Kaufman, M. T., Seely, J. S., Sussillo, D., Ryu, S. I. & Shenoy, K. V. The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eneuro* (2016). doi:10.1523/ENEURO.0085-16.2016
 25. Yuste, R. From the neuron doctrine to neural networks. *Nat Rev Neurosci* (2015). doi:10.1038/nrn3962
 26. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
 27. Ames, K. C., Ryu, S. I. & Shenoy, K. V. Neural Dynamics of Reaching following Incorrect or Absent Motor Preparation. *Neuron* **81**, 438–451 (2014).
 28. Afshar, A. *et al.* Single-Trial Neural Correlates of Arm Movement Preparation. *Neuron* **71**, 555–564 (2011).
 29. Michaels, J. A., Dann, B., Intveld, R. W. & Scherberger, H. Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. *J. Neurosci.* **35**, 11415–11432 (2015).
 30. Stokes, M. G. *et al.* Dynamic coding for cognitive control in prefrontal cortex. *Neuron* **78**, 364–375 (2013).
 31. Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 1–6 (2013).
 32. Kobak, D., Brendel, W., Constantinidis, C. & Feierstein, C. E. Demixed principal component analysis of neural population data. *Elife* (2016). doi:10.7554/eLife.10989.001
 33. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**, 1–10 (2014).
 34. Dann, B., Michaels, J. A., Schaffelhofer, S. & Scherberger, H. Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates. *Elife* **5**, 2870 (2016).
 35. Shadlen, M. N. & Kiani, R. Perspective. *Neuron* **80**, 791–806 (2013).
 36. Cisek, P. Making decisions through a distributed consensus. *Curr. Opin. Neurobiol.* **22**, 927–936 (2012).
 37. Buzsáki, G. & Mizuseki, K. The log-dynamic brain: how skewed distributions affect network operations. *Nat Rev Neurosci* **15**, 264–278 (2014).
 38. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Publishing Group* 1–12 (2015). doi:10.1038/nn.4042
 39. Stokes, M. G. ‘Activity-silent’ working memory in prefrontal cortex: a dynamic coding framework. *Trends in Cognitive Sciences* **19**, 1–12 (2015).
 40. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the

2.2. Three information subspaces explain the category-free population dynamics

- null space: permitting preparation without movement. *Nature Publishing Group* 1–12 (2014). doi:10.1038/nn.3643
41. Sussillo, D. ScienceDirectNeural circuits as computational dynamical systems. *Curr. Opin. Neurobiol.* **25**, 156–163 (2014).
 42. Bosco, A., Breveglieri, R., Chinellato, E., Galletti, C. & Fattori, P. Reaching Activity in the Medial Posterior Parietal Cortex of Monkeys Is Modulated by Visual Feedback. *J Neurosci* **30**, 14773–14785 (2010).
 43. Kraskov, A., Dancause, N., Quallo, M. M., Shepherd, S. & Lemon, R. N. Corticospinal neurons in macaque ventral premotor cortex with mirror properties: a potential mechanism for action suppression? *Neuron* **64**, 922–930 (2009).
 44. Fries, P. Perspective. *Neuron* **88**, 220–235 (2015).
 45. Buzsáki, G. Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron* **68**, 362–385 (2010).
 46. Engel, A. K. & Fries, P. Beta-band oscillations--signalling the status quo? *Curr. Opin. Neurobiol.* **20**, 156–165 (2010).
 47. Research, C. O. G. F. T. U. O. A. I. N. A. B., National Research Council, Studies, D. O. E. A. L. Institute for Laboratory Animal Research. *Guidelines for the Care and Use of Mammals in Neuroscience and Behavioral Research*. (National Academies Press, 2003).
 48. Musial, P. G., Baker, S. N., Gerstein, G. L., King, E. A. & Keating, J. G. Signal-to-noise ratio improvement in multiple electrode recording. *Journal of Neuroscience Methods* **115**, 29–43 (2002).
 49. Quiroga, R. Q., Nadasdy, Z. & Ben-Shaul, Y. Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput* **16**, 1661–1687 (2004).
 50. Gozani, S. N. & Miller, J. P. Optimal discrimination and classification of neuronal action potential waveforms from multiunit, multichannel recordings using software-based linear filters. *Biomedical Engineering, IEEE Transactions on* **41**, 358–372 (1994).
 51. Hartigan, J. A. & Hartigan, P. M. The dip test of unimodality. *The Annals of Statistics* (1985).
 52. Maris, E. & Oostenveld, R. Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods* **164**, 177–190 (2007).
 53. Yu, B. M. *et al.* Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**, 614–635 (2009).
 54. Newman, M. E. J. Modularity and community structure in networks. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 8577–8582 (2006).
 55. Sussillo, D. & Abbott, L. F. Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron* **63**, 544–557 (2009).
 56. Schaffelhofer, S. & Scherberger, H. A new method of accurate hand- and arm-tracking for small primates. *J. Neural Eng.* **9**, 026025–14 (2012).
 57. Schaffelhofer, S., Sartori, M., Scherberger, H. & Farina, D. Musculoskeletal Representation of a Large Repertoire of Hand Grasping Actions in Primates. *IEEE Trans. Neural Syst. Rehabil. Eng.* **23**, 210–220 (2015).
 58. Martens, J. & Sutskever, I. Learning recurrent neural networks with hessian-free optimization. in (2011).
 59. Rifai, S., Muller, X., Glorot, X., Mesnil, G. & Bengio, Y. Learning invariant features through local space contraction. *arXiv.org* (2011).

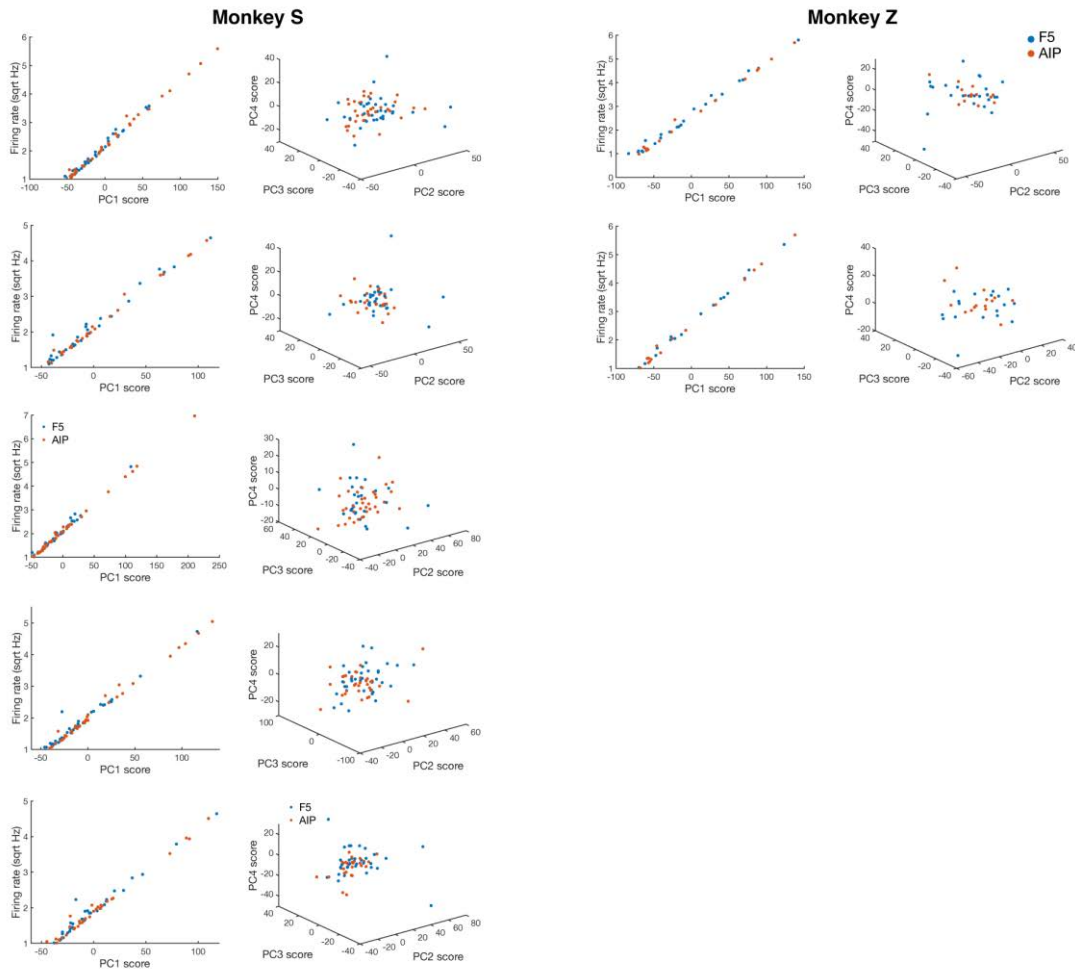
2.2. Three information subspaces explain the category-free population dynamics

Supplementary Figures



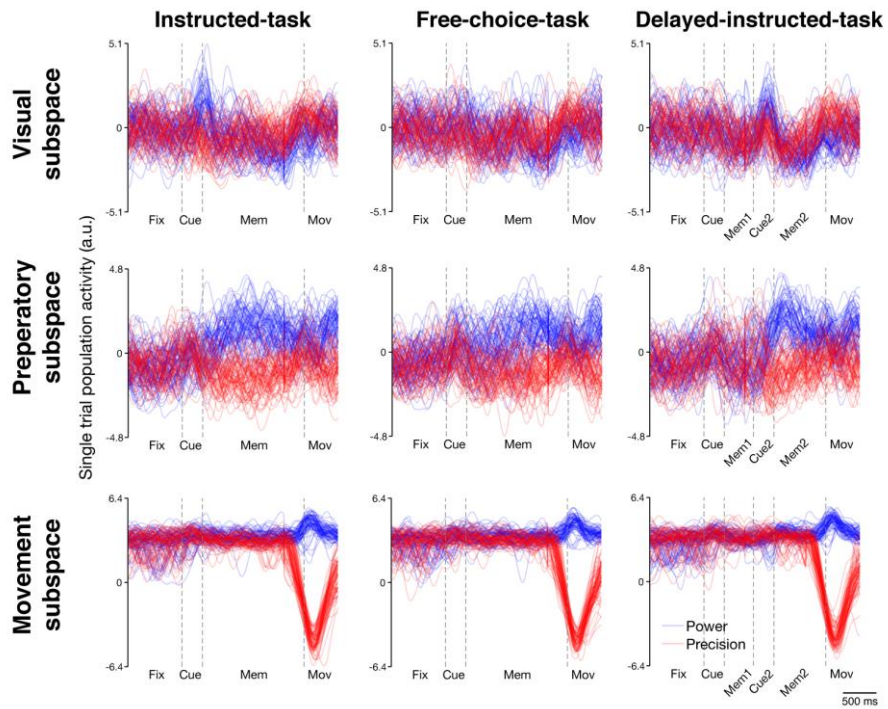
Supplementary Figure 1 Reaction and movement times separate for monkeys, grip- and task-types. The average of all successful trials of all datasets per monkey is shown (4385 and 2167 for monkey S and Z, respectively) with standard deviation across trials.

2.2. Three information subspaces explain the category-free population dynamics



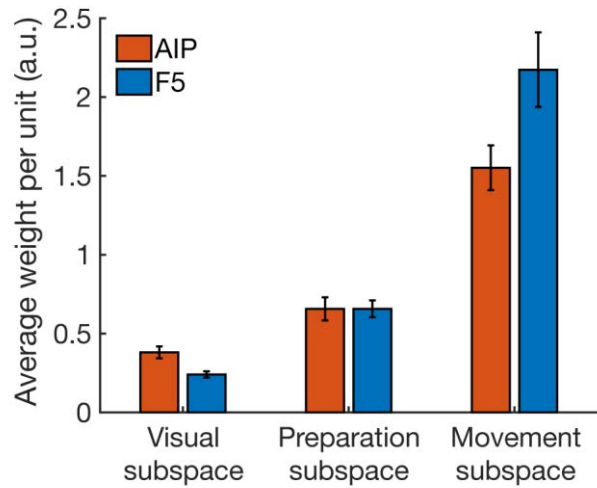
Supplementary Figure 2 Projection of all units from all datasets but the two displayed in Fig. 3 onto the first 4 PCs. PCA was applied on the trial averaged responses of all single units recorded in parallel per dataset with conditions \times time as variables and units as observations. The square root of average firing rates was taken before calculating PCA that is also why the first PC is plotted against the square root of the average firing rate. The results of the PCAs is valid despite the number of variables being higher than the number of observations, since the first 4 PCs already explain more than 90% of the condition and temporal variance.

2.2. Three information subspaces explain the category-free population dynamics



Supplementary Figure 3 **Single trial trajectories for the projections into the three subspaces of one example dataset from monkey S. 150 trials (75 trials per grip-type) were randomly chosen per task-type. Not all single trial trajectories were displayed for a better overview. The range per subspace projection was fixed for better comparison between the different task-types.**

2.2. Three information subspaces explain the category-free population dynamics



Supplementary Figure 4 **Contribution the three information subspaces per area. The subspace projection vectors per dataset were first split per area, then the absolute value was taken since negative weights contribute as much as positive weights, and finally the individual weights per single unit were averaged separate for AIP and F5. Displayed is the average across datasets with errorbars representing standard error across datasets.**

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

Jonathan A Michaels^{1,*}, Benjamin Dann^{1,*}, Rijk W Intveld¹, Hansjörg Scherberger^{1,2}

¹Deutsches Primatenzentrum GmbH, Kellnerweg 4, 37077 Göttingen, Germany

²Faculty of Biology, Georg-August-Universität Göttingen, 37073 Göttingen, Germany

*Equal contribution

Acknowledgements: We would like to thank Natalie Bobb, Ricarda Lbik, and Matthias Dörge for technical assistance, and Roman Eppinger for preliminary analysis. This work was supported by the Deutsche Forschungsgemeinschaft (DFG; SCHE 1575/1-1 & 3-1).

2.3. Probing the continuum of immediate to withheld grasping movements

Abstract

Our voluntary grasping actions lie on a continuum between immediate action and waiting for the right moment, depending on the context. Therefore, studying grasping requires investigating how preparation time affects this process. Two macaque monkeys performed a grasping task with a short instruction followed by an immediate or delayed go cue (0-1300 ms) while we recorded in parallel from neurons in the hand area (F5) of the ventral premotor cortex and the anterior intraparietal area (AIP). Initial population dynamics followed a fixed trajectory in the neural state space unique to each grip type, reflecting unavoidable preparation, then diverged depending on the delay. Although similar types of single unit responses were present in both areas, population activity in AIP stabilized within a unique memory state while F5 activity continued to evolve, tracking anticipation. Intriguingly, activity during movement initiation clustered into two trajectory clusters, corresponding to movements that were either 'as fast as possible' or withheld movements, demonstrating a widespread state shift in the fronto-parietal grasping network when movements must be withheld. Our results reveal how dissociation between static and dynamic components of movement preparation as well as differentiation between cortical areas is possible through population level analysis.

2.3. Probing the continuum of immediate to withheld grasping movements

Introduction

Some actions, such as reacting to a spilling 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 (Wise 1985; Riehle and Requin 1989), 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 and disentangle object properties and hand shapes (Baumann et al. 2009; Fluet et al. 2010; Townsend et al. 2011; Schaffelhofer et al. 2015; Schaffelhofer and Scherberger 2016), 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. 2006; Afshar et al. 2011) and parietal cortex (Snyder et al. 2006; 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). A recent study showed that task features are randomly distributed over many neurons of an area, questioning the neuron doctrine (Raposo et al. 2014). These recent studies are made possible by the increasing implementation of large-scale sequential and parallel recordings 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-

2.3. Probing the continuum of immediate to withheld grasping movements

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. One study showed how preparatory activity in motor cortex acts as an initial state for subsequent movement dynamics (Churchland et al. 2012). However, another 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 delayed movements (Ames et al. 2014), suggesting that successful preparation of the same reach may be achieved through different neural trajectories. After adequate preparation time activity stabilized in the state space, while other studies have shown that premotor cortex may track time or expectation (Carnevale et al. 2015). Only analyzing the full continuum of preparation from immediate to fully planned movements can provide an understanding of the complex interaction between planning and movement. Furthermore, it has been proposed that delayed and immediate movements are controlled quite differently (Braver 2012), a feature that has not been investigated in premotor cortex. Crucially, to understand how diversely the motor system encodes and executes movements, multiple distributed brain regions known to be involved in the preparation of the same movement must be investigated together.

To address these questions, we recorded neural populations from the grasping circuit (Luppino et al. 1999) consisting of the hand area (F5) of the ventral premotor cortex (PMv) and the anterior intraparietal area (AIP) while 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 the neural states achieved during longer delays were bypassed during immediately cued grasps. However, the initial trajectory was specific to each grip type, but the same regardless of delay, providing evidence that this activity may be required for successful movement. Activity in AIP stabilized during long delays, but activity in F5 was highly dynamic and well matched the subjective probability of a cue throughout the memory period, implicating differing functional roles of the two areas. Interestingly, activity in both areas formed distinct long and short delay trajectory clusters following the go cue, demonstrating that a network-wide shift occurs when movements are withheld and executed from memory. Crucially, our findings highlight the dissociation of static and dynamic components of movement preparation as well as the function of cortical areas through population analysis.

2.3. Probing the continuum of immediate to withheld grasping movements

Results

Task and behavior

To investigate the continuum of grasp movement preparation, we trained two macaque monkeys (B and S) 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 (Materials and Methods). Monkeys fixated a central point (red), received a grip cue (300 ms) corresponding to either precision (white) or power grip (green), and were cued to perform this grip following a variable delay when the central fixation point turned off (Fig. 1a-b). 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 S1). In addition to the normal task, we also randomly inserted no-movement trials to ensure that monkeys waited for the go cue before acting. Both monkeys completed these trials successfully (monkey B: 100%; monkey S: 97.7%).

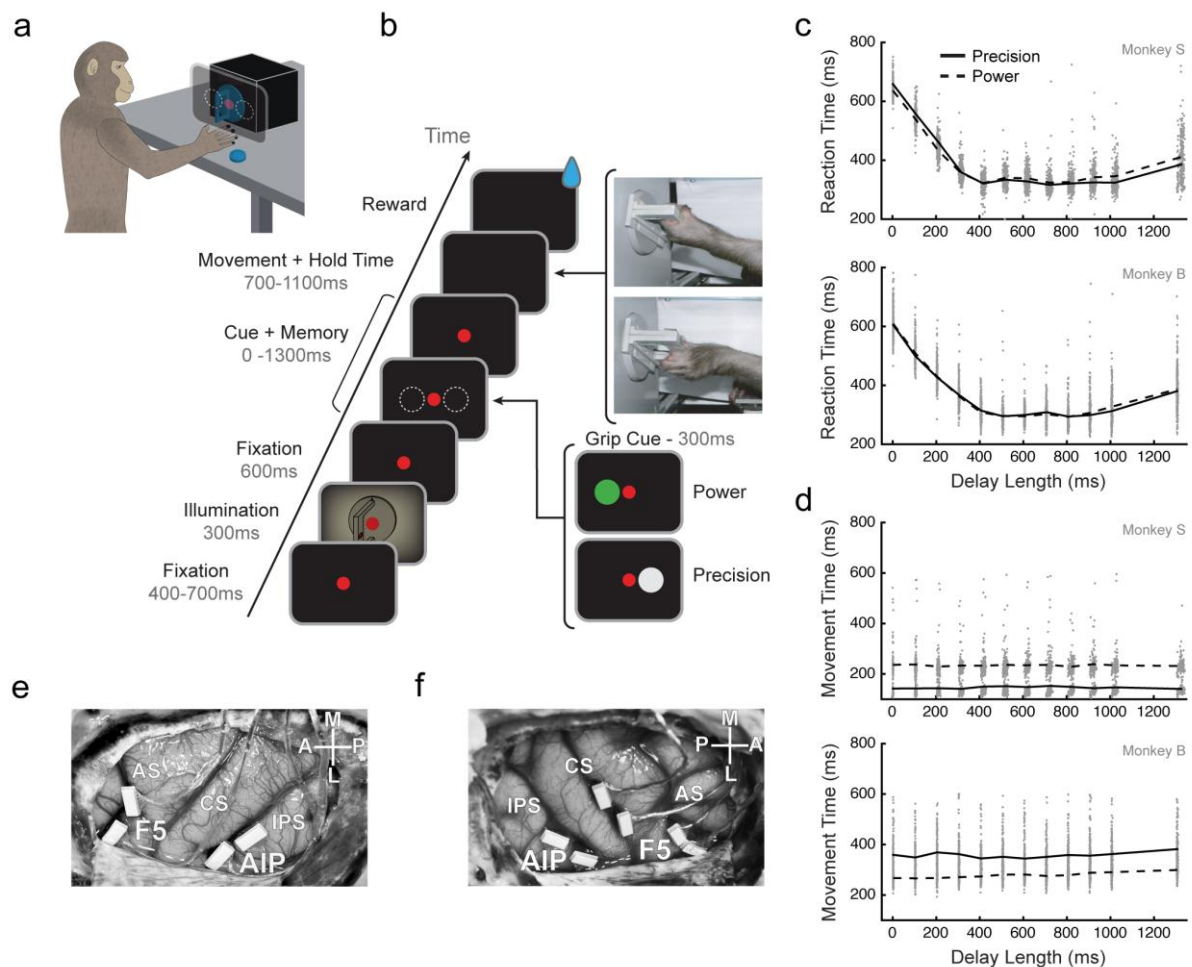


Fig 1. Task design, implantation, and behavior. (a) Illustration of a monkey in the experimental setup. The cues were presented on a masked monitor and reflected by a mirror such that cues appeared super-imposed on the

2.3. Probing the continuum of immediate to withheld grasping movements

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 and movement time against delay length for both monkeys. The solid line represents the mean for each delay bin. **(e and f)** Array locations for monkey S **(e)** and B **(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), but not used in this study. 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.

In addition to number of correctly executed trials, reaction times (RTs) and movement times (MTs) of the monkeys provided useful insight into the performance of the task. RT decreased steadily with increasing amounts of preparation (Rosenbaum 1980), approaching a minimum after approximately 400 ms of preparation (Fig. 1c), well in line with previous findings (Churchland et al. 2006). RT increased slightly for the longest delay. For monkey S, MT did not correlate with length of the delay period (Fig. 1d, $p = 0.9$), indicating that although RT was slower for short delays, movements were only initiated once they were fully prepared. In monkey B there was a small positive correlation between delay and MT (Fig. 1d, $r = 0.11$). Movement kinematics were likely similar regardless of delay, since the variability in mean movement times between different delay lengths were extremely small. The standard deviations in mean movement times (Monkey S, precision grip: 3.5 ms SD, power grip: 1.8 ms SD; Monkey B, precision grip: 14.2 ms SD, power grip: 10.8 ms SD) provide evidence that the kinematics of the movements did not vary between delays, especially for monkey S. The number of errors showed no clear relationship to the length of the delay period, and the number of errors was extremely low, providing evidence that the monkeys could complete all conditions equally well.

Neural responses

We recorded six sessions of each monkey using floating microelectrode arrays for a total of 128 channels (64 in each area) simultaneously in F5 and AIP (Fig. 1e,f) and single- and multi-unit activity was isolated (Materials and Methods). 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 individual session information see Table 1. 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

2.3. Probing the continuum of immediate to withheld grasping movements

AIP, 28% and 26% in the 200 ms preceding go cue, and 55% and 45% in the 200 ms following movement onset (t -test, $p < 0.05$), conservatively measured only for movements with a distinct memory period (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), confirming their involvement in grasp coding.

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 (Fig. 2). In many cases the initial cue response was suppressed when the go cue appeared concurrently with the grip information (Fig. 2a,d), while in other cases the initial cue response was present regardless of delay (Fig. 2b,e). Other interesting responses were observed, such as a peak in activity during the memory period (Fig. 2c), and activity during the movement period which differed between delayed and non-delayed grasps (Fig. 2c,f). All of these diverse types of responses were present in both F5 and AIP. The broad variety of unit responses reveals a complex interaction between differing amounts of preparation, making strict categorization of individual neurons difficult.

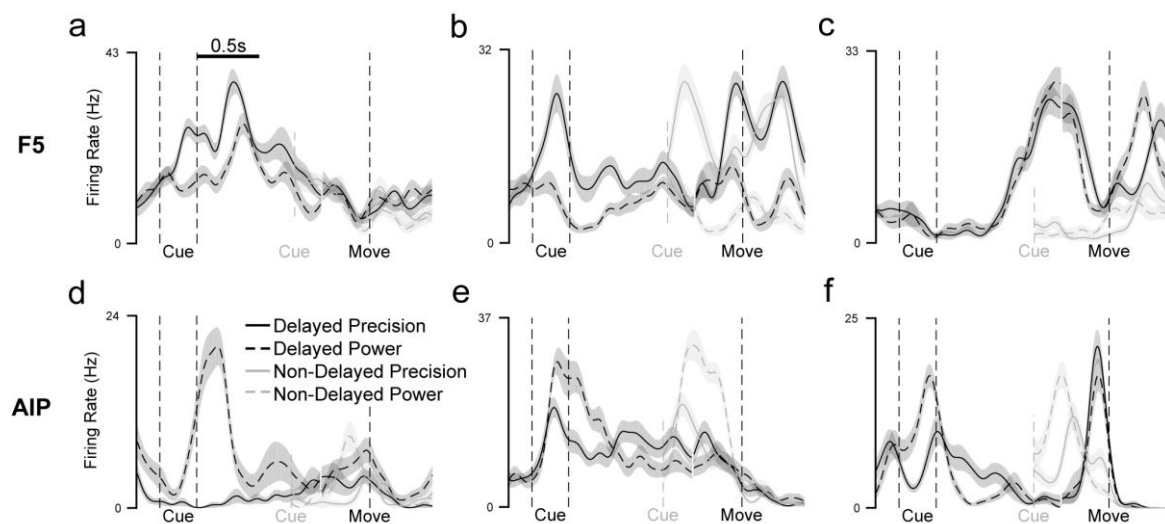


Fig 2. Example average firing rate curves of single-units for delayed (1300 ms) vs. non-delayed (0 ms) grasps. (a-c) Example single-units from area F5 of monkey B showing (a) a completely suppressed cue response during non-delayed grasps, (b) an identical cue response for either delay, (c) differing movement period activity between delayed and non-delayed grasps. (d-f) Similar single-unit examples from AIP of monkeys B and S. Delayed data were aligned to two events, grip cue onset and movement onset and are separated by a gap, which marks the go cue. Non-delayed data were only aligned to movement onset. Dotted gray line represents

2.3. Probing the continuum of immediate to withheld grasping movements

approximate time of cue onset and go cue for non-delayed grasps. The cue was always presented for 300 ms regardless of delay. Curves and shaded bands represent mean and standard error of the mean, respectively.

Visualizing the population response

An alternative approach to categorizing single units is the state space framework, in which all units 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 projected the population activity of all units across both areas for all trials into a lower dimensional space of 10 latent dimensions using Gaussian Process Factor Analysis (GPFA; Materials and Methods). These 10 latent dimensions well captured the variance of both areas. Once the latent dimensions were found, the activity of each area was independently projected into these dimensions in order to compare the contribution of each area. Fig. 3a,c shows the neural trajectories of exemplar data of each monkey (sessions B4, S2) from 100 ms before grip cue onset to 400 ms after movement onset.

In both monkeys the first dimension was a mostly condition-independent movement signal, especially large in F5, a feature observed previously in motor cortex (Kaufman et al. 2016). The other dimensions show varying levels of grip-specific cue responses, delay- or grip-specific memory responses, and strong movement activity. Particularly interesting is latent 3 in Fig. 3a and latent 4 in Fig. 3c, which showed in both monkeys sustained grip selectivity through memory into movement. Plotting latents 2-4 against each other revealed other features (Fig. 3b,d, 100 ms before cue onset to 50 ms after movement onset). Trajectories began in a tight cluster at grip cue onset and remained overlapped for the initial response (200-300 ms) regardless of delay, but specific to each grip type. The trajectories for longer delays continued to evolve for hundreds of milliseconds, but the short delays proceeded to movement onset, bypassing the part of the space achieved by long delays. Interestingly, while activity in AIP congregated in a stable state 500-600 ms after the grip cue, activity in F5 continued to evolve for the entire memory period, never congregating in an area of low variability. Finally, for each grip type short and long delays grouped into two clusters during movement initiation (Fig. 3b, AIP; Fig 3d, F5).

Unique memory state for delayed grasping movements

As we saw in Fig. 3, unique memory states were traversed by the neural trajectory during trials with long delays. To test this possibility statistically, we used a continuous

2.3. Probing the continuum of immediate to withheld grasping movements

distance analysis (Materials and Methods). We measured the minimum Euclidean distance (known as point-to-curve) between each time point on the trajectory of a delayed condition (1000 ms delay condition in steps of 50 ms) and the entire non-delayed trajectory (0 ms delay condition). This was done for the 10 latent dimensions of each area to determine which points in the state space were traversed by both conditions and which were unique to longer delayed movements, separately for each recording session and each grip type. After the cue, distance between delayed and non-delayed trajectories rose and remained significantly above chance level until around movement onset or later in example data sets of both areas and monkeys (Fig. 4a; sessions B3, S2; Bootstrapping procedure with 1000 resamples, $p < 0.05$, cluster-based permutation test; Materials and Methods). Over all grip types and data sets the same effect is present (Fig. 4b), showing that distance between the trajectories was most prevalent until shortly before movement onset. The amount of divergence between the delayed and non-delayed trajectories was very similar in F5 and AIP, indicating that when grasps are cued without a delay the neural population of both areas bypass the states achieved by longer delays. Performing the same analysis on the full neural space without dimensionality reduction produced similar results (data not shown).

2.3. Probing the continuum of immediate to withheld grasping movements

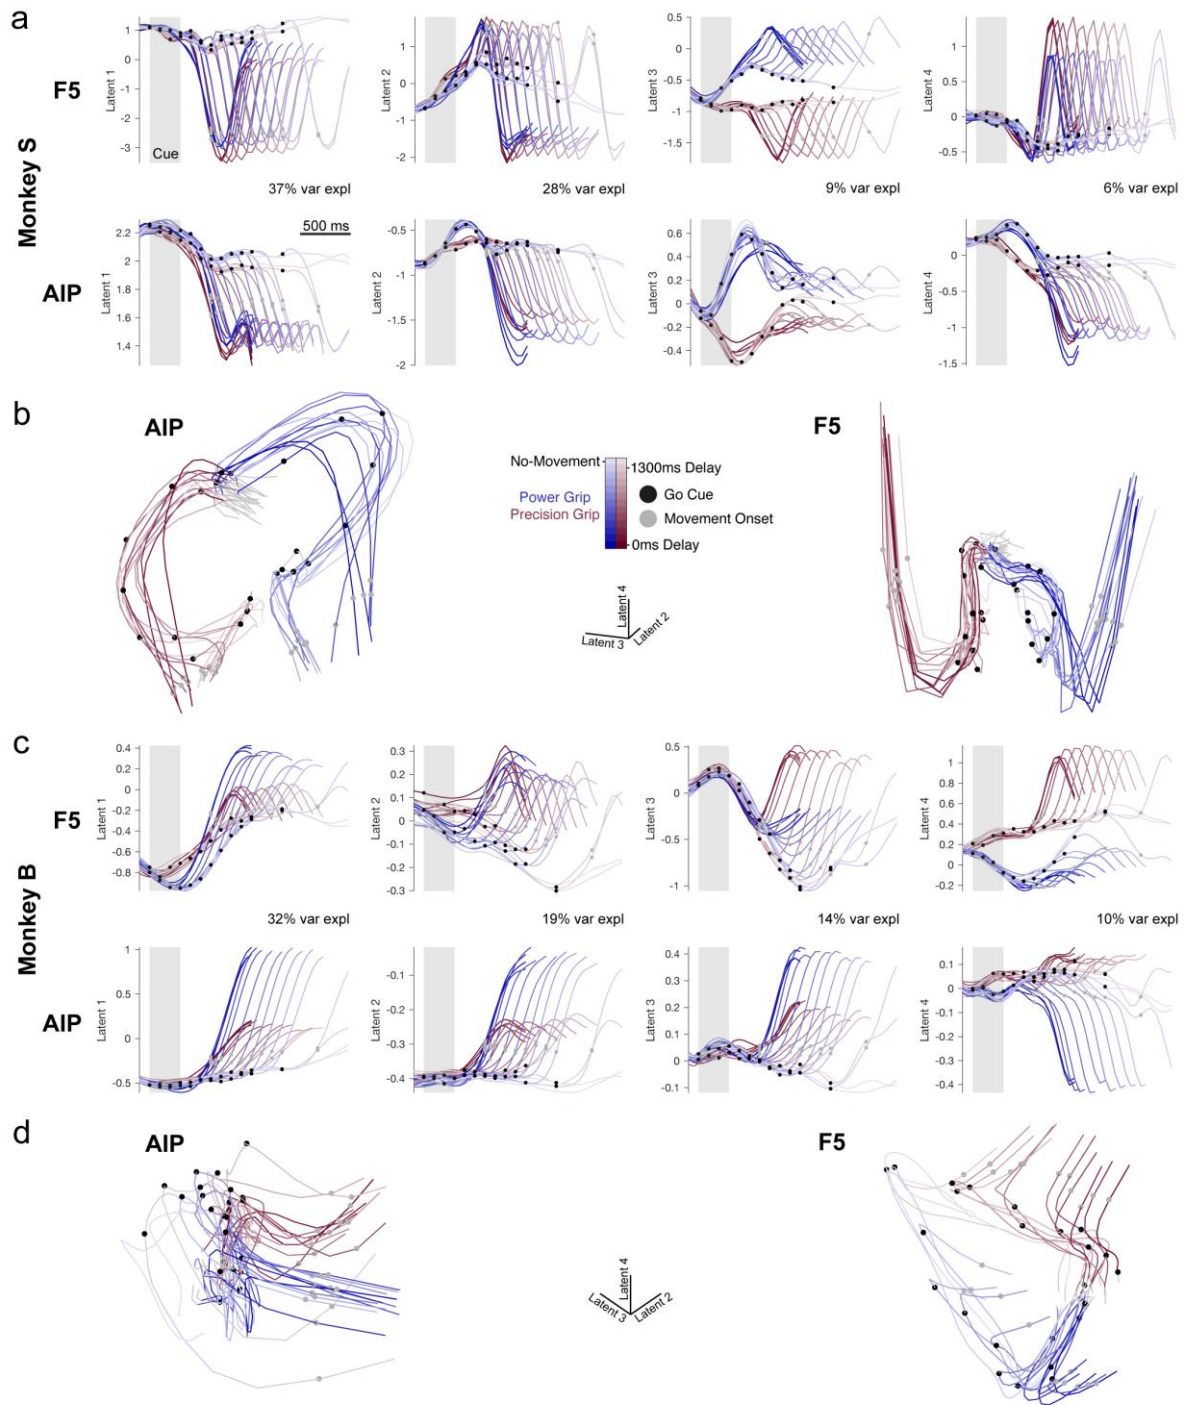


Fig 3. Low-dimensional latent space trajectories of F5 and AIP. Population data of all conditions were projected into a 10 dimensional latent space as determined by GPFA. **(a)** A single session trial-averaged example from monkey S is shown for the first 4 latent dimensions (S4). Trajectories begin 100 ms before the grip cue and end 400 ms after movement onset. **(b)** A 3D plot of the second to fourth latent dimensions plotted from 100 ms before cue onset to 50 ms after movement onset. **(c-d)** same as **(a-b)** for a single session from monkey B (B2).

As mentioned earlier, it appeared in Fig. 3 that the difference between grip types was present before the difference between delays. In other words, the effect of the grip cue appeared before the effect of the go cue. To test this, we repeated the distance analysis with

2.3. Probing the continuum of immediate to withheld grasping movements

a finer time resolution around cue onset (GPFA using steps of 20 ms) and additionally tested the Euclidean distance between grip conditions (Fig. 4c, Materials and Methods). Comparing the first onset of significance between delay and grip effects for each data set separately revealed that grip separation consistently appeared before delay separation in both areas and monkeys (Wilcoxon sign-rank test, F5 monkey S, $p < 0.001$; AIP monkey S, $p < 0.001$; F5 monkey B, $p = 0.003$; AIP monkey B, $p = 0.016$). On average across monkeys and areas, grip separation occurred 128 ms after cue onset and delay separation occurred 352 ms after cue onset.

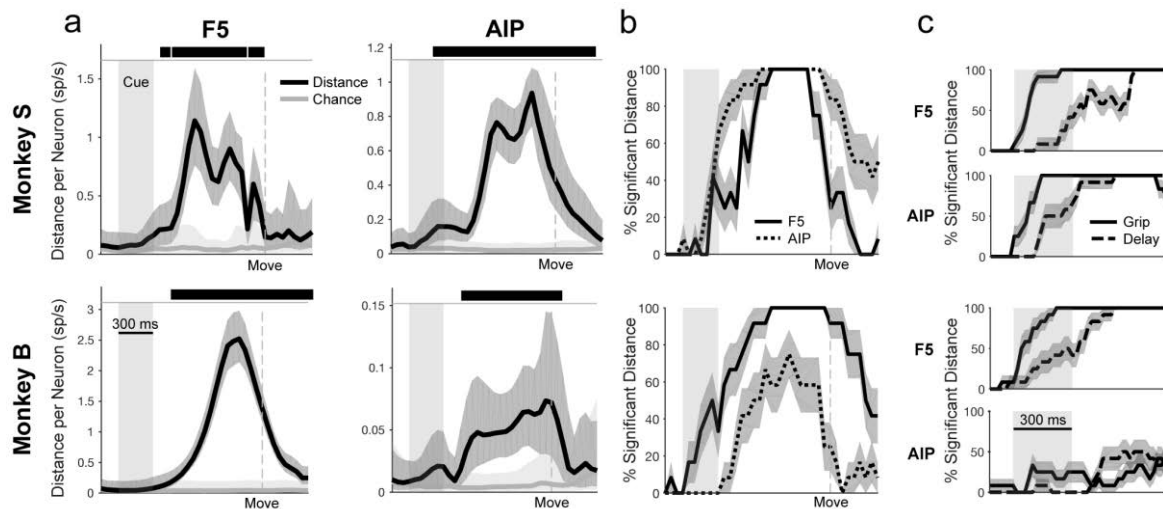


Fig 4. Point-to-curve distance between delayed (1000 ms) and non-delayed (0 ms) trajectories. (a) Minimum Euclidean distance in the latent space between each time point on the delayed trajectory (in steps of 50 ms) and the entire non-delayed trajectory over time for 2 example data sets (B2-Power, S3-Power) from both areas and monkeys. The black line represents the minimum point-to-curve distance between the delayed and non-delayed trajectory, while the gray lines represent the chance level (Materials and Methods). Black bars along the top of plots denote times when the distance is significantly greater than chance level (Bootstrapping procedure with 1000 resamples, $p = 0.05$, Cluster-based permutation test; Materials and Methods). Error bars represent the 5th and 95th percentiles of the distances generated by the bootstrapping procedure. (b) Fraction of significant distances over all data sets and grip types (6 data sets x 2 grip types). Error bars represent the standard error of the mean over data sets and grip types. (c) Difference in onset of grip and delay separation over all data sets and grip types (6 data sets x 2 grip types) at a higher temporal resolution (20 ms bins).

Taken together, these results provide evidence 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, and the activity in the first ~300 ms likely represents unavoidable processing.

Static and dynamic memory states

Given that the trajectories of delayed and non-delayed grasps only overlap for the first ~300 ms of preparation, what is the function and dynamics of the memory period activity? A

2.3. Probing the continuum of immediate to withheld grasping movements

striking feature of the visualization in Fig. 3 was that the F5 activity continually evolved throughout the course of the memory period, while activity in AIP congregated in an area of low variability. To analyze when and if the neuronal trajectory of the two areas stabilized, we systematically compared the Euclidean distance between all pairs of time points along the trajectories for the no-movement trajectories (Fig. 5a, example data sets S6 and B5). Dynamic activity should appear as large distances between trajectories everywhere except the diagonal (points close in time), while static activity should appear as a ‘block’ of activity with a small distance between trajectories.




Fig 5. Neural trajectory stability over the course of no-movement trials. (a) Mean Euclidean distance in the latent space for the no-movement trials between all pairs of time points over both grip types for example data sets in each monkey (sessions B5, S6). For each pair of time points, distance results were tested for a significant difference using a bootstrapping procedure (10000 resamples in steps of 50 ms, $p = 0.01$). The abbreviations *Cue*, *Mem*, and *Rew*, correspond to the cue, memory, and reward epochs, respectively. All plots are clipped at 1 sp/s for visualization. The times where a significant difference was found (in no conditions, one grip type, or both grip types) are shown in (b). (c) Percentage of time points showing a significant difference over all data

2.3. Probing the continuum of immediate to withheld grasping movements

sets and grip types (6 data sets x 2 grip types) of each monkey separately. (d) Mean distance between all time points during the stable portion of the memory period (600 ms – 1800 ms after cue onset) for all individual data sets and grip types (6 data sets x 2 grip types) across areas and paired according to recording session. Stars indicate a significant difference (Wilcoxon sign-rank test, $p < 0.001$).

The strongest differences occurred shortly after cue onset and near reward. Most remarkably, the neuronal trajectory during the memory period in F5 continuously and uniformly progressed in the absence of behavioral events. On the contrary, the neuronal trajectory in AIP stabilized 200-300 ms after cue offset. The effect becomes clearer when visualizing the time points that significantly differed (Fig. 5b, Materials and Methods), showing a stereotypical ‘block’ pattern in AIP and also visible over all data sets (Fig. 5c). Taking the average distance between all time points during the portion of the memory period unaffected by cue or reward (600 ms – 1800 ms after cue onset) showed a significantly more dynamic representation in F5 than AIP (Fig. 5d; Wilcoxon signed-rank test, $p < 0.001$). Similar results were obtained using the full neural space (data not shown). These results indicate a considerably different code at the population level in AIP and F5.

It is also important to consider that the probability of having to perform a movement did not remain constant, since the probability of being in the no-movement condition 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 the change in necessity of the motor plan. To rule out this possibility, we repeated the current analysis on data of a similar experiment in which movements were required in all conditions (Michaels et al. 2015). We found that the same inter-area difference reported here were present (S1 Fig.), lending support to the observed dissociation between areas.

Memory period dynamics

Given the dynamic nature of activity during the memory period, does this activity follow any predictable pattern? As mentioned earlier, some units appeared to change their activity strictly during the memory period (Fig. 2c), even in the absence of behavioral cues. The observed pattern appears similar to the hazard rate, which in the current experiment is the probability of a go cue occurring at any moment, given that the go cue has not appeared yet (Janssen and Shadlen 2005). The form of the hazard rate during no-movement trials and corresponding subjective anticipation function, which takes the monkey’s uncertainty about time into account (Materials and Methods), is shown in Fig. 6a. We fit the average activity of

2.3. Probing the continuum of immediate to withheld grasping movements

each latent dimension (over both areas) to subjective anticipation. The best fitting dimension per data set had an average adjusted R-square of 0.73 for monkey S and 0.88 for monkey B, indicating that anticipation may be significantly represented (mean time shift: -11 ms, w_3 in Eq. 3). Example data sets are shown in Fig. 6b,e (data from session S2 and B4).

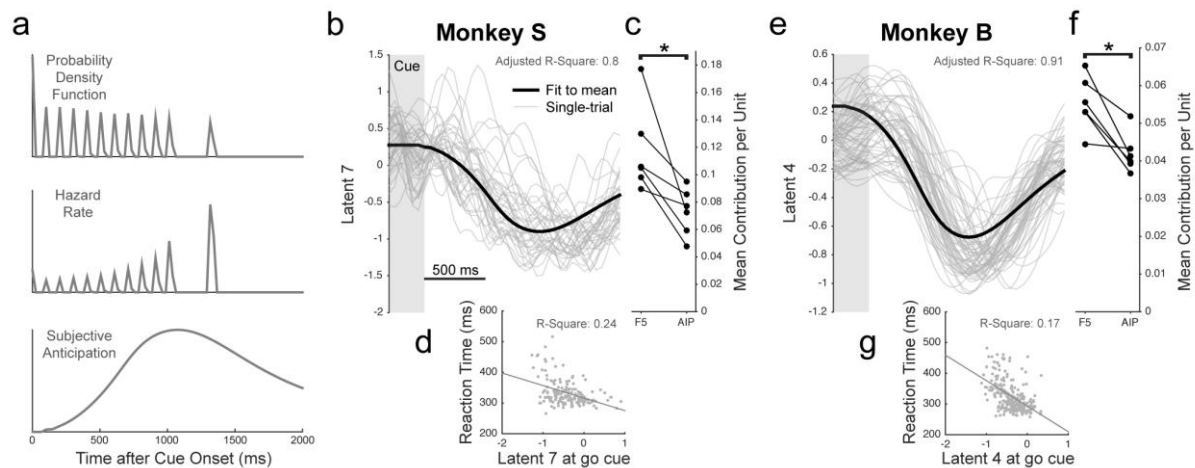


Fig 6. Representation of subjective anticipation across F5 and AIP. (a) Illustration of the probability of a go cue at all times during the delay, the hazard rate (Eq. 1), and the subjective anticipation function (Eq. 2 substituted into Eq. 1). (b) subjective anticipation (Eq. 3) fit to an example latent dimension during the no-movement condition (session S2). (c) Mean contribution per unit in each area to the best latent dimension of each data set. Stars indicate a significant difference (Wilcoxon sign-rank test, $p < 0.001$). (d) Example latent dimension at go cue correlated with single-trial reaction time for delays of at least 800 ms. (e-g) Same as (b-d) for monkey B (session B4).

When comparing the mean contribution per unit (weight in GPFA loading matrix) between areas across data sets to the best fitting latent dimensions, F5 clearly contributes more (Fig. 6c,f, Wilcoxon signed-rank test, $p < 0.001$), with an average of 1.5 times the contribution per neuron, supporting the finding that F5 memory activity was much more dynamic. On average across data sets, the best fitting latent dimension explained the 4th most variance of the 10 dimensions extracted for each data set, corresponding to on average 11% variance explained.

Interestingly, activity on single trials in the ideal latent dimensions at the go cue was correlated with reaction time (Fig. 6d,g; trials with a delay of at least 800 ms), with a mean R-square of 0.17 in monkey S and 0.16 in monkey B, similar to results obtained in F5 with other state space methods (Michaels et al. 2015). For this analysis only the causal portion of all GPFA smoothing kernels were used so that activity at the go cue conservatively reflected only past spikes. Given that the activity in this latent dimension is predictive of reaction time, does being closer or farther away from the movement state predict reaction time in a consistent way? When the absolute difference between the go cue activity and mean

2.3. Probing the continuum of immediate to withheld grasping movements

activity during movement initiation (100 ms before movement onset) was correlated with reaction time, 11 out of 12 data sets produced a positive correlation (mean R-square of 0.1), providing evidence that being closer to the movement initiation state on a given trial led to shorter reaction times.

Converging on movement

As was clearly visible in Fig. 3, the population state at the time of go cue varied greatly between delays, especially in F5. However, activity converged towards a state of lower variability at movement onset. Taking a closer look at a few single units over all delay lengths (Fig. 7a), we can see a large variety of paths before movement initiation.

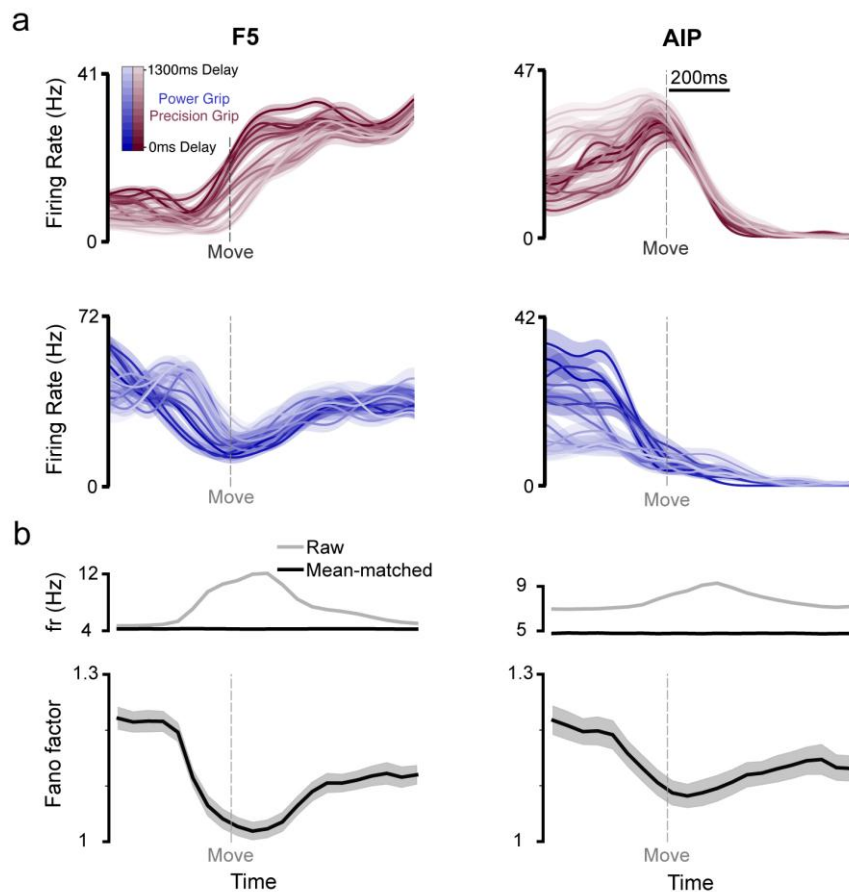


Fig 7. Rapid decrease in trial-to-trial variability during movement initiation. (a) Example average firing rates of single-units in F5 and AIP from both monkeys showing large firing rate differences between the various delay conditions of a single grip type (sessions – top left: B1-precision, top right: S4-precision, bottom left: S6-power, bottom right: S1-power). Error bars represent standard error of the mean across trials. (b) (Top) Mean firing rate before (gray) and after (black) mean-matching for all units (pooled over monkeys, data sets, and conditions). (Bottom) Mean-matched Fano Factor over all units (pooled over monkeys, data sets, and conditions), showing a decrease to near Poisson spiking variability in the 150 ms before movement onset. Error bars represent 95% confidence interval from regression.

2.3. Probing the continuum of immediate to withheld grasping movements

To quantify how trial-to-trial variability changed leading up to movement onset, we calculated the Fano factor over this interval for all trials of all delays, but separately for each grip type (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, Yu, et al. 2010). Since the firing rate increases during movement (Fig. 7b upper panel, gray line), which could possibly affect variability due to saturation of neurons at high firing rates, data were mean-matched (Fig. 7b upper panel, black line) before calculating the Fano factor. Variability based on Fano factor is rapidly reduced 150-200 ms prior to movement onset (Fig. 7b bottom panel), reaching levels almost equivalent to the spontaneous spiking patterns of neurons, which inherently do not spike in a completely predictable way, following a Poisson process. When comparing the Fano factor 300 ms before and 100 ms after movement onset, Fano factor was significantly lower after movement onset for both areas and monkeys tested separately ($p < 0.001$, confidence interval of regression; Materials and Methods), with a stronger effect in F5. These results show that although the pre-movement activity is initially quite variable, this variability is significantly decreasing around movement onset, implicating an internal mechanism that brings trajectories onto a similar path while the movement is being initiated.

Clustering of immediate and withheld movements from memory

In the population visualization in Fig. 3 we saw that the trajectories of short and long delays formed two distinct clusters leading up to movement onset. To visualize the clustering for example data sets in F5, we plotted the activity of all linearly spaced delays (0-1000 ms) of a single grip type around movement onset in an example latent dimension (Fig 8a). Looking specifically at around 100 ms before movement onset, trajectories from the conditions with a delay of 0-400/500 ms and from the conditions with a delay of 400/500-1000 ms seem to form two clusters. This effect is also present in AIP, where trajectories deflect into two distinct groups in a similar fashion (S2 Fig.).

To quantify clustering at the population level, we calculated the Euclidean distance between all pairs of delay lengths for each grip type separately in the space of all latent dimensions (Fig. 8b) and looked for clusters in the distance matrices without assuming clustering a priori (Materials and Methods). Two clusters were identified for the example

2.3. Probing the continuum of immediate to withheld grasping movements

data set (Fig. 8c), showing a split around the 400-500 ms delay point that lasts until shortly before movement onset (permutation test, $p < 0.01$; Materials and Methods). This pattern was very similar over all data sets (Fig. 8d, S2 Fig.), did not differ between grip types, and was present in both areas and monkeys, indicating that the state change that occurs between short and long delays spans both the frontal and parietal lobes.

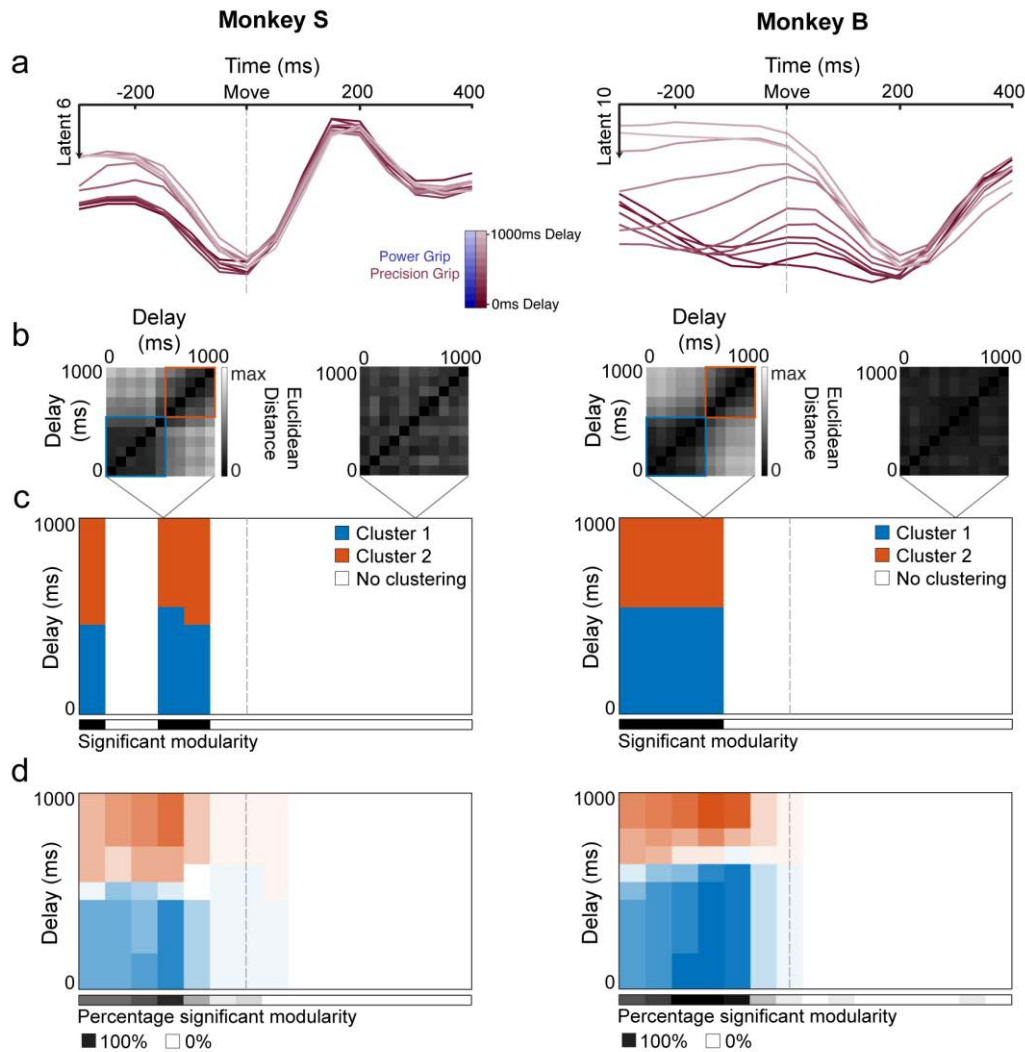


Fig 8. Clustering of movement initiation activity in F5. (a) Example latent projection population activity in F5 over all linearly spaced delays (0-1000 ms) for precision grip trials for an example data set from each monkey (sessions S4, B2), aligned to movement onset. (b) Euclidean distance between all pairs of delays in the full latent space for two example time points of the example data set including identified clustering using a clustering analysis that finds community structure (Materials and Methods). (c) Clusters identified in the distance matrices over time (in steps of 50 ms) for the example data set. Black significance bar shows time points where the modularity statistic exceeded chance level (permutation test, $p < 0.01$). (d) Same analysis as (c) averaged over all data sets and grip types (6 data sets x 2 grip types).

Clustering is not likely due to different movement kinematics, since the movement times were nearly identical for all delay lengths (Fig. 1d), especially for monkey S. However, since the time of movement onset is determined by the monkey's behavior, the time that

2.3. Probing the continuum of immediate to withheld grasping movements

has elapsed since the visual grip cue was presented could introduce a potential confound. Yet, 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 behavioral data, i.e. the mean time 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.

Discussion

To systematically probe the interplay between planning and movement in the grasping network, we recorded 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). Firstly, the initial part (~300 ms) of the neural space traversed was the same for all delays, but was grip specific, providing evidence that this activity was an unavoidable part of preparing the correct movement. Next, population activity shifted into a separate state that was not achieved during short delays. The memory state was more dynamic in F5 than in AIP, tracking subjective movement anticipation over time. Lastly, activity during movement initiation formed two distinct clusters that were eliminated after movement onset, demonstrating a network-wide shift when movements need to be withheld. Our findings reinforce the notion that more global aspects of movements, such as the movement plan, as well as dynamic aspects, such as cue anticipation, can be well extracted at the population level.

As shown in Fig. 4, separation between the neural trajectories occurred more than 200 ms earlier between the two grips than between long and short delays. This novel result indicates that while grip information is swiftly encoded in F5 and AIP following the cue, responses to the go cue are delayed at least 200 ms relative to the grip information in order to facilitate the completion of the motor plan, after which areas of the state space traversed by longer delays are not strictly necessary to produce successful movements, similar to the results of Ames et al. (2014) in dorsal premotor cortex (PMd).

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 differs to the results of Ames et al. (2014) in nearby PMd, who postulated that delay period activity may act as an attractor state into which all trials would congregate

2.3. Probing the continuum of immediate to withheld grasping movements

given enough preparation time. It is possible that PMd activity would be more dynamic if an experimental design with a memory period were utilized, 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). However, given current evidence our results support the notion that strongly dynamic memory period activity is a unique feature of F5.

It could be that the temporal dynamics during the memory period are a result of an internalized representation of the likelihood of task events occurring at specific times throughout the memory period, known as hazard rate and previously observed in the lateral intraparietal cortex (LIP) (Leon and Shadlen 2003; Janssen and Shadlen 2005). We observed significant fits of latent dimensions to the subjective anticipation rate across both areas, although F5 contributed significantly more to this activity. Furthermore, activity in these dimensions was predictive of reaction time, supporting the role of this activity in increasing or decreasing sensitivity to an external stimulus.

Time dependence has been identified in prefrontal areas (Genovesio et al. 2006), and 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). 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. The strongest evidence for this view comes from a recent study in which the presence or absence of a sensory stimulus on a given trial had to be reported (Carnevale et al. 2015). The authors found that the neural state space of premotor cortex evolved over the course of the trial and was more sensitive to incoming sensory information during the fixed window that the monkeys knew would or would not contain the stimulus. Importantly, Carnevale et al. (2015) showed that a recurrent neural network model trained for optimal response sensitivity well explained the behavior of the monkey. 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), suggesting that F5 is able to track the course of time internally and use this information to predict when an action is likely to be required. Furthermore, even though activity continues to change throughout memory, a stable representation of the desired action remains at the population level (Druckmann and Chklovskii 2012), consistent with the constant separation between grip types observed in some latent dimensions (Fig. 3).

2.3. Probing the continuum of immediate to withheld grasping movements

One of the most striking features in both areas, but especially in 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 (Fig. 7), 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 directly be explained by a rotational dynamical system (Churchland, Cunningham, et al. 2010; Churchland et al. 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. 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 (Sussillo et al. 2015; Michaels et al. 2016). Under this framework, selecting between multiple movement plans would only require the neural population to be within a general region of activity. Such a framework is also in line with the finding that preparatory activity in PMd/M1 projects into the null-space of upper limb muscles and transitions into the potent-space during movement (Kaufman et al. 2014), as this transition likely takes place during movement initiation when variability between movement plans is heavily reduced (Elsayed et al. 2016). Once movement is initiated, activity would fall onto a common trajectory unique to each action plan and rotational dynamics could proceed. Future work must tackle the question of to what degree local circuit features or extrinsic inputs can account for the rapid decrease in trial-to-trial variability taking place before movement execution.

While variability decreased leading up to movement onset, trajectories clustered into two distinct groups splitting between delay conditions less than or greater than 400-500 ms (Fig. 8). Given that full preparation likely takes ~400 ms, evidenced by the leveling of the RT curve after ~400 ms (Fig. 1d), the two clusters could correspond to movements executed 'as fast as possible' and movements executed from memory where the monkey must first wait for the go signal. Our results indicate that shifting between immediate movements and

2.3. Probing the continuum of immediate to withheld grasping movements

withheld movements from memory may cause a state shift in the fronto-parietal network that produces the two clusters during movement initiation. Once the state has been changed, the trajectories continue to cluster for the entirety of movement initiation (up to movement onset). Specifically, the underlying cause of the shift is likely the transition from reactive to proactive control, i.e., the increased ability to properly anticipate a go cue after sufficient preparation times (Braver 2012). This sensitivity to task timing is inherent in highly trained tasks, and has been shown in supplementary motor area (SMA; Chen et al. 2010) and medial frontal cortex (Stuphorn and Emeric 2012). 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).

It remains a possibility that systematic differences in hand-shaping latencies or final posture between different delay lengths could contribute to the observed clustering. However, clustering of delay conditions was almost non-existent after movement onset, especially in F5, making differences in final posture improbable. Although differences in hand-shaping during movement cannot be ruled out, the extreme similarity in movement times between delays (Results), especially for monkey S, make this possibility unlikely.

Given that the current task also involved a large reaching component, reach planning is likely a significant part of the observed activity. Still, the presence of grip type tuning in all epochs (Results), 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), indicates that F5 encodes grasping quite independently of reaching. Furthermore, reversibly inactivating F5 (Fogassi et al. 2001) or AIP (Gallese et al. 1994) selectively impairs hand-shaping and not reaching, providing evidence that our results are an accurate representation of the grasping network.

In summary, our results provide novel insights building on delayed reaching and grasping literature in premotor (Cisek et al. 2003; Lucchetti et al. 2005; Fluet et al. 2010) and parietal cortex (Murata et al. 1996; Snyder et al. 2006; Baumann et al. 2009). We show that dissociation of global and dynamic aspects of movement, such as the movement plan and the anticipation over time, respectively, can be coherently extracted at the level of neural populations and allow for comparison and dissociation between interacting cortical areas.

2.3. Probing the continuum of immediate to withheld grasping movements

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; Dann et al. 2016). 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, preventing any premature movement of the hands. 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

2.3. Probing the continuum of immediate to withheld grasping movements

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, additionally ensuring that monkeys did not prematurely move their hands or arms.

Task Design

The trial course of the delayed grasping task is shown in Fig. 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 150 ms 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, in discrete memory period bins of 0, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000, or 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). 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. A failed trial occurred if the monkeys stopped fixating the central point before movement onset, moved their hand from the hand rest sensor, performed the incorrect grip, or took longer than 1100 ms to complete the movement following the go cue. 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, apart from the 300 ms handle illumination period, in total darkness.

2.3. Probing the continuum of immediate to withheld grasping movements

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; 0.3-7500Hz hardware filter; see Supplementary Methods). 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 (Musial et al. 2002; Dann et al. 2016). 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. 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.

2.3. Probing the continuum of immediate to withheld grasping movements

Data preprocessing

Although units were classified as single- or multi-units, all recorded units were used for all 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. For individual unit plotting (Fig. 2), spike trains were smoothed with a Gaussian window ($\sigma = 50$ ms), but for all analyses spike trains were further reduced to a set of latent dimensions (see next section). 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 button. The cue alignment proceeded from 200 ms before cue onset until the go cue, and the movement onset alignment from movement onset minus the median reaction time for each delay condition until 400 ms after movement onset. These two alignments were combined to produce a continuous signal. 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.

Dimensionality reduction

In order to extract a set of smooth single-trial neural trajectories in our neural populations we applied Gaussian Process Factor Analysis (GPFA; Yu et al. 2009) to all neurons of both areas over all successful trials from 200 ms before cue onset to 400 ms after movement onset for each recording session separately. Performing a single dimensionality reduction over both areas allows a direct comparison of each area's contribution to the common signals. Units within each session were recorded simultaneously across both areas. GPFA is similar to factor analysis in that it finds an explanatory set of orthogonal dimensions based on the covariance structure between units that is a linear combination of binned neural data. However, in GPFA, each dimension de-noises data with a Gaussian smoothing kernel of unique width learned from the data. For our GPFA analysis, neural spiking data on single trials were binned into 50 ms bins and square-rooted before being transformed through linear combination into 10 latent dimensions. Units with an average firing rate less than 1 Hz were discarded before the analysis. These 10 dimensions, each based on an individual smoothing kernel, were further orthonormalized to produce a set of 10 orthogonal dimensions, each containing a combination of all smoothing kernels. Cross-validation procedures were undertaken to determine the optimal number of latent dimensions (Yu et al. 2009). Beyond 10 latent dimensions very little shared variance was explained by further

2.3. Probing the continuum of immediate to withheld grasping movements

addition of dimensions (<3% per dimension), and visualization of these dimensions showed almost no modulation. Since GPFA was carried out across both recorded areas simultaneously, the neural data of each area were separately transformed into the previously found latent dimensions to identify the specific contribution of each area to each latent dimension. For most analyses the extracted single trials were then cut into two alignments (previous section) and averaged over all trials of the same condition. In general at the boundary of alignments the signals matched very well to each other, showing almost no jumps in activity.

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 space of the 10 latent dimensions extracted through GPFA separately for each area. Three versions of this analysis were performed. For the distance in Fig. 4a, we iterated through all time points on delayed trajectory (in steps of 50 ms) and calculated the Euclidean point-to-curve distance from the delayed (1000 ms) trajectory to the non-delayed (0 ms) trajectory, where the point-to-curve distance is the minimum distance from a specific time 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 a similar point in neural space at some point in time, while large distances indicate that the two trajectories do not pass through a similar point in the high dimensional space. Euclidian distances were normalized by the square root of the number of neurons in order to make spaces with different number of neurons comparable.

For the distance analysis in Fig. 4c, GPFA was recalculated on a smaller portion of the data (200 ms before cue onset to 800 ms after) with a shorted bin width of 20 ms. Distance was then calculated as before between the delayed and non-delayed trajectories. In addition, to determine when grip information becomes present in the population, distance between the delayed trajectories (1000 ms) of each grip type was calculated in the same manner.

2.3. Probing the continuum of immediate to withheld grasping movements

For the distance analysis in Fig. 5, the Euclidean distance was calculated between all pairs of time points on the same trajectory (no-movement) 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 was in general the same, but with slight variations for the different distance analyses presented above. 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 average firing rates for each condition 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.

To determine when the observed distance distribution was significantly greater than the self-sampled distribution, we used a cluster-based permutation test (CBPT; Maris and Oostenveld 2007). Briefly, we used a modification of the original test that evaluates the area under the receiver operator characteristic curve (AUC) between the distance distribution and the self-sampled distribution over all time points and extracts clusters (consecutive time segments) of activity whose AUC exceeds a predefined threshold ($\alpha = 0.1$), then the absolute AUCs within each cluster were summed to produce cluster-level statistics. To generate a chance-level distribution from which the cluster-level statistics could be calculated, trials were randomly partitioned between the two conditions and the AUC and clustering redone (1000 partitions). From every partition the largest cluster was used to generate a largest chance cluster distribution. Cluster-level statistics were calculated by comparing the real cluster-levels against the largest chance cluster distribution. Real clusters were considered significant if they exceeded 95% of all largest chance cluster values corresponding to a $p = 0.05$. In this way, sensitivity to short or small time-scale differences is greatly reduced, but

2.3. Probing the continuum of immediate to withheld grasping movements

the overall false-alarm rate across time points remains below the designated 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 Fig. 5, resampling of trials was carried out 10000 times, with replacement, for each condition and data set. For each of the 10000 resampling steps the same trajectory was resampled twice, termed \mathbf{p} and \mathbf{p}' . Then, for every pair of time points (t_1 and t_2), the resampled distance along the first trajectory $d = d(\mathbf{p}(t_1), \mathbf{p}(t_2))$ was compared to the two inter-trajectory distances at time t_1 and t_2 : $d_1 = d(\mathbf{p}(t_1), \mathbf{p}'(t_1))$ and $d_2 = d(\mathbf{p}(t_2), \mathbf{p}'(t_2))$. We determined the percentile of resamples (across all 10000) for which the along-trajectory distance d exceeded both inter-trajectory distances: $d > \max(d_1, d_2)$. This percentile determined a specific p -value for each time pair (t_1, t_2). The resampled distance, d , was then considered significant if $p < 0.01$. In this way, the significance level was dependent on which time points were compared along the trajectory, establishing a conservative estimate of the underlying trial-to-trial variability.

Hazard rate

To classify the temporal evolution of activity during the memory period, the mean firing rate of each latent dimension for the no-movement condition from cue onset until reward onset was fit with an anticipation function, which can be described as the conditional probability that a movement will be required at a given moment, given that it has not occurred until this point. This type of anticipation has been termed the hazard rate, and we present it here precisely as in Janssen and Shadlen (Janssen and Shadlen 2005). The hazard rate can be expressed as

$$h(t) = \frac{f(t)}{1 - F(t)} \quad (13)$$

where $f(t)$ is the probability that a go cue will come at a given time after cue onset, and $F(t)$ is the cumulative distribution, $\int_{s=0}^t f(s) ds$.

As in Janssen and Shadlen (2005), to obtain an estimate of the monkey's internal representation of anticipation we calculate 'subjective anticipation' based on the assumption that the animal is uncertain about time and that this uncertainty scales with time since an event. Therefore, before calculating hazard rate we smoothed our probability

2.3. Probing the continuum of immediate to withheld grasping movements

density function, $f(t)$, with a normal distribution where standard deviation is proportional to elapsed time.

$$\tilde{f}(t) = \frac{1}{\phi t \sqrt{2\pi}} \int_{-\infty}^{\infty} f(\tau) e^{\frac{-(\tau-t)^2}{2\phi^2 t^2}} d\tau \quad (14)$$

The coefficient of variation, ϕ , is a Weber fraction under the assumption that the experience of elapsed time carries uncertainty that is proportional to the true duration (Weber's Law).

For all analyses we used a value of 0.26, as has been calculated from behavioral experiments and used previously (Leon and Shadlen 2003; Janssen and Shadlen 2005). To obtain the final subjective anticipation function, $\tilde{f}(t)$ was then substituted into Eq. 1, along with its cumulative distribution, $\tilde{F}(t)$.

$$r(t) = w_1 + w_2 \tilde{h}(t - w_3) \quad (15)$$

All fitting procedures were performed by fitting Eq. 3 to the average activity of each latent dimension over both areas, where w are constant terms obtained during the fitting procedure (Matlab function: *fit*), and \tilde{h} is Eq. 2 substituted into Eq. 1.

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. (2010). 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), pooled over all delays, but separately for each grip type. 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 all trials of each grip type 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.

2.3. Probing the continuum of immediate to withheld grasping movements

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, spikes 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 50 ms) in the 10 latent dimensions determined by GPFA 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 0.75. Each distance matrix was normalized to the maximum value over all time and subtracted from a matrix of ones in order to prepare them for analysis. 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 index Q between the empirical and permuted data. The percentile of instances where the permuted distribution values exceeded the empirical value corresponds to the p -value.

2.3. Probing the continuum of immediate to withheld grasping movements

2.3. Probing the continuum of immediate to withheld grasping movements

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.
- 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.
- Braver TS. 2012. The variable nature of cognitive control: a dual mechanisms framework. *Trends in Cognitive Sciences*. 16:106–113.
- Buonomano DV, Laje R. 2010. Population clocks: motor timing with neural dynamics. *Trends in Cognitive Sciences*. 14:520–527.
- Carnevale F, de Lafuente V, Romo R, Barak O, Parga N. 2015. Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty. *Neuron*. 86:1067–1077.
- Chen X, Scangos KW, Stuphorn V. 2010. Supplementary motor area exerts proactive and reactive control of arm movements. *J Neurosci*. 30:14657–14675.
- 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. 2010. 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, 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.
- Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV. 2006. Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci*. 26:3697–3712.
- 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.

2.3. Probing the continuum of immediate to withheld grasping movements

- 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.
- 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.*
- 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.
- Druckmann S, Chklovskii DB. 2012. Neuronal circuits underlying persistent representations despite time varying activity. *Curr Biol.* 22:2095–2103.
- Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP. 2016. Reorganization between preparatory and movement population responses in motor cortex. *Nat Commun.* 7:13239.
- 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.
- 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.
- Genovesio A, Tsujimoto S, Wise SP. 2006. Neuronal activity related to elapsed time in prefrontal cortex. *J Neurophysiol.* 95:3281–3285.
- 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.
- 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

2.3. Probing the continuum of immediate to withheld grasping movements

judgments. *eLife*.

- 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.
- 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.
- 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.
- Janssen P, Shadlen MN. 2005. A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci.* 8:234–241.
- 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, Seely JS, Sussillo D, Ryu SI, Shenoy KV, Churchland MM. 2016. The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eNeuro.* 3.
- 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.
- 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.
- 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.
- 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.
- Maris E, Oostenveld R. 2007. Nonparametric statistical testing of EEG- and MEG-data. *J*

2.3. Probing the continuum of immediate to withheld grasping movements

- Neurosci Methods. 164:177–190.
- 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.
- 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.
- Murata A, Gallese V, Kaseda M, Sakata H. 1996. Parietal neurons related to memory-guided hand manipulation. *J Neurophysiol.* 75:2180–2186.
- 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.
- 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.
- 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.
- Rosenbaum DA. 1980. Human movement initiation: specification of arm, direction, and extent. *J Exp Psychol Gen.* 109:444–474.
- 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.
- Schaffelhofer S, Scherberger H. 2016. Object vision to hand action in macaque parietal,

2.3. Probing the continuum of immediate to withheld grasping movements

premotor, and motor cortices. *eLife*. 5:6436.

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

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.

Stuphorn V, Emeric EE. 2012. Proactive and reactive control by the medial frontal cortex. *Front Neuroeng*. 5:9.

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.

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

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

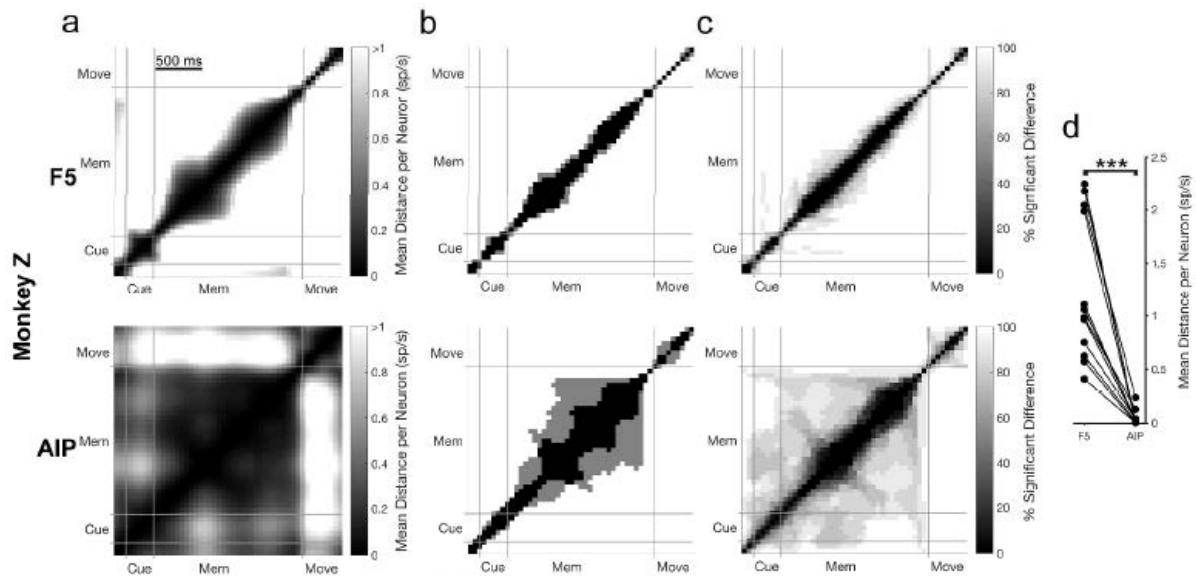
Wise SP, Kurata K. 1989. Set-Related Activity in the Premotor Cortex of Rhesus Monkeys: Effect of Triggering Cues and Relatively Long Delay Intervals. *Somatosens Mot Res*. 6:455–476.

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.

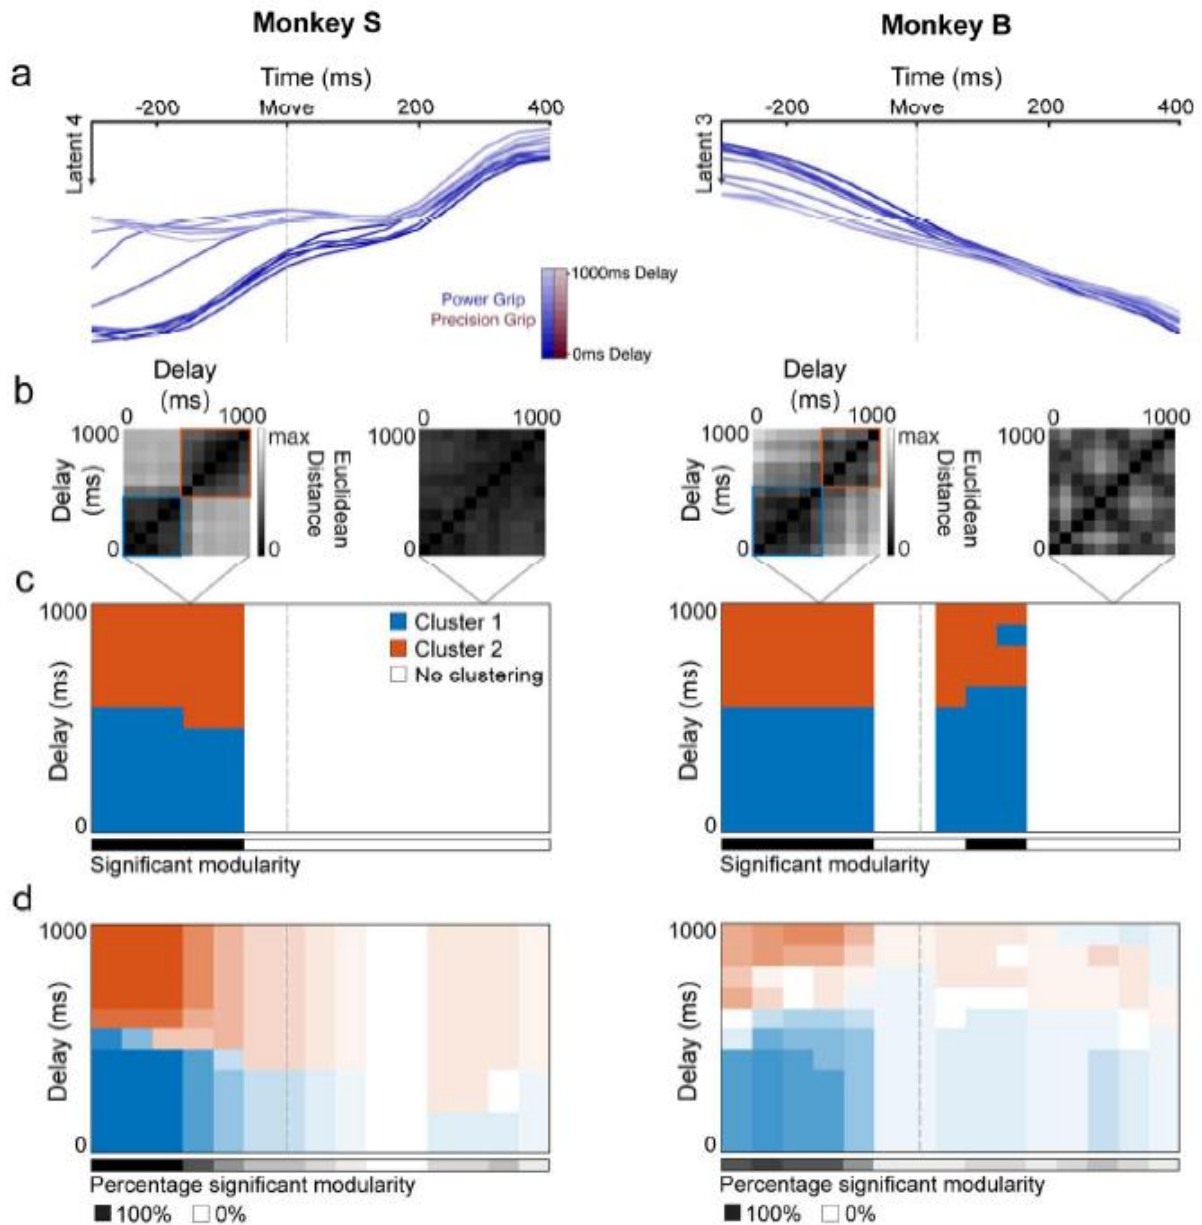
2.3. Probing the continuum of immediate to withheld grasping movements

Supplementary Figures and Tables



S1 Fig. Neural trajectory stability over the course of instructed trials for an additional experiment. Same layout as Fig 5. (a) Mean Euclidean distance in the latent space for the Instructed trials between all pairs of time points over both grip types for an example data set in monkey Z. For each pair of time points, distance results were tested for a significant difference using a bootstrapping procedure (10000 resamples in steps of 50 ms, $p = 0.01$). The abbreviations Cue, Mem, and Move, correspond to the cue, memory, and movement epochs, respectively. All plots are clipped at 1 sp/s for visualization. The times where a significant difference was found are shown in (b). (c) Percentage of time points showing a significant difference over all data sets and grip types (6 data sets x 2 grip types). (d) Mean distance over the stable portion of the memory period (600 ms after cue onset – go cue) for all individual data sets and grip types (6 data sets x 2 grip types) across areas and paired according to recording session. Stars indicate a significant difference (Wilcoxon signrank test, $p < 0.001$). 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 in F5 is maintained, suggesting that this variability is not due to the varying chance of subsequent movement, but represents features of the examined areas.

2.3. Probing the continuum of immediate to withheld grasping movements



S2 Fig. Clustering of movement initiation activity in AIP. (a) Example latent projection population activity in AIP over all linearly spaced delays (0-1000 ms) for precision grip trials for an example data set from each monkey (S3, B4), aligned to movement onset. (b) Euclidean distance between all pairs of delays in the full latent space for two example time points of the example data set including identified clustering using a clustering analysis that finds community structure (Materials and Methods). (c) Clusters identified in the distance matrices over time (in steps of 50 ms) for the example data set. Black significance bar shows time points where the modularity statistic exceeded chance level (permutation test, $p < 0.01$). (d) Same analysis as (c) averaged over all data sets and grip types (6 data sets x 2 grip types).

2.3. Probing the continuum of immediate to withheld grasping movements

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

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

3. General Discussion

In this dissertation the encoding, transformation, and coordination of information across the fronto-parietal grasping network was investigated while monkeys performed two different tasks. In the first task monkeys were either instructed or free to choose to grasp a target in two different ways, allowing for an investigation of internal decision making. In the second task monkeys performed the transition of immediate and delayed grasp movements, allowing for a detail investigation of this transition. In order to analyse the exact nature of the neuronal process within and across the fronto-parietal network including area AIP and F5 (and in chapter 2.1 also M1) large populations of neurons were recorded in parallel across all areas. Especially the possible to analyses the simultaneous activity of this area-spanning neuronal population gave new insights into the encoding, transformation and coordination of the behavioural relevant information within the network. In the following paragraph the results are summarized in detail.

3.1. Summary

In chapter 2.1 it was analyzed how the information flow is coordinated across the fronto-parietal single neuron network. Large numbers of single neurons were recorded in parallel across AIP, F5, and M1 while monkeys performed a delayed grasping task and the functional connectivity between all pairs of neurons was calculated based on cross-correlation histograms. To achieve a reliable estimate of the functional network connectivity, a new statistical procedure that corrected for multiple comparisons across different temporal delays and neuronal pairings was developed. This procedure allowed us to analyze the form of synchronization together with the functional network topology. The functional fronto-parietal single neuron network was nowhere near randomly organized, but appeared as a complex network, with a modular and small-world topology. Interestingly, the centrality distributions of all datasets were highly heterogeneous based on degree centrality as well as betweenness centrality, which could not be explained by distance-dependent connectivity. This indicated that functional hub neurons likely coordinated the network activity. The hub neurons were equally distributed across all three areas and strongly interconnected, forming an area-spanning coordinative rich-club. Surprisingly, when we analyzed the form of

3. General Discussion

synchronization, neurons were either synchronized by oscillatory synchrony in the beta-band, in the low-frequency range, or synchronized in a non-oscillatory manner. Intriguingly, the hub neurons forming a rich-club were oscillatory synchronized nearly without exception, while large parts of the rest of the network were non-oscillatory synchronized. When we analyzed the rhythmicity of the spiking of hub neurons, they were nearly exclusively rhythmically active in the beta- or low-frequency band, defining them as oscillators. Thus, the findings of this study suggest that the information flow of the fronto-parietal grasping network is coordinated by an area-spanning oscillatory-synchronized rich-club.

In chapter 2.2 it was investigated how information is encoded and transformed in the fronto-parietal grasping network while monkeys were either visually instructed or freely choosing to grasp a handle with one of two grip types. When analyzing the neuronal population from the classical representational view, describing activity of individual neurons as a function of various parameters, a large number of neurons were significantly tuned in AIP and F5 of the fronto-parietal grasping network and during all time points of the task. However, tuning changed dynamically over time and tuning parameters were uniformly distributed across the population; both findings were at odds with the classical representational view. In contrast, when considering the whole neuronal population as one strongly interconnected network, in which neural population activity evolves dynamically through space-space over time and conditions as suggested by the dynamical system perspective, a clear low dimensional structure became apparent. All task specific single trial activity could be explained by an evolution through just three independent informational subspaces representing visual, preparatory, and movement activity. Interestingly, for free-choice trials, where no specific visual information was given, all task specific activity during the decision process was explained by the preparatory space, suggesting that decision related activity and preparatory activity were the same for this task. Furthermore, changes of mind, e.g. when enforced by a later given second visual instruction, were clearly visible in the preparatory space. Crucially, contributions to all three informational spaces were randomly distributed across neurons with no significant category structure. A regularized recurrent neuronal network trained to produce muscle activity for the two grasps could accurately reproduce the neuronal dynamics both at the single unit and the population level. These results indicate that instead of addressing the attributes of individual neurons, neuronal activity can be more completely understood at the population level, where a

3. General Discussion

neuronal population can encode different processes at different and overlapping times. These processes can be dynamically transformed according to the behavioral demands, including free choices.

In chapter 2.3 the neuronal population encoding in AIP and F5 of the transition between immediate and withheld movement was examined. Single neuron responses of both areas were complex and difficult to characterise from the representation view. However, when considered on the population-level and visualized by dimensionality reduction techniques, a clearly describable temporal and conditional population dynamics became apparent. Neuronal population dynamics of both areas first followed a grip specific defined trajectory indistinguishable for immediate up to long delayed grasps. These trajectories properly represented unavoidable processing from visual to preparatory information. However, after this initial phase, population activity in AIP tended to stabilize, whereas activity in F5 continued to evolve through state space, likely reflecting movement anticipation. Interestingly, population activity of both areas evolved through two distinct and significantly separate spaces for immediate movements and withhold delayed movements, suggesting a unique state for movements performed from memory. However, trajectories for the different grasp movements were maintained in separate spaces. These findings suggest that the complex interplay of dynamical and static aspects of movement preparation, such as anticipation and planning of a particular grasp type, can be understood as an evolution of neuronal population activity through specific dimensions of a higher dimensional state space.

In the work presented in Appendix A we evaluated how representational models based on single neuron characterizations, and dynamical system models based on the neuronal population activity describing the generation of reach movements in PMd and M1, can be integrated and better tested for their validity. This study builds upon the results of Churchland M. et al. 2012 showing that population dynamics during reach movements can be described by a dynamical system model, with the preparatory state serving as an initial state of a rotation dynamic. However, by simulating simple velocity-tuned neurons for a center-out reaching task and incorporating variable latencies between kinematics and individual neuronal activities, rotational dynamics appeared on the population level. Yet, meaningful rotational dynamics should depend on the conditional population structure, while this should be irrelevant for representational models. To distinguish between these

3. General Discussion

two possibilities, we developed a covariance-matched permutation test (CMPT) that reassigned neural data between task conditions independently for each neuron while maintaining overall neuron-to-neuron relationships. While the rotations of representational models of neuronal activity did not depend on the conditional structure, they did strongly depend on the conditional structure for recorded data as well as a RNN trained to produce kinematics. These findings speak in favour of the dynamical systems perspective in describing motor cortex population dynamics. Interestingly, directional tuning was an emergent property of our RNN model simply as a consequence of the generated output parameters. Yet, the directional tuning was found to change over time and neuronal tuning was often only roughly matched by a cosine tuning function, similar to recorded neurons. These observations suggest that, even if representational models can describe single neuron data to a certain extent, their results can nonetheless be misleading, and the neuronal population dynamics can potentially be better explained by a dynamical system model.

Finally, in the study described in Appendix B we showed that the reaction time to initiate a grasp movement could be predicted from the activity of large numbers of simultaneously recorded neurons in AIP and F5. Single-trial preparatory activity of both areas was predictive of reaction time, although results differed strongly based on the method of analysis used. Population-based methods for predicting reaction time were found to give better and more reliable results than single neuron based predictions for both areas. Interestingly, in comparing different population-based methods, those which were not based on the assumption that shorter reaction times are associated with higher firing rates performed much better. Furthermore, the predictive information was distributed across the whole population of neurons of both areas with no evidence for distinct subpopulations tuned to reaction time. However, neuronal populations of F5 were more predictive than populations of AIP, suggesting that F5 populations are more directly related to grasp initiation. These observations indicate that aspects of movement initiation are distributed across neuronal populations and even across different brain areas.

3.2. Outlook

A great deal of new insight into how ensembles of neurons generate emergent functional states has been obtained by considering the activity of large populations of neurons as one dynamical trajectory, which evolves in time within a low dimensional state space. This leads

3. General Discussion

to the generation of movements, resulting in a transition from the representational view of neuronal activity to a dynamical systems perspective (Shenoy et al., 2013; Yuste, 2015). On the other hand, there is also increasing evidence that, globally as well as locally, information flow is coordinated by oscillatory synchronization in distinct frequency bands (Buzsáki, 2010; Engel and Fries, 2010; Fries, 2015; Schomburg, 2015). The studies of this thesis emphasize the view that neuronal population activity of the fronto-parietal network can be best described as a dynamical process evolving through a limited number of subspaces. Each subspace represents different aspects such as visual, preparatory, and moment related information. Information about the anticipation of an upcoming event or the timing of movement initiation is also represented by these processes. Analyses performed on the same single neuron population, which revealed the different information subspaces, showed that the information flow within and across the fronto-parietal network was coordinated by a rich-club of oscillatory synchronized neurons. Yet, how these two findings are interrelated is currently unclear. A possible explanation is given by two groups of investigators who independently described a similar concept of information transformation on the population level (Womelsdorf et al., 2013; Elsayed et al., 2016). However, one comes from the field of neuronal state space analyses, while the other comes from the field of oscillatory synchrony analyses. The first study (Elsayed et al., 2016) suggests that the same neural population acts, at different times, as two separate circuits with very different properties spanning orthogonal but lawfully related subspaces. This relationship was shown for the transition from preparatory to movement related activity (Figure 1a).

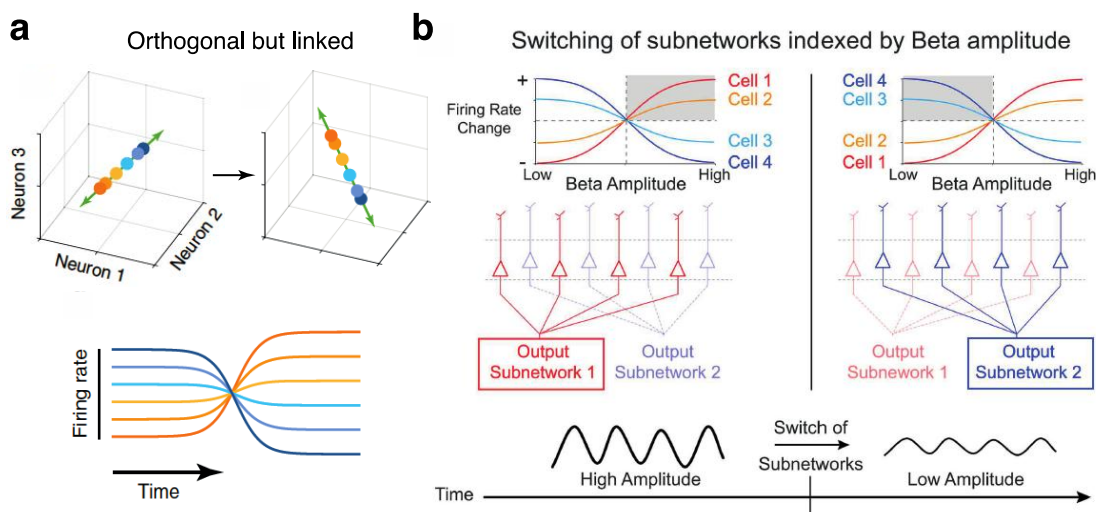


Figure 1 Information transformation on the population level (a) Activity of three hypothetical neurons for the transition from preparatory to movement related activity. Each axis represents the firing rate of one neuron and each dot represents the neural state for one of six conditions. The activity of the three neurons occupies a subspace

3. General Discussion

of the full state space for preparatory activity, which is orthogonal to the subspace occupied during movement related activity. However, the relationship between conditional responses is lawfully linked between the two states. The panel below is a one-dimensional illustration of the firing rate change at the transition between the two states. The colors correspond to the dots in the panels above and indicate the different condition identities. Adapted from Elsayed et al. (2016). **(b)** In the second model, the firing rates of single neurons are coupled in a sigmoidal relation to the amplitude of beta-band activity of the LFP. Some neurons fire stronger during high beta amplitudes (gray shading, left panel), while other neurons fire more weakly during high beta amplitudes (gray shading, right panel). These findings suggest that a high beta amplitude cortical state (left panels) activates a selected subnetwork of neurons, while a low beta amplitude cortical state activates another selected subnetwork of neurons (right panels). This was found for the transition from preparatory to movement related activity. Adapted from Womelsdorf (2013).

The coordinative mechanism introduced in the second article (Womelsdorf et al., 2013) is that subnetworks are selected by a change in beta rhythmic activity, serving as a true switch in the local network by causally modulating single neuron firing rates. This process is called cross-level coupling (Figure 1b). Surprisingly, this was shown as well for the transition from preparatory to movement related activity (Canolty et al., 2012). What if the described subnetworks are the orthogonal subspaces of the first study and a rich-club of coordinative beta synchronized neurons is causing the subspace transition by cross-level coupling?

4. Bibliography

- Achard S, Salvador R, Whitcher B, Suckling J, Bullmore E (2006) A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *J Neurosci* 26:63–72.
- 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.
- Ahrens MB, Keller PJ (2013) Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature Publishing Group*:1–12.
- Andersen RA, Cui H (2009) Intention, action planning, and decision making in parietal-frontal circuits. *Neuron* 63:568–583.
- Barabási A-L (2009) Scale-free networks: a decade and beyond. *Science* 325:412–413.
- Barabási A-L, Albert R, Jeong H (1999) Internet: Diameter of the World-Wide Web. *Nature* 401:130–131.
- Barabási A-L, Oltvai ZN (2004) Network biology: understanding the cell's functional organization. *Nat Rev Genet* 5:101–113.
- Barrese JC, Aceros J, Donoghue JP (2016) Scanning electron microscopy of chronically implanted intracortical microelectrode arrays in non-human primates. *J Neural Eng* 13:1–27.
- Barrese JC, Rao N, Paroo K, Triebwasser C, Vargas-Irwin C, Franquemont L, Donoghue JP (2013) Failure mode analysis of silicon-based intracortical microelectrode arrays in non-human primates. *J Neural Eng* 10:066014.
- Bastos AM, Schoffelen J-M (2016) A Tutorial Review of Functional Connectivity Analysis Methods and Their Interpretational Pitfalls. *Front Syst Neurosci* 9:413.
- Bastos AM, Vezoli J, Bosman CA, Schoffelen J-M, Oostenveld R, Dowdall JR, De Weerd P, Kennedy H, Fries P (2015) Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron* 85:390–401.
- Baumann MA, Fluet MC, Scherberger H (2009) Context-Specific Grasp Movement Representation in the Macaque Anterior Intraparietal Area. *J Neurosci* 29:6436–6448.
- Berger D, Warren D, Normann R, Arieli A, Grün S (2007) Spatially organized spike correlation in cat visual cortex. *Neurocomputing* 70:2112–2116.
- Bonifazi P, Goldin M, Picardo MA, Jorquera I, Cattani A, Bianconi G, Represa A, Ben-Ari Y, Cossart R (2009) GABAergic Hub Neurons Orchestrate Synchrony in Developing Hippocampal Networks. *Science* 326:1419–1424.
- Bosman CA, Schoffelen J-M, Brunet N, Oostenveld R, Bastos AM, Womelsdorf T, Rubehn B, Stieglitz T, De Weerd P, Fries P (2012) Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* 75:875–888.

4. Bibliography

- Brovelli AA, Ding MM, Ledberg AA, Chen YY, Nakamura RR, Bressler SLS (2004) Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by Granger causality. *Proc Natl Acad Sci USA* 101:9849–9854.
- Bullmore E, Sporns O (2009) Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* 10:186–198.
- Buzsáki G (2004) Large-scale recording of neuronal ensembles. *Nat Neurosci* 7:446–451.
- Buzsáki G (2010) Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron* 68:362–385.
- Buzsáki G, Mizuseki K (2014) The log-dynamic brain: howskewed distributions affect networkoperations. *Nat Rev Neurosci* 15:264–278.
- Buzsáki G, Wang X-J (2012) Mechanisms of Gamma Oscillations. *Annu Rev Neurosci* 35:203–225.
- Buzsáki GG, Anastassiou CAC, Koch CC (2011) The origin of extracellular fields and currents - EEG, ECoG, LFP and spikes. *Nat Rev Neurosci* 13:407–420.
- Canolty RT, Ganguly K, Carmena JM (2012) Task-Dependent Changes in Cross-Level Coupling between Single Neurons and Oscillatory Activity in Multiscale Networks Sporns O, ed. *PLoS Comput Biol* 8:e1002809–e1002823.
- 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 (2010) Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* 68:387–400.
- Churchland MM, Shenoy KV (2007) Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *J Neurophysiol* 97:4235–4257.
- Cisek P (2012) Making decisions through a distributed consensus. *Curr Opin Neurobiol* 22:927–936.
- Cisek P, Kalaska JF (2005) Neural Correlates of Reaching Decisions in Dorsal Premotor Cortex: Specification of Multiple Direction Choices and Final Selection of Action. *Neuron* 45:801–814.
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nature Publishing Group* 14:811–819.
- Colizza V, Flammini A, Serrano MA, Vespignani A (2006) Detecting rich-club ordering in complex networks. *Nat Phys* 2:110–115.
- Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings. *Nat Neurosci* 17:1–10.

4. Bibliography

- Dean HL, Hagan MA, Pesaran B (2012) Only coherent spiking in posterior parietal cortex coordinates looking and reaching. *Neuron* 73:829–841.
- Dhamala M, Rangarajan G, Ding M (2008) Analyzing information flow in brain networks with nonparametric Granger causality. *Neuroimage* 41:354–362.
- Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS (2010) Decorrelated Neuronal Firing in Cortical Microcircuits. *Science* 327:584–587.
- Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP (2016) Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications*:1–15.
- Engel AK, Fries P (2010) Beta-band oscillations--signalling the status quo? *Curr Opin Neurobiol* 20:156–165.
- Engel AK, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 2:704–716.
- Fluet MC, 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.
- Freedman DJ, Assad JA (2006) Experience-dependent representation of visual categories in parietal cortex. *Nature* 443:85–88.
- Freedman DJ, Assad JA (2016) Neuronal Mechanisms of Visual Categorization: An Abstract View on Decision Making. *Annu Rev Neurosci* 39:129–147.
- Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291:312–316.
- Freeman LC (1977) A set of measures of centrality based on betweenness. *Sociometry*.
- Freiwald WA, Tsao DY (2010) Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science* 330:845–851.
- Freiwald WA, Tsao DY, Livingstone MS (2009) A face feature space in the macaque temporal lobe. *Nature Publishing Group* 12:1187–1196.
- Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences* 9:474–480.
- Fries P (2009) Neuronal Gamma-Band Synchronization as a Fundamental Process in Cortical Computation. *Annu Rev Neurosci* 32:209–224.
- Fries P (2015) Perspective. *Neuron* 88:220–235.

4. Bibliography

- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*.
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11:823–833.
- Gail A (2006) Neural Dynamics in Monkey Parietal Reach Region Reflect Context-Specific Sensorimotor Transformations. *J Neurosci* 26:9376–9384.
- Gallese V, Murata A, Kaseda M, Niki N, Sakata H (1994) Deficit of hand preshaping after muscimol injection in monkey parietal cortex. *Neuroreport*.
- Georgopoulos AP, Kalaska JF, Caminiti R (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of ...*
- Gerhard F, Pipa G, Lima B, Neuenschwander S, Gerstner W (2011) Extraction of Network Topology From Multi-Electrode Recordings: Is there a Small-World Effect? *Front Comput Neurosci* 5.
- Gold C (2006) On the Origin of the Extracellular Action Potential Waveform: A Modeling Study. *J Neurophysiol* 95:3113–3128.
- Gollo LL, Mirasso C, Sporns O, Breakspear M (2014) Mechanisms of Zero-Lag Synchronization in Cortical Motifs Gutkin BS, ed. *PLoS Comput Biol* 10:e1003548.
- Gray CM, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc Natl Acad Sci USA* 86:1698–1702.
- Gregoriou GG, Gotts SJ, Zhou H, Desimone R (2009) High-Frequency, Long-Range Coupling Between Prefrontal and Visual Cortex During Attention. *Science* 324:1207–1210.
- Haegens S, Nacher V, Hernandez A, Luna R, Jensen O, Romo R (2011) Beta oscillations in the monkey sensorimotor network reflect somatosensory decision making. *Proc Natl Acad Sci USA* 108:10708–10713.
- Harriger L, van den Heuvel MP, Sporns O (2012) Rich Club Organization of Macaque Cerebral Cortex and Its Role in Network Communication Kaiser M, ed. *PLoS ONE* 7:e46497.
- Harvey CD, Coen P, Tank DW (2012) Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484:62–68.
- Hebb DO (1949) *The organization of behavior*. New York.
- Hernandez A, Nacher V, Luna R, Zainos A, Lemus L, Alvarez M, Vazquez Y, Camarillo L, Romo R (2010) Decoding a Perceptual Decision Process across Cortex. *Neuron* 66:300–314.
- Hilgetag CC, Burns GA, O'Neill MA, Scannell JW, Young MP (2000) Anatomical connectivity defines the organization of clusters of cortical areas in the macaque monkey and the cat. *Philos Trans R Soc Lond, B, Biol Sci* 355:91–110.

4. Bibliography

- Honey CJ, Kötter R (2007) Identification and Classification of Hubs in Brain Networks. *PLoS ONE*.
- Honey CJ, Kötter R, Breakspear M, Sporns O (2007) Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proc Natl Acad Sci USA* 104:10240–10245.
- Hubel DH (1957) Tungsten microelectrode for recording from single units. *Science*.
- Hubel DH, Wiesel TN (1968) Receptive Fields and Functional Architecture of Monkey Striate Cortex.
- Janssen P, Scherberger H (2015) Visual Guidance in Control of Grasping. *Annu Rev Neurosci* 38:69–86.
- Jeong H, Tombor B, Albert R, Oltvai ZN, Barabási AL (2000) The large-scale organization of metabolic networks. *Nature* 407:651–654.
- 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.
- Kiani R, Shadlen MN (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* 324:759–764.
- Kim S, Putrino D, Ghosh S, Brown EN (2011) A Granger causality measure for point process models of ensemble neural spiking activity. *PLoS Comput Biol* 7:e1001110.
- Klaes C, Westendorff S, Chakrabarti S, Gail A (2011) Choosing Goals, Not Rules: Deciding among Rule-Based Action Plans. *Neuron* 70:13–13.
- Kobak D, Brendel W, Constantinidis C, Feierstein CE (2016) Demixed principal component analysis of neural population data. *Elife*.
- Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673.
- König P, Engel AK, Singer W (1995) Relation between oscillatory activity and long-range synchronization in cat visual cortex. *Proc Natl Acad Sci USA*.
- Kraskov A, Stögbauer H, Grassberger P (2004) Estimating mutual information. *Phys Rev E* 69:066138–16.
- la Rocha de J, Doiron B, Shea-Brown E, Josić K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. *Nature* 448:802–806.
- Lehmann SJ, Scherberger H (2013) Reach and Gaze Representations in Macaque Parietal and Premotor Grasp Areas. *J Neurosci* 33:7038–7049.
- Lindner M, Vicente R, Priesemann V, Wibral M (2011) TRENTOOL: a Matlab open source toolbox to analyse information flow in time series data with transfer entropy. *BMC Neuroscience* 12:119.

4. Bibliography

- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157.
- 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.
- Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. *Neuron* 73:1–19.
- Markov NT, Ercsey-Ravasz MM, Lamy C, Vezoli J, Falchier A, Quilodran R, Gariel MA, Gamanut R, Huissoud C, Clavagnier S, Dehay C, Toroczkai Z, Van Essen DC (2014) A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cerebral Cortex* 24:17–36.
- 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* 12:1–19.
- Miller EK (2000) The prefrontal cortex and cognitive control. *Nat Rev Neurosci* 1:59–66.
- 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.
- Musallam S, Bak MJ, Troyk PR, Andersen RA (2007) A floating metal microelectrode array for chronic implantation. *Journal of Neuroscience Methods* 160:122–127.
- NAcher V, Ledberg A, Deco G, Romo R (2013) Coherent delta-band oscillations between cortical areas correlate with decision making. *Proc Natl Acad Sci USA* 110:15085–15090.
- Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature* 341:52–54.
- Nicolelis MAL, Dimitrov D, Carmena JM, Crist R, Lehew G, Kralik JD, Wise SP (2003) Chronic, multisite, multielectrode recordings in macaque monkeys. *Proc Natl Acad Sci USA* 100:11041–11046.
- Nigam S, Shimono M, Ito S, Yeh FC, Timme N, Myroshnychenko M, Lapish CC, Tosi Z, Hottowy P, Smith WC, Masmanidis SC, Litke AM, Sporns O, Beggs JM (2016) Rich-Club Organization in Effective Connectivity among Cortical Neurons. *J Neurosci* 36:670–684.
- Novembre J, Stephens M (2008) Interpreting principal component analyses of spatial population genetic variation. *Nat Genet* 40:646–649.
- Okatan M, Wilson MA, Brown EN (2005) Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput* 17:1927–1961.
- Padoa-Schioppa C (2011) Neurobiology of Economic Choice: A Good-Based Model. *Annu Rev Neurosci* 34:333–359.

4. Bibliography

- Perin R, Berger TK, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. *Proc Natl Acad Sci USA* 108:5419–5424.
- Pesaran B (2010) Neural correlations, decisions, and actions. *Curr Opin Neurobiol* 20:166–171.
- Pesaran B, Nelson MJ, Andersen RA (2008) Free choice activates a decision circuit between frontal and parietal cortex. *Nature* 453:406–409.
- Pnevmatikakis EA, Soudry D, Gao Y, Machado TA, Merel J, Pfau D, Reardon T, Mu Y, Lacefield C, Yang W, Ahrens M, Bruno R, Jessell TM, Peterka DS, Yuste R, Paninski L (2016) Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. *Neuron*:1–16.
- Quinn CJC, Coleman TPT, Kiyavash NN, Hatsopoulos NGN (2011) Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. *J Comput Neurosci* 30:17–44.
- Quiroga RQ, Nadasdy Z, Ben-Shaul Y (2004) Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput* 16:1661–1687.
- Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435:1102–1107.
- Ramalingam N, McManus JNJ, Li W, Gilbert CD (2013) Top-down modulation of lateral interactions in visual cortex. *J Neurosci* 33:1773–1789.
- Raposo D, Kaufman MT, Churchland AK (2014) A category-free neural population supports evolving demands during decision-making. *Nature Publishing Group* 17:1784–1792.
- Ravasz E, Somera AL, Mongru DA, Oltvai ZN, Barabási AL (2002) Hierarchical organization of modularity in metabolic networks. *Science* 297:1551–1555.
- 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* 497:1–6.
- Rishel CA, Huang G, Freedman DJ (2013) Independent Category and Spatial Encoding in Parietal Cortex. *Neuron* 77:969–979.
- Roelfsema PR, Engel AK, König P, Singer W (1997) Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385:157–161.
- Romo R, Salinas E (2003) Flutter Discrimination: neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4:203–218.
- Rossant C, Kadir SN, Goodman DFM, Schulman J, Hunter MLD, Saleem AB, Grosmark A, Belluscio M, Denfield GH, Ecker AS, Tolias AS, Solomon S, Buzsáki G, Carandini M, Harris KD (2016) Spike sorting for large, dense electrode arrays. *Nat Neurosci* 19:634–641.
- Rousche PJ, Normann RA (1998) Chronic recording capability of the Utah Intracortical Electrode Array in cat sensory cortex. *Journal of Neuroscience Methods* 82:1–15.

4. Bibliography

- Salazar RF, Dotson NM, Bressler SL, Gray CM (2012) Content-specific fronto-parietal synchronization during visual working memory. *Science* 338:1097–1100.
- 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.
- Schaffelhofer S, Scherberger H (2016) Object vision to hand action in macaque parietal, premotor, and motor cortices. *Elife*.
- Scherberger H, Andersen RA (2007) Target Selection Signals for Arm Reaching in the Posterior Parietal Cortex. *J Neurosci* 27:2001–2012.
- Scherberger H, Jarvis MR, Andersen RA (2005) Cortical Local Field Potential Encodes Movement Intentions in the Posterior Parietal Cortex. *Neuron* 46:347–354.
- Schomburg EW (2015) What does gamma coherence tell us about inter-regional neural communication? *Nat Neurosci* 18:484–489.
- Schroeter MS, Charlesworth P, Kitzbichler MG, Paulsen O, Bullmore ET (2015) Emergence of Rich-Club Topology and Coordinated Dynamics in Development of Hippocampal Functional Networks In Vitro. *J Neurosci* 35:5459–5470.
- Schröter M, Paulsen O, Bullmore ET (2017) Micro-connectomics: probing the organization of neuronal networks at the cellular scale. *Nat Rev Neurosci*:1–16.
- Sejnowski TJ, Churchland PS, Movshon JA (2014) Putting big data to good use in neuroscience. *Nat Neurosci* 17:1440–1441.
- Seth AK (2010) A MATLAB toolbox for Granger causal connectivity analysis. *Journal of Neuroscience Methods* 186:262–273.
- Shadlen MN, Kiani R (2013) Perspective. *Neuron* 80:791–806.
- Shenoy KV, Sahani M, Churchland MM (2013) Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annu Rev Neurosci* 36:337–359.
- Shimono M, Beggs JM (2014) Functional Clusters, Hubs, and Communities in the Cortical Microconnectome. *Cerebral Cortex*.
- Siegel M, Buschman TJ, Miller EK (2015) Cortical information flow during flexible sensorimotor decisions. *Science* 348:1352–1355.
- Smith MA, Kohn A (2008) Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex. *J Neurosci* 28:12591–12603.
- Stephan KE, Hilgetag CC, Burns GA, O'Neill MA, Young MP, Kötter R (2000) Computational analysis of functional connectivity between areas of primate cerebral cortex. *Philos Trans R Soc Lond, B, Biol Sci* 355:111–126.

4. Bibliography

- Steriade M (2001) Impact of network activities on neuronal properties in corticothalamic systems. *J Neurophysiol* 86:1–39.
- Sugrue LP, Corrado GS, Newsome WT (2005) Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat Rev Neurosci* 6:363–375.
- Sussillo D (2014) ScienceDirectNeural circuits as computational dynamical systems. *Curr Opin Neurobiol* 25:156–163.
- Sussillo D, Churchland MM, Kaufman MT, Shenoy KV (2015) A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Publishing Group*:1–12.
- Takagaki K, Lippert MT, Dann B, Wanger T, Ohl FW (2008) Normalization of Voltage-Sensitive Dye Signal with Functional Activity Measures Mansvelder HD, ed. *PLoS ONE* 3:e4041–12.
- Todorov E (2000) Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat Neurosci* 3:391–398.
- Ts'o DY, Gilbert CD, Wiesel TN (1986) Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J Neurosci* 6:1160–1170.
- Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286:1943–1946.
- van den Heuvel MP, Kahn RS, Goñi J, Sporns O (2012) High-cost, high-capacity backbone for global brain communication. *Proc Natl Acad Sci USA* 109:11372–11377.
- van den Heuvel MP, Sporns O (2013) Network hubs in the human brain. *Trends in Cognitive Sciences* 17:683–696.
- Vicente R, Gollo LL, Mirasso CR, Fischer I, Pipa G (2008) Dynamical relaying can yield zero time lag neuronal synchrony despite long conduction delays. *Proc Natl Acad Sci USA* 105:17157–17162.
- Waldert S, Lemon RN, Kraskov A (2013) Influence of spiking activity on cortical local field potentials. *The Journal of Physiology* 591:5291–5303.
- Watts DJ, Strogatz SH (1998) Collective dynamics of “small-world” networks. *Nature* 393:440–442.
- Womelsdorf T, Everling S (2015) Long-Range Attention Networks: Circuit Motifs Underlying Endogenously Controlled Stimulus Selection. *Trends Neurosci* 38:682–700.
- Womelsdorf T, Schoffelen J-M, Oostenveld R, Singer W, Desimone R, Engel AK, Fries P (2007) Modulation of neuronal interactions through neuronal synchronization. *Science* 316:1609–1612.
- Womelsdorf T, Westendorff S, Ardid S (2013) Subnetwork selection in deep cortical layers is mediated by beta-oscillation dependent firing. *Front Syst Neurosci* 7:25.

4. Bibliography

Wong YT, Fabiszak MM, Novikov Y, Daw ND, Pesaran B (2016) Coherent neuronal ensembles are rapidly recruited when making a look-reach decision. *Nat Neurosci* 19:327–334.

Yamins DLK, DiCarlo JJ (2016) Using goal-driven deep learning models to understand sensory cortex. *Nat Neurosci* 19:356–365.

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.

Yu S, Huang D, Singer W, Nikolic D (2008) A Small World of Neuronal Synchrony. *Cerebral Cortex* 18:2891–2901.

Yuste R (2015) From the neuron doctrine to neural networks. *Nat Rev Neurosci*.

Curriculum Vitae

Benjamin Dann



Personal details

Address	Wilhelm-Weber-Str.27a, Göttingen, Germany
Email	Benjamin.dann@googlemail.com
Birthday	September 25 th 1982 in Darmstadt, Germany
Nationality	German
Family Status	Married with Tanja Strecker, one daughter: Julia Sophia Dann (August 11st 2016)
Current position	Graduate Student at the German Primate Center, Goettingen, Germany

Education

2011 -	Graduate Student, Neurobiology Lab of Hans Scherberger, German Primate Center (DPZ), Göttingen, German
2009 – 2011	Research Assistance, Neurobiology Lab of Hans Scherberger, German Primate Center (DPZ), Göttingen, German
2007 – 2009	Diploma Student, Diploma theses/ Master theses, Department of Neurophysiology of Wolf Singer, Max-Planck Institute for Brain Research, Frankfurt, Germany
2005 – 2007	Undergraduate, Hauptstudium/ Master classes in Neurobiology/Neuroscience, Otto-von-Guericke-University, Magdeburg, Germany

5. Curriculum Vitae

2006 – 2007	Research Assistant, Neuroprotheses Lab of Frank Ohl, Leibniz-Institute for Neurobiology (LIN), Magdeburg, Germany
2003-2005	Undergraduate, Grundstudium/ Bachelor classes in Biology, Technical University Darmstadt, Darmstadt, Germany
2002	Abitur, Eleonorenschool, Darmstadt, Germany

Research Interest

My main research focus is on system neuroscience, in particular multi-electrode array electrophysiological recordings of behaving animals. I am interested in population coding and functional connectivity to understand the information encoding and coordination across cortex. The analytical methods I am working with are: functional connectivity measures, parametric and nonparametric statistics including modern multiple comparison corrections, clustering algorithms, frequency analyses, functional connectivity measures in the frequency domain, dimensionality reduction techniques, and Graph theoretical analyses.

Skills

Languages

German (native language), English (fluent reading and writing)

Programming

Matlab, Lab View

Software

Words, Excel, Power Point, Adobe Illustrator

Mathematics

Signal processing, statistics, network analyses with graph theory, frequency analyses, dimensionality reduction methods, signal decoding

Work with animals

5. Curriculum Vitae

Monkey care, Rat and Gerbil care, Monkey training, Rat and Gerbil training, surgical procedures to implant electrodes into a monkey brain, surgical procedures to implant electrodes into a rat or gerbil brain

Supervised theses

Ph.D. Theses

2015- Swathi Sheshadri. Amplitude free reconfiguration of the beta and theta fronto-parietal network for different grasps and tasks (Georg-August-University Göttingen)

Master Theses

2017 Michael Lutz. Different role of broad and narrow spiking neurons in the front-parietal grasping network (??? University)

2014 - 2015 Yves Baetz, Fronto-parietal synchronization for hand grasping in macaque monkey (Georg-August-University Göttingen)

2009 - 2011 Alexandra Wellner. LFP activity in AIP and F5 encode grip type and object orientation (Westfälische Wilhelms-University Münster).

Bachelor Theses

2013 Carolina Focke. Spike-field coherence between and within F5 and AIP for different grasp conditions

2012 Yves Baetz. LFP activity for decision making and grasping (Georg-August-University Göttingen)

Long Internships

2011 - 2012 Steve Suway, Reward representation in AIP and F5 during decision-making in a grasping task (German Primate Center Göttingen)

Papers in peer-reviewed journals

Michaels JA, **Dann B**, Scherberger H (2016c) Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning Yu B, ed. PLoS Comput Biol 12:e1005175–22.

5. Curriculum Vitae

Dann B, Michaels JA, Schaffelhofer S, Scherberger H (2016b) Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates. *Elife* 5:2870.

Dann B, Michaels JA, Scherberger H (2016a) Separable decoding of cue, intention, and movement information from the fronto-parietal grasping-network. *Proceedings of the Sixth International Brain-Computer Interface Meeting*, pp 1–261.

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.

Takagaki K, Lippert MT, **Dann B**, Wanger T, Ohl FW (2008) Normalization of Voltage-Sensitive Dye Signal with Functional Activity Measures Mansvelder HD, ed. *PLoS ONE* 3:e4041–12.

In-progress publications

Michaels JA*, **Dann B***, Intveld RW, Scherberger H. *Equal contribution (in prep.). Probing the continuum of immediate to withheld grasping movements in the macaque fronto-parietal network.

Dann B*, Michaels JA*, Scherberger H. *Equal contribution (in prep.). Three information subspaces explain the category-free population dynamics in the macaque fronto-parietal 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.

Scherberger H, **Dann B**, Kronen P (in prep.) Population single unit recording from primate sub-surface cortical areas using floating multi-electrode arrays.

Talks

2017 Three population dynamical states describe single trial activity in the fronto-parietal network. 10th Primate Neurobiology Meeting. March 6th, Göttingen, Germany.

2016 Separable decoding of visual, intention, and movement information from the fronto-parietal grasping-network. 6th International Brain-

5. Curriculum Vitae

- 2016 Computer Interface Meeting. May 31st, Pacific Grove, CA, USA.
- 2016 Functional rich-club, hub neurons of the front-parietal network are predominantly oscillators. Ernst-Strüngmann Institute (ESI), May 23rd, Frankfurt, Germany.
- 2014 Delta and Beta dynamics of the fronto-parietal spiking-network. 7th Primate Neurobiology Meeting. March 26th, Tübingen, Germany.

Posters and conference participations

- Dann B**, Michaels JA, Scherberger H (2016). Separable decoding of visual, intention, and movement information from the fronto-parietal grasping-network. 6th International Brain-Computer Interface Meeting. Pacific Grove, CA, USA.
- 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, Stefan Schaffelhofer S, Scherberger H (2015). The single unit network for hand grasping has a small-world and rich-club topology with oscillators as hubs. 6th biennial NEURIZONS Conference. Göttingen, Germany.
- Dann B**, Michaels, JA, Stefan Schaffelhofer S, Scherberger H. Small world and rich club dynamics of the single unit motor network and their correlation to oscillations. 8th Primate Neurobiology Meeting. Göttingen, Germany.
- Dann B**, Michaels, JA, Stefan Schaffelhofer S, Scherberger H (2015). Small world and rich club dynamics of the single unit motor network and their correlation to oscillations. 11th Göttingen Meeting of the German Neuroscience Society. Göttingen, Germany.
- Wellner 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.
- Wellner 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. ESI-SyNC - Workshop on Inter-areal interactions 2014. Frankfurt, Germany.

5. Curriculum Vitae

- 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.
- Wellner B**, Suway SB, Scherberger H (2014). Neuronal network dynamics within and between frontal and parietal cortex in a massively parallel recording approach in the macaque monkey. Computational and Systems Neuroscience (Cosyne) 2014. Salt Lake City, UT, USA.
- Wellner B**, Suway SB, Scherberger H (2013). Network dynamics of spike-spike interactions within and between frontal and parietal cortex. 43rd Annual Meeting of the Society for Neuroscience. San Diego, CA, USA.
- 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.
- 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**, Suway SB, Scherberger H (2013). Neuronal network dynamics within and between frontal parietal cortex in a massively parallel recording approach. The Assembly and Function of Neuronal Circuits 2013. Ascona, Switzerland.
- Wellner B**, Michaels JA, Wellner A, Scherberger H (2013). Single trial neuronal correlates of decision-making for hand grasping in macaque area F5 and AIP. 10th Göttingen Meeting of the German Neuroscience Society. Göttingen, Germany.
- Wellner B**, Michaels JA, Wellner A, Scherberger H (2013). Single trial neuronal correlates of decision-making for hand grasping in macaque area F5 and AIP. 6th Primate Neurobiology Meeting. Göttingen, Germany.
- Suway SB, **Wellner B**, Wellner A, Scherberger H (2012). Encoding of reward value in AIP and F5 during decision-making in a grasping task. AREADNE 2012. Santorini, Greece.
- Wellner B**, Wellner A, Suway SB, Scherberger H (2012). Differential neuronal activity during freely chosen and instructed hand grasping movements. Internal Conference on Brain Dynamics and Decision Making 2012. Ascona, Switzerland.
- Suway SB, **Wellner B**, Wellner, A, Scherberger, H (2012). Representation of reward value in area AIP and F5 during a grasping task. 5th Primate Neurobiology Meeting. Tübingen,

5. Curriculum Vitae

Germany.

Wellner B, Wellner A, Suway SB, Scherberger H (2012). Different activity for choice and instructed trials for grasping in AIP and F5. 5th Primate Neurobiology Meeting. Tübingen, Germany.

Wellner B, Wellner A, Scherberger, H (2011). Neuronal correlates of decision-making for hand grasping. 40th Annual Meeting of the Society for Neuroscience. San Diego, Washington, DC, USA.

Wellner A, Townsend B, **Wellner B**, Scherberger H (2011). Oscillatory power in macaque areas F5 and AIP reflect grasping movements. 40th Annual Meeting of the Society for Neuroscience. San Diego, Washington, DC, USA.

Wellner A, **Wellner B**, Scherberger, H (2011). Representation of grasp movements in oscillatory activity of macaque's area F5 and AIP. 9th Göttingen Meeting of the German Neuroscience Society. Göttingen, Germany.

Wellner B, Wellner A, Scherberger, H (2011). Decision-making between two grasp types modulated by different reward values in Area AIP and F5 of macaque monkey. 9th Göttingen Meeting of the German Neuroscience Society. Göttingen, Germany.

Wellner A, **Wellner B**, Scherberger, H (2011). Representation of grasp movements in oscillatory activity of macaque's area F5 and AIP. 4th Primate Neurobiology Meeting. Göttingen, Germany.

Wellner B, Wellner A, Scherberger, H (2011). Decision-making between two grasp types modulated by different reward values in Area AIP and F5 of macaque monkey. 4th Primate Neurobiology Meeting. Göttingen, Germany.

Wellner A, **Wellner B**, Scherberger H (2011). Oscillatory power of neurons appear to encode for different grasp movements in macaque's area F5 and AIP reflect grasping movements. FENS-IBRO HERTIE Winter School, The systems neuroscience of primate hand function: models, mechanisms, rehabilitation and mirror systems. Obergurgl, Austria.

Wellner B, Wellner A, Scherberger, H (2011). Decision-making between different grasp types in AIP and F5 of macaque monkey modulated by different reward values. FENS-IBRO HERTIE Winter School, The systems neuroscience of primate hand function: models, mechanisms, rehabilitation and mirror systems. Obergurgl, Austria.

Wellner B, Wellner A, Scherberger, H (2010). Decision-making between different grasp types

5. Curriculum Vitae

in AIP and F5 of macaque monkey. 39th Annual Meeting of the Society for Neuroscience. San Diego, CA, USA.

Wellner A, **Wellner B**, Scherberger, H (2010). Long term recording using carbon nanotube coated floating microarrays. 3rd Primate Neurobiology Meeting. Tübingen, Germany.

Teaching experience

- | | |
|------|--|
| 2016 | Experimental design in Neurobiology, Laboratory Animal Course on Primates (LAS), German Primate Center (DPZ), Göttingen, Germany |
| 2008 | Introduction to Neurophysiology, lecture for the master program biology, Max-Planck Institute for Brain Research, Frankfurt, Germany |
| 2008 | Methods of Neurophysiology, lecture for the master program biology, Max-Planck Institute for Brain Research, Frankfurt, Germany |

Workshop participations

- | | |
|-----------|--|
| 2016 | Primate Neurobiology Methods: Behavior, Experiments, Analysis, and Ethics |
| 2015 | EUPRIM-NET Course on General Primatology, Göttingen, Germany |
| 2014 | ESI-SyNC - Workshop on Inter-areal interactions, Frankfurt, Germany |
| 2014 | Laboratory Animal Science Course on Primates, Tübingen Germany |
| 2014 | Fieldtrip Workshop, Göttingen, Germany |
| 2012 | Fieldtrip Workshop, Frankfurt, Germany |
| 2012 | Scientific integrity & the responsible conduct of research, Göttingen, Germany |
| 2011 | Press communication training of the Klaus Tschira Stiftung, Göttingen, Germany |
| 2010-2011 | FENS-IBRO HERTIE Winter School, The systems neuroscience of primate hand function: models, mechanisms, rehabilitation and mirror systems, Obergurgl, Austria |
| 2009 | FELASA Compact course: experimental animals and replacement methods, Course on Laboratory Animal Science, Berlin, Germany |

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

Jonathan A Michaels¹, Benjamin Dann¹, Hansjörg Scherberger^{1,2,†}

¹Deutsches Primatenzentrum GmbH, Kellnerweg 4, 37077 Göttingen, Germany

²Faculty of Biology, Georg-August-Universität Göttingen, 37073 Göttingen, Germany

[†]Corresponding author. Email: hscherberger@dpz.eu

Acknowledgements: We would like to thank Clodhna Quigley for comments on an earlier version of the manuscript.

RESEARCH ARTICLE

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

Jonathan A. Michaels¹, Benjamin Dann¹, Hansjörg Scherberger^{1,2*}

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

* hscherberger@dpz.eu



 OPEN ACCESS

Citation: 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(11): e1005175. doi:10.1371/journal.pcbi.1005175

Editor: Byron Yu, Carnegie Mellon University, UNITED STATES

Received: March 31, 2016

Accepted: September 24, 2016

Published: November 4, 2016

Copyright: © 2016 Michaels et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant modeling data are within the paper and its supporting information files. We also compared our model with third party data (Churchland et al, 2012; <http://doi.org/10.1038/nature11129>) that was downloaded from: <http://churchlandlab.neuroscience.columbia.edu/links.html>. This data can be accessed there or by contacting the authors of the original study: Mark Churchland (email: mc3502@cumc.columbia.edu); Krishna Shenoy (email: shenoy@stanford.edu).

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.

Funding: This work was supported by Deutsche Forschungsgemeinschaft (SCHE 1575/1-1 & SFB 889, Project C09). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

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]).

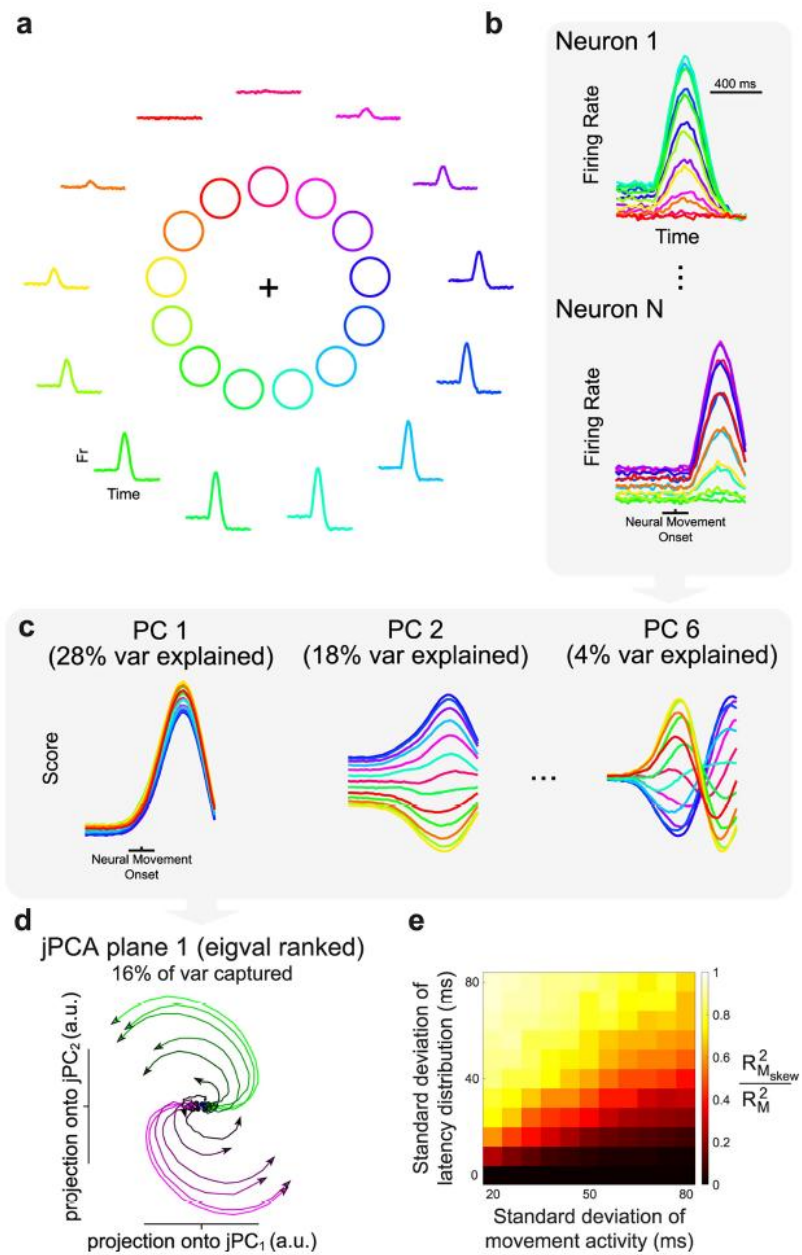


Fig 1. Simulation of a velocity-tuning based model with variable neuron-kinematic latencies. (a) Task design of a 13-direction center-out reaching task. The firing of a simulated neuron is plotted around the reach directions. (b) Two example neurons with differing latencies. (c) Principal components (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.

doi:10.1371/journal.pcbi.1005175.g001

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 $R_{M_{low}}^2$ and R_M^2 (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

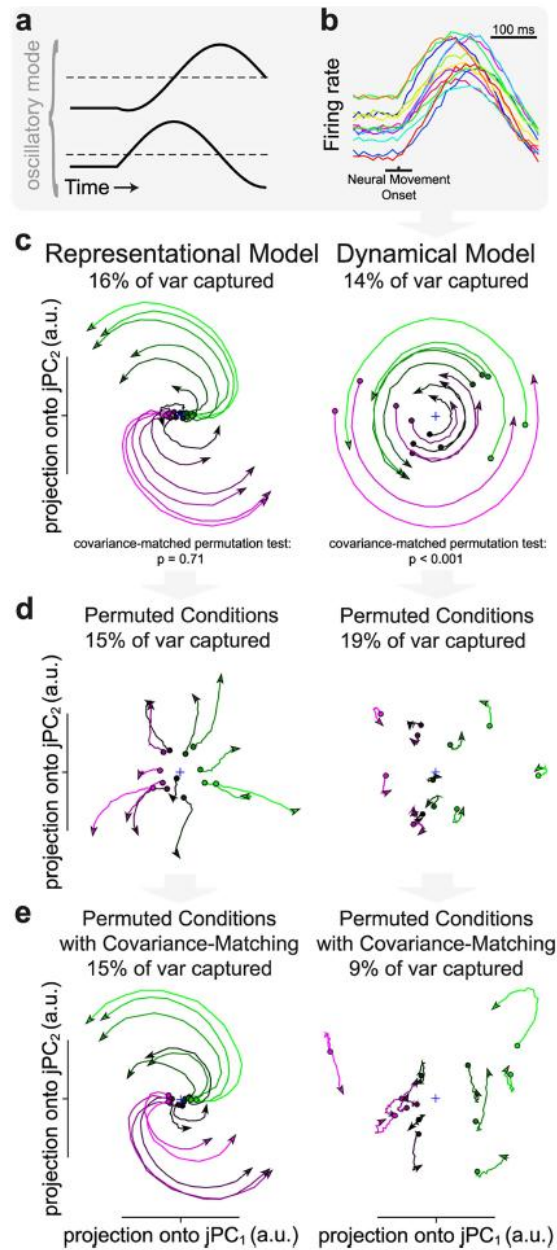


Fig 2. Comparing rotational structure between the representational and the dynamical models. (a) One of the two oscillatory modes (2.8 Hz) used to generate the simulated muscle activity of all conditions (2.8 Hz and 0.3 Hz). (b) Firing rate of an example neuron of the dynamical model for all 13 conditions. 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: $R_{\text{observed}}^2/R_{\text{fit}}^2$). (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].

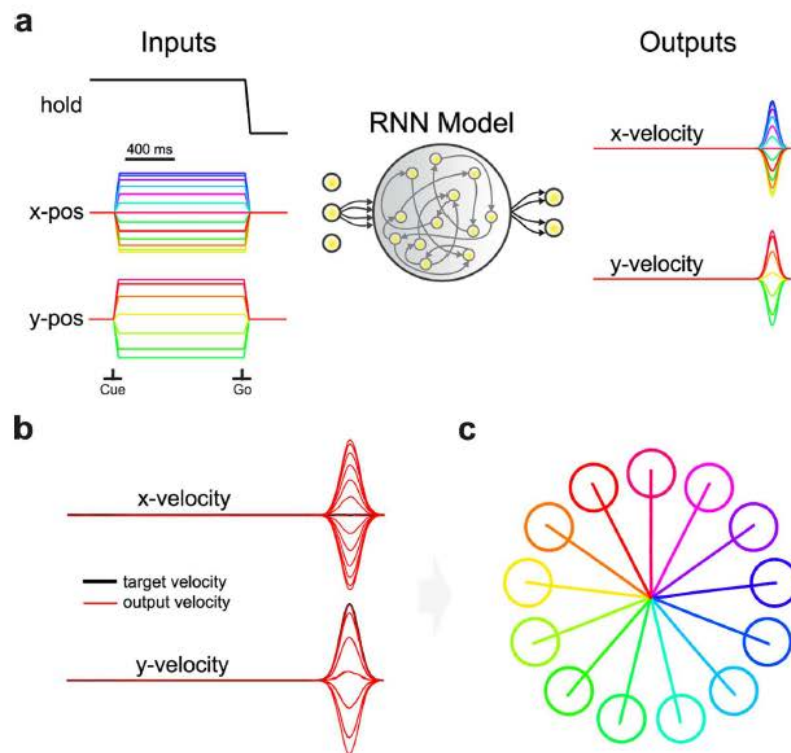


Fig 3. Schematic of recurrent neural network performing center-out reaching. (a) Schematic of RNN, with input layer, hidden layer, and output layer. The three inputs were a condition-independent hold signal that was released at the go cue and two inputs representing the target angle. The two outputs 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.

doi:10.1371/journal.pcbi.1005175.g003

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).

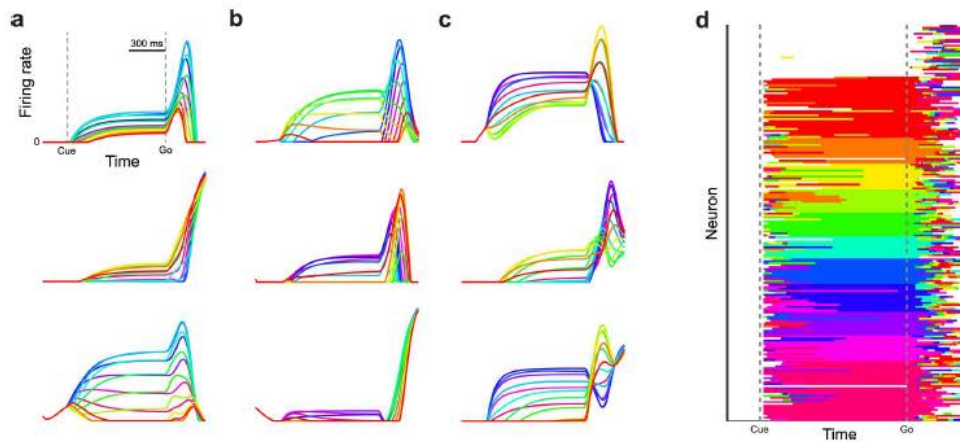


Fig 4. Tuning properties of RNN neurons. (a) Three example units for which the pattern of directional tuning remained highly correlated between the delay period and movement. (b) Same as a, but for example units that have delay tuning that is not correlated with movement activity. (c) Same as a, but 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.

doi:10.1371/journal.pcbi.1005175.g004

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.

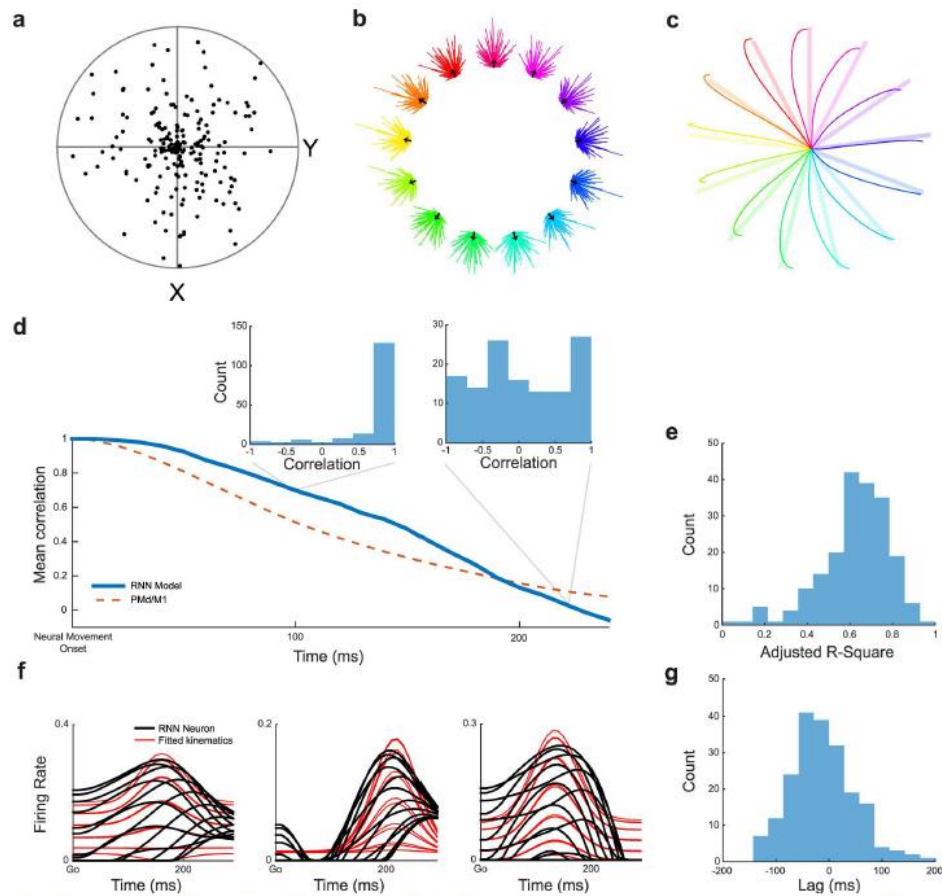


Fig 5. Representational tuning in an RNN for center-out reaching. (a) Preferred movement direction in Cartesian space of all units, corresponding to the magnitude of $b_{1,2}$ and $b_{1,3}$ in Eq 9. (b) Summary of contribution vectors of all individual neurons (one vector each) over the entire movement, 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.

doi:10.1371/journal.pcbi.1005175.g005

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 5E), 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 necessarily lead to mechanistic insight [42], as the observation of rotational structure arose under multiple

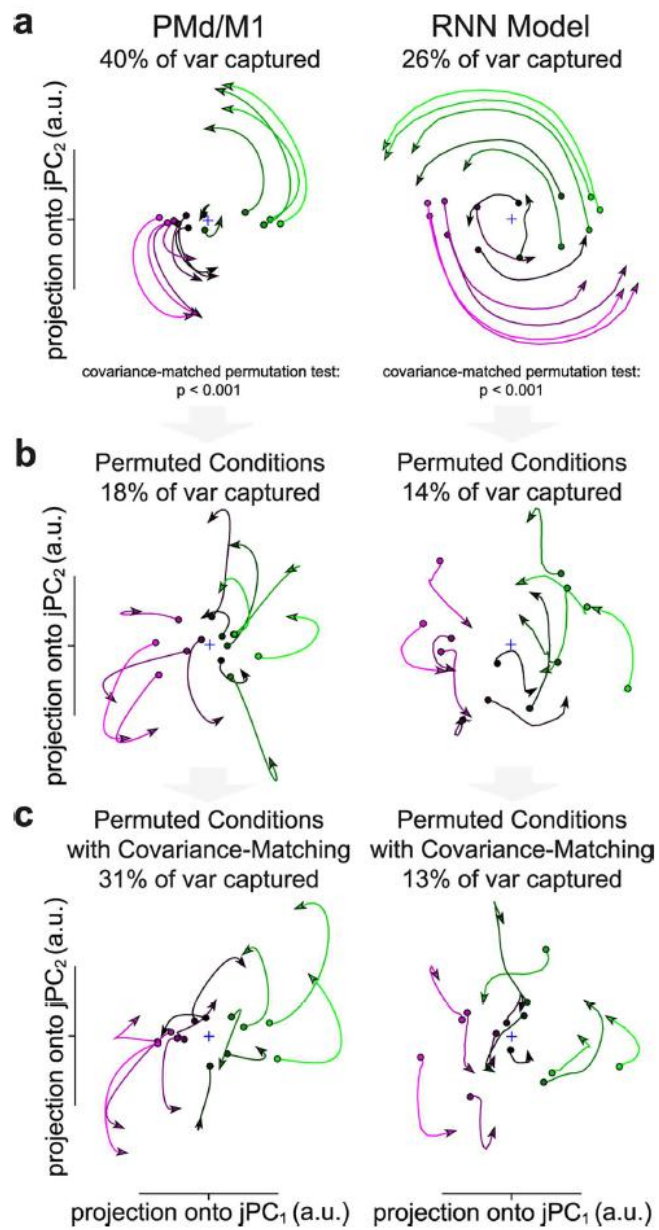


Fig 6. Significant rotational structure in PMd/M1 data and RNN model. Comparison of rotational dynamics for (a) observed, (b) permuted without covariance matching, and (c) covariance-matched data in the first jPCA plane. *p*-values in a are from the CMPT for the rotational goodness-of-fit ratio (RGR: R_{obs}^2/R_{perm}^2) 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,

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




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_{obs}^2/R_{perm}^2). (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 $\mu_0 = \sigma \sqrt{-2 \ln \varphi}$, 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 $c = 1, \dots, 13$ 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 $f_1 = 2.8 \text{ Hz}$ and $f_2 = 0.3 \text{ Hz}$ (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 $n \in \{1, \dots, N\}$ and condition c was generated as a neuron-specific combination of the condition-specific oscillations and offset ($F_{c,1}, F_{c,2}, o_c$)

$$r_{n,c}(t) = \text{Re}(w_{n,1} F_{c,1}(t) + w_{n,2} F_{c,2}(t)) + s_n o_c + \epsilon_{n,c}(t) \tag{3}$$

with the real and imaginary components of the complex coefficients, $w_{n,1}$ and $w_{n,2}$, 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 $\epsilon_{n,c}(t)$ 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: $R_{\text{Meas}}^2 / R_{\text{3d}}^2$) 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

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

$$\text{effect size} = \frac{RGR_{\text{observed}} - \overline{RGR}_{\text{permuted}}}{\sigma_{RGR_{\text{permuted}}}} \quad (4)$$

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

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

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

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

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>)

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

$$C_i = \begin{bmatrix} \frac{b_{i,2}}{k_i} \\ \frac{b_{i,3}}{k_i} \end{bmatrix} \tag{9}$$

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,

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

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 \vec{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: $R_{M_{obs}}^2 / R_M^2$) 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)

Acknowledgments

We would like to thank KV Shenoy for generous use of data and C Quigley for comments on an earlier version of the manuscript.

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](https://doi.org/10.1038/35093102) PMID: [11557980](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1152/jn.00110.2006) PMID: [17005623](https://pubmed.ncbi.nlm.nih.gov/17005623/)
5. Naselaris T, Merchant H, Amirkian 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](https://doi.org/10.1152/jn.00487.2006) PMID: [16971681](https://pubmed.ncbi.nlm.nih.gov/16971681/)
6. Georgopoulos AP, Naselaris T, Merchant H, Amirkian 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](https://doi.org/10.1016/j.neuron.2012.10.041) PMID: [23312524](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.conb.2015.01.012) PMID: [25646932](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1146/annurev-neuro-062111-150509) PMID: [23725001](https://pubmed.ncbi.nlm.nih.gov/23725001/)
10. Yuste R. From the neuron doctrine to neural networks. *Nat Rev Neurosci*. 2015; 16: 487–497. doi: [10.1038/nrn3962](https://doi.org/10.1038/nrn3962) PMID: [26152865](https://pubmed.ncbi.nlm.nih.gov/26152865/)
11. Tanji J, Evars EV. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *J Neurophysiol*. 1976; 39: 1062–1068. PMID: [824409](https://pubmed.ncbi.nlm.nih.gov/824409/)
12. Evars EV. Pyramidal tract activity associated with a conditioned hand movement in the monkey. *J Neurophysiol*. 1966; 29: 1011–1027. PMID: [4961643](https://pubmed.ncbi.nlm.nih.gov/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](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.conb.2015.04.003) PMID: [25932978](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.conb.2007.11.001) PMID: [18093826](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.neuron.2011.05.047) PMID: [21835350](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1523/JNEUROSCI.1714-15.2015) PMID: [26269647](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1152/jn.00808.2006) PMID: [17005608](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.neuron.2013.11.003) PMID: [24462104](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1038/nn.4042) PMID: [26075643](https://pubmed.ncbi.nlm.nih.gov/26075643/)
21. Pandarinath C, Gilja V, Blabe CH, Nuyujukian P, Sarma AA, Soric 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](https://doi.org/10.7554/eLife.07436) PMID: [26099302](https://pubmed.ncbi.nlm.nih.gov/26099302/)

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](https://doi.org/10.1038/nature11129) PMID: [22722855](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1101/sqb.2014.79.024703) PMID: [25851506](https://pubmed.ncbi.nlm.nih.gov/25851506/)
24. Moran DW, Schwartz AB. Motor cortical representation of speed and direction during reaching. *J Neurophysiol*. 1999; 82: 2676–2692. PMID: [10561437](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1038/73964) PMID: [10725930](https://pubmed.ncbi.nlm.nih.gov/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](https://pubmed.ncbi.nlm.nih.gov/7703686/)
27. Schwartz AB. Motor cortical activity during drawing movements: population representation during sinusoid tracing. *J Neurophysiol*. 1993; 70: 28–36. PMID: [8360717](https://pubmed.ncbi.nlm.nih.gov/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](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1038/ng.139) PMID: [18425127](https://pubmed.ncbi.nlm.nih.gov/18425127/)
31. Menozzi P, Piazza A, Cavalli-Sforza L. Synthetic maps of human gene frequencies in Europeans. *Science*. 1978; 201: 786–792. PMID: [356262](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.neuron.2014.04.045) PMID: [24945778](https://pubmed.ncbi.nlm.nih.gov/24945778/)
33. Renart A. Bringing the dynamics of movement under control. *Neuron*. 2014; 82: 1193–1195. doi: [10.1016/j.neuron.2014.06.002](https://doi.org/10.1016/j.neuron.2014.06.002) PMID: [24945762](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1038/nn.3643) PMID: [24487233](https://pubmed.ncbi.nlm.nih.gov/24487233/)
36. Churchland MM, Shenoy KV. Temporal complexity and heterogeneity of single-neuron activity in pre-motor and motor cortex. *J Neurophysiol*. 2007; 97: 4235–4257. doi: [10.1152/jn.00095.2007](https://doi.org/10.1152/jn.00095.2007) PMID: [17376854](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.neuron.2010.09.015) PMID: [21040842](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.conb.2015.01.006) PMID: [25637959](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1523/JNEUROSCI.2643-14.2015) PMID: [26109672](https://pubmed.ncbi.nlm.nih.gov/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.

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](https://doi.org/10.1038/nature11076) PMID: [22596161](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1038/nature06996) PMID: [18509337](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1038/ncomms8759) PMID: [26220660](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1088/1741-2560/9/2/026027) PMID: [22427488](https://pubmed.ncbi.nlm.nih.gov/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](https://doi.org/10.1016/j.neuron.2009.07.018) PMID: [19709635](https://pubmed.ncbi.nlm.nih.gov/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.

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

Jonathan A Michaels¹, Benjamin Dann¹, Rijk W. Intveld¹, Hansjörg Scherberger^{1,2,+}

¹Deutsches Primatenzentrum GmbH, Kellnerweg 4, 37077 Göttingen, Germany

²Faculty of Biology, Georg-August-Universität Göttingen, 37073 Göttingen, Germany

⁺Corresponding author. Email: hscherberger@dpz.eu

Acknowledgements: We thank R. Lbik, and N. Bobb for animal support, M. Dörge for technical support, and B. Lamplmair and Stefan Schaffelhofer for providing illustrations. This work was supported by the Deutsche Forschungs Gemeinschaft (SCHE 1575/1-1).

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}¹German Primate Center, D-37077 Göttingen, Germany, and ²Faculty 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 method 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.

Received April 30, 2015; revised July 3, 2015; accepted July 10, 2015.

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

This work was supported by German Research Foundation Research Grant SCHE 1575/1-1. We thank Natalie Nazarens, Ricarda Ahlert, and Matthias Dörge for technical assistance.

The authors declare no competing financial interests.

Correspondence should be addressed to Hansjörg Scherberger, Neurobiology Laboratory, German Primate Center, Kellnerweg 4, D-37077 Göttingen, Germany. E-mail: HScherberger@dpz.eu.

DOI:10.1523/JNEUROSCI.1714-15.2015

Copyright © 2015 the authors 0270-6474/15/3511415-18\$15.00/0

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; Fluett 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 non-acting 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 non-acting 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.

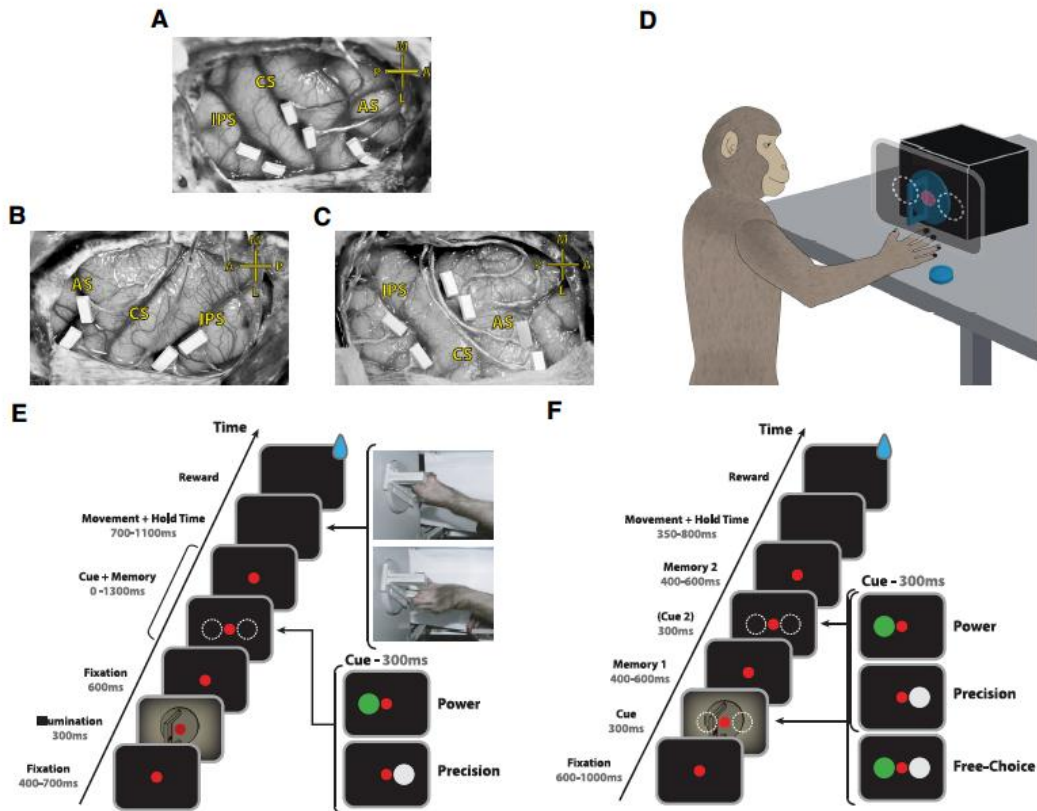


Figure 1. FMA implantation and task design. *A–C*, Array locations for animals B, S, and Z, respectively. Two arrays were placed in F5 on the bank of the arcuate sulcus (AS). Two additional arrays were placed in AIP toward the lateral end of the intraparietal sulcus (IPS). In animals B and Z, two more 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 buprenorphine, *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 CO₂, ~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 method 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 $go \pm \text{some } \Delta t$. 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

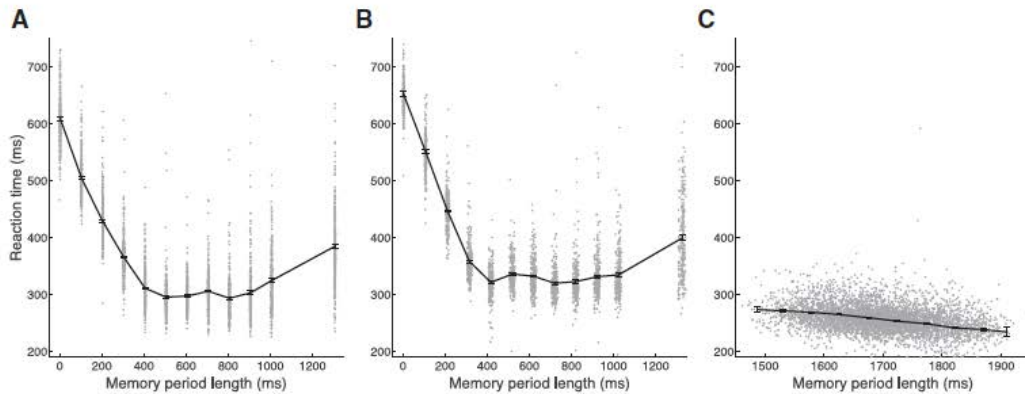


Figure 2. Scatter plots of RT versus memory period length. **A–C**, The RT of animals B, S, and Z, respectively, as a function of memory period length for all task conditions and datasets. The solid line represents the mean, and error bars indicate SEM within nonoverlapping 50 ms bins.

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

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

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:

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

where 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 \pm SD of 63 ± 17 and 28 ± 18 units were recorded in F5 and AIP, respectively, as well as in animal S (mean \pm 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

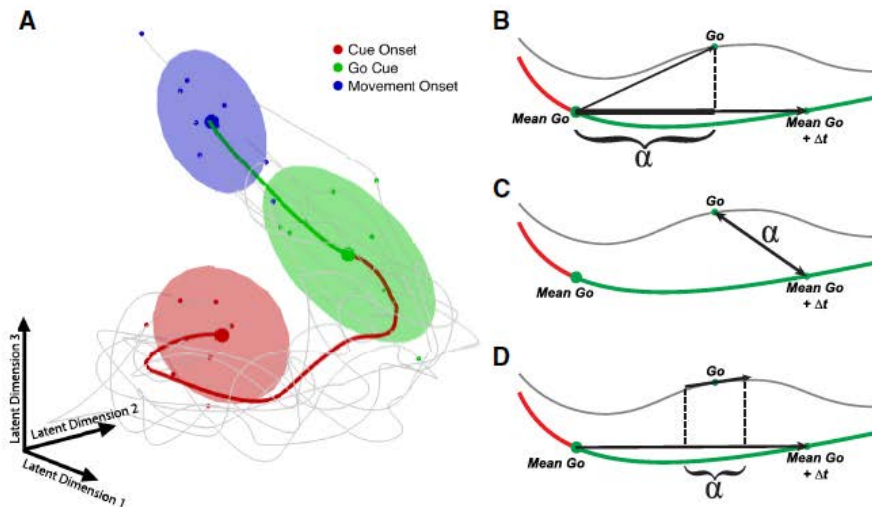


Figure 3. Neural data and RT prediction methods visualized as low-dimensional trajectories. **A**, Neural data of both areas of an exemplar condition reduced to a low-dimensional representation of the trial course (determined by GPFA). Thick trace represents the mean of trials for one condition (instructed 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.

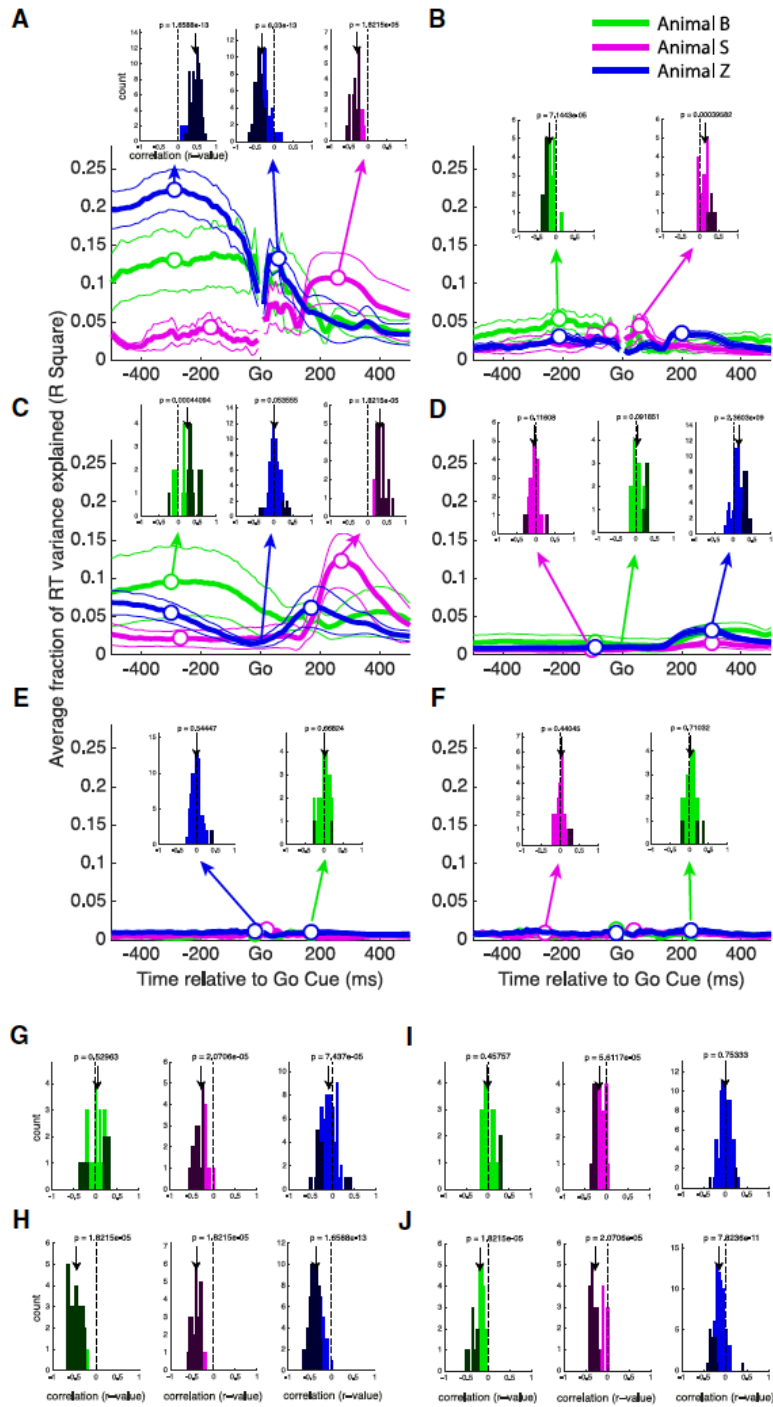
The Euclidian distance method performed similarly to the projection method but explained overall less variance in RT (Fig. 4C,D). In both F5 and AIP, reference points generally produced correlation histograms that were shifted significantly from zero. In most cases the pre-go distribution was shifted 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

Appendix B Predicting reaction time from the neural state space

11422 • J. Neurosci., August 12, 2015 • 35(32):11415–11432

Michaels et al. • Reaction Time Prediction in the Neural State Space



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-

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 Δt 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 4. Determination of the optimal reference time (Δt) 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 Δt 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.

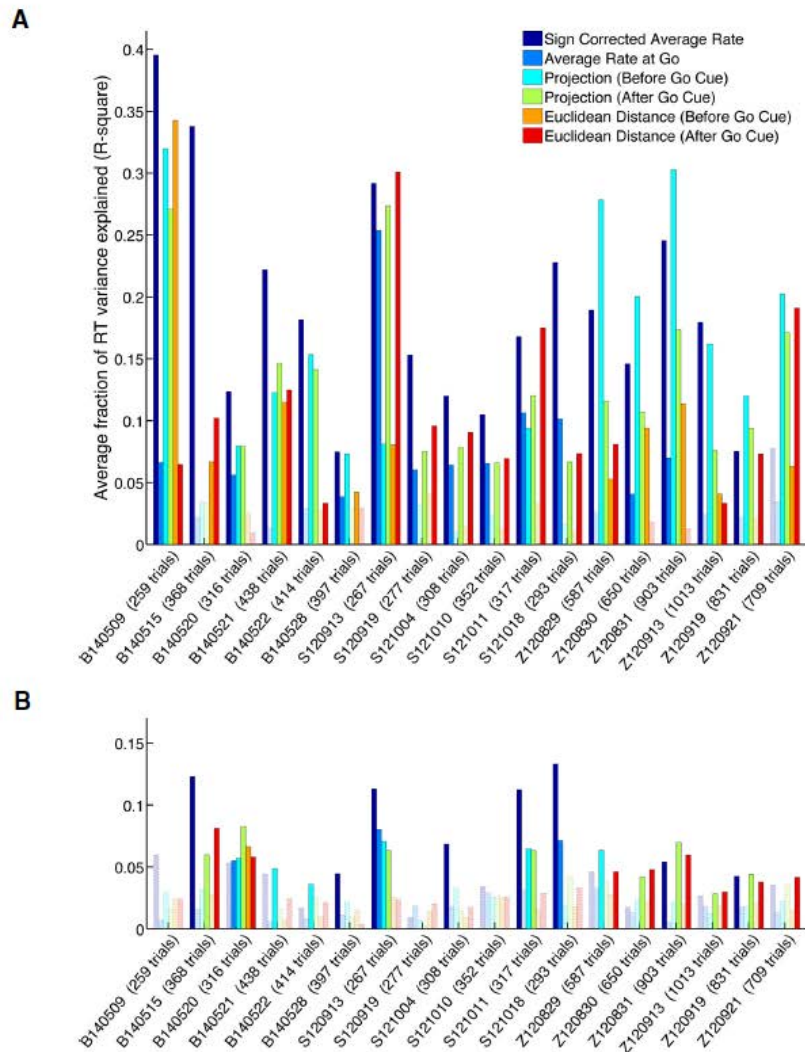


Figure 5. Average fraction of RT variance explained for all methods and datasets (averaged across conditions and cross-validation folds). **A**, Average fraction explained by F5 data. **B**, Average fraction for AIP. Note the clear advantage of area F5 over AIP. Chance-level calculation is based on shuffling 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

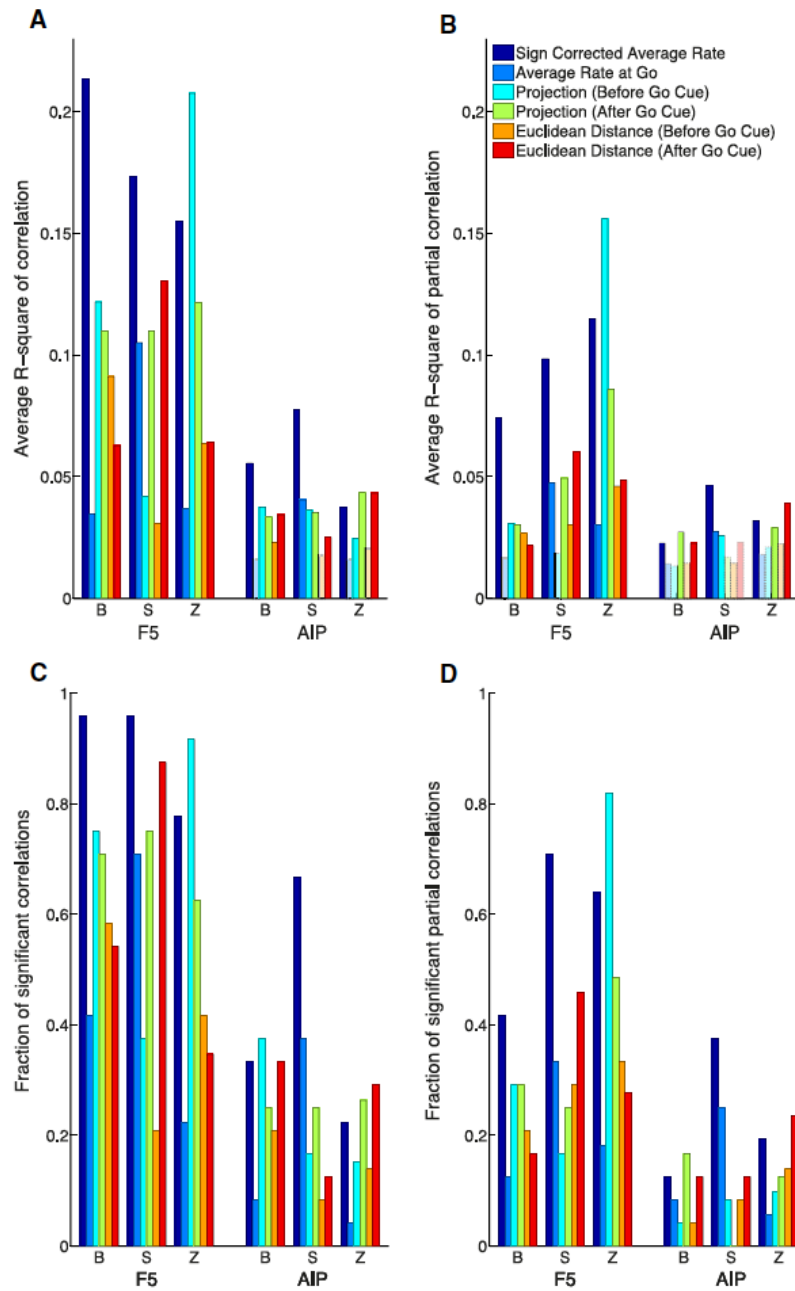


Figure 6. Comparison of prediction performance and fraction of significant full/partial correlations between predictors and RT over all task conditions, datasets, and brain areas. **A**, Average fraction of RT variance explained by correlation. **B**, Average fraction of RT variance explained by partial 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 outnumbered 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

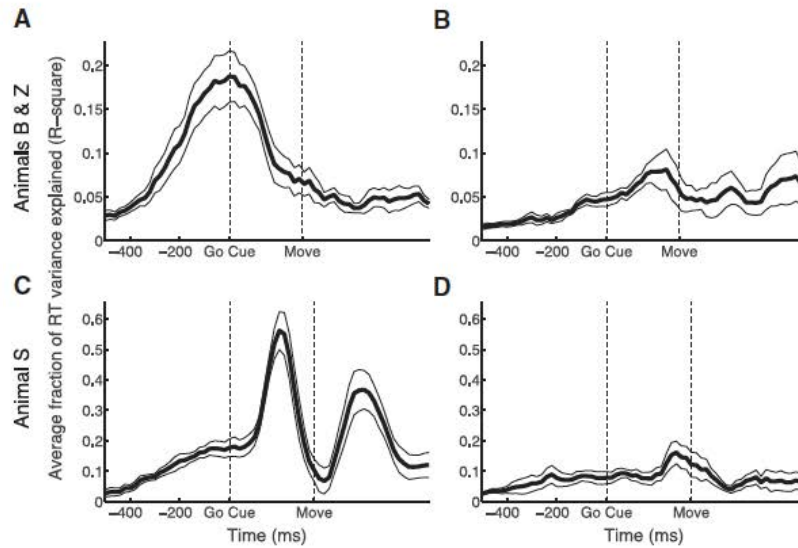


Figure 7. Stability of the sign-correction vector determined at the go cue by the SCAR method. RT prediction is calculated using sign-corrected neural activity around each time point. **A, B**, SCAR method as a function of time for animals B and Z in areas F5 (**A**) and AIP (**B**). **C, D**, SCAR method 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 (Verhoeve 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-

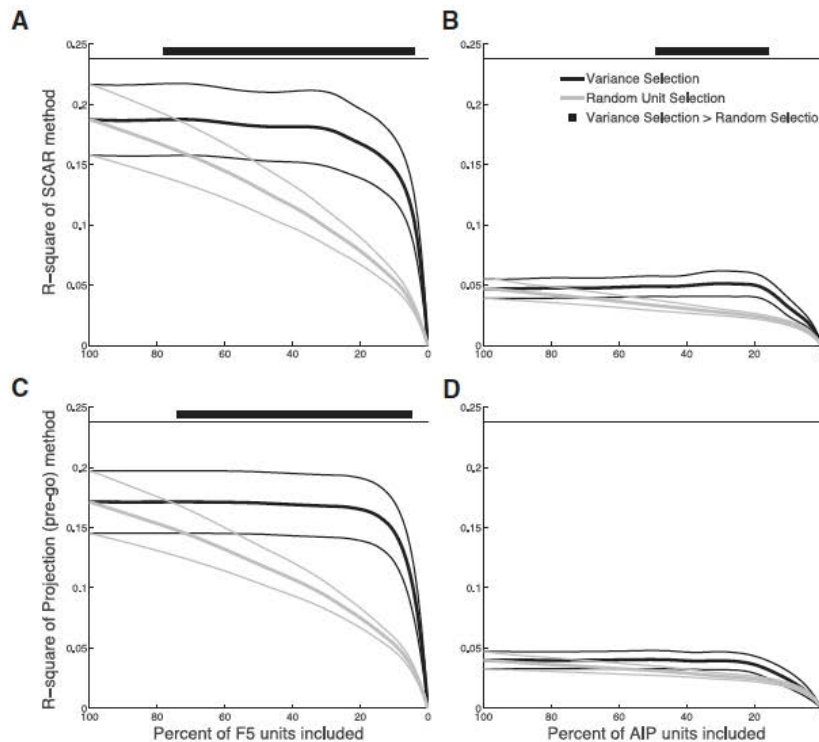


Figure 8. Selection of units by firing rate variance at the go cue for the two best performing methods in animals B and Z. *A, B*, Variance selection of units versus random selection using the SCAR method in F5 (*A*) and AIP (*B*). *C, D*, Variance selection of units versus random selection using the 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 2*A*, 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 9*A*, 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 6*B*, 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 9*A*, 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 9*B* 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 9*A*. However, the orientation of the projection axis is approximately orthogonal to that of an axis running through the slow and fast RT trials, sug-

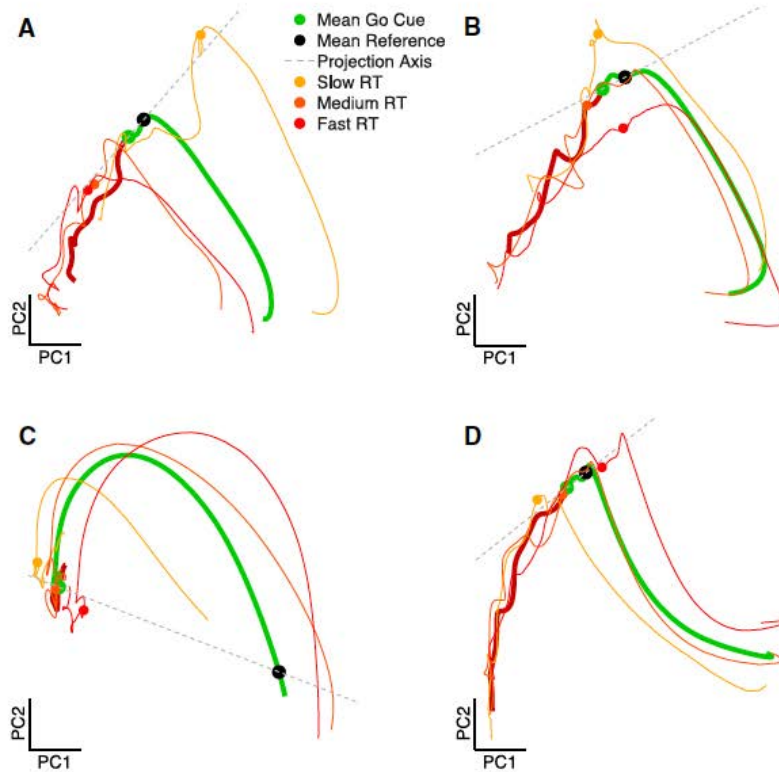


Figure 9. Low-dimensional representation (PCA) of neural state space illustrating day-to-day and animal-to-animal variability. Trajectories are plotted in the two principal components of area F5 explaining the most variance. *A*, Trajectory of a power grip from dataset B140509 in which RT is most variable 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 S120913. 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 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, Qi 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, Umiltà 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, Umiltà 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