NEURAL STATE CHANGES IN PRIMATE MOTOR CORTEX DURING ARM MOVEMENTS WITH DISTINCT CONTROL REQUIREMENTS

by

Steven B Suway

BA, Kenyon College, 2011

Submitted to the Graduate Faculty of

School of Medicine in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2019

UNIVERSITY OF PITTSBURGH

SCHOOL OF MEDICINE

This dissertation was presented

by

Steven B Suway

It was defended on

January 23, 2019

and approved by

Peter L. Strick, Distinguished Professor, Neurobiology

Robert E. Kass, Professor, Statistics, Machine Learning, CMU

Carl R. Olson, Professor, Center for the Neural Basis of Cognition

William R. Stauffer, Assistant Professor, Neurobiology

Giuseppe Pellizzer, Associate Professor, Neuroscience, University of Minnesota

Dissertation Director: Andrew B. Schwartz, Distinguished Professor, Neurobiology

NEURAL STATE CHANGES IN PRIMATE MOTOR CORTEX DURING ARM MOVEMENTS WITH DISTINCT CONTROL REQUIREMENTS

Steven B Suway, PhD

University of Pittsburgh, 2019

The primary motor cortex (M1) is an important structure of the motor system that contributes to many aspects of movement. Firing patterns of M1 neurons can be surprisingly complex, and there is substantial interest in understanding these patterns and their relation to behavior. Here, we characterize the temporal structure of M1 activity during reaching in several ways. First, we show that single neurons encode movement information in a series of discrete segments. Information is stably encoded during each brief segment, and the firing patterns of most neurons transition between segments at similar times during movement. This pattern may therefore reflect transitions between different neural "states." Next, we establish that the sequence of states observed during behavior is related to a sequence of distinct drivers, including visuospatial information and visual feedback from a movement. If no feedback is provided, neurons may produce a truncated response sequence. Last, we link the temporal structure of firing patterns to the structure of reaches and demonstrate that the classical two-component model of reaching is reflected in M1 activity. Our findings may help establish a useful framework for interpreting seemingly complex neural activity during behavior.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1.0 INTRODUCTION

1.1 BACKGROUND

Studies of the primary motor cortex (M1) have played a critical role in our understanding of the neural basis of movement. Early experiments involving electrical stimulation of the cortex established that M1 is of special importance for generating motor output (Fritsch & Hitzig 1870). Investigators found they could elicit movement of different body parts by stimulating different parts of M1, leading to the development of somatotopic maps from cortex to the body's effectors (Leyton & Sherrington 1917; Penfield & Boldrey 1937; Penfield & Rasmussen 1950; Woolsey 1958). Though cortical stimulation studies are often difficult to interpret (Asanuma & Sakata 1967; Jankowska, Padel, & Tanaka 1975b), the anatomical linkage from M1 to spinal circuitry that commands muscle contraction has long been appreciated (Holmes & May 1909; Levin 1936; Lassek 1942; Lawrence & Kuypers 1965; Wiesendanger 1969). Efferents from M1 constitute a large portion of the corticospinal tract (CST, Dum & Strick 1991), and a species' manual dexterity is related to the density of CST terminations in the ventral horn (Heffner & Masterton 1983; Bortoff & Strick 1993). The primary motor cortex is thus undoubtedly an important structure of the motor system.

With the advent of single-neuron recording in awake behaving monkeys, investigators began to characterize the functional relation between M1 and volitional behavior in great detail.

Evarts (1965) quickly established a link between M1 pyramidal tract neuron (PTN) discharge and naturalistic behaviors. However, he recognized at the time that those types of complex movements were "unsuitable for precise quantitative analysis." Seeking a more controlled experiment, he then examined M1 PTN activity during simple wrist extensions (Evarts 1966). Firing rates tended to change shortly before contralateral movement, and response latency was correlated with reaction time, suggesting a causal relationship to behavior. Evarts initially suspected that wrist displacement was the primary determinant of firing rate, but his subsequent work demonstrated a strong relation to flexion/extension force (Evarts 1968), even for isometric muscle contraction (Evarts 1969). He concluded that "output of motor cortex PTNs is related to the patterns of muscular contraction" rather than to "joint displacement or the joint position" (Evarts 1969). Still, he found that single neuron firing properties could be idiosyncratic– some units had a non-monotonic relation to force and some showed force-dependent response latency (Evarts 1968). Complexities in the data and limitations of experimental design made it difficult to draw definitive conclusions. Other investigators suggested the functional link between M1 and motor output was rather flexible (Fetz 1969; Fetz & Finocchio 1971), or that firing patterns of M1 neurons covaried with several behavioral features (Thach 1978).

Despite some disagreement between studies, most data highlighted the "directional" nature of M1 activity: firing rates commonly increased for motor output (displacement, force, etc.) in one direction, but decreased for output in the opposite direction. A seminal work by Georgopoulos et al (1982) expanded this relationship to whole-arm movements using a "center-out" reaching task. In this now-classical paradigm, monkeys reached from a central starting point outward to a peripheral target, with targets arranged in a circle around the center. Using this paradigm, it was shown that firing rates in M1 varied smoothly and systematically with

reach direction. A broad cosine-shaped curve was a remarkably good fit for time-averaged firing rates as a function of movement direction. Cosine-like directional "tuning" of neurons was found to be valid for movements in three dimensions (Schwartz, Kettner, & Georgopoulos 1988) and for isometric force production (Georgopoulos et al 1992). Each recorded neuron seemed to have a different "preferred" direction (PD) in space for which peak firing was observed, a property that proved useful for "reading out" or "decoding" the movement from brain activity (Georgopoulos et al 1983; Schwartz 1994; Schwartz & Moran 1999). Common decoders such as the population vector algorithm (PVA) are based on the directional tuning properties of M1 neurons, and these algorithms hold promise as the basis of restorative prosthetic devices that interface between brain activity and robotic limbs (Wessberg et al, 2000; Serruya et al, 2002; Taylor, Helms Tillery, and Schwartz, 2002; Hochberg et al, 2006; Velliste et al, 2008; Wodlinger et al, 2015).

1.2 THE RELATION BETWEEN M1 ACTIVITY AND MOVEMENT PARAMETERS

1.2.1 Information carried by M1 neurons is diverse

Despite the success of movement decoders and brain-computer interfaces (BCIs) in lab settings, there are still significant controversies surrounding the "function" of M1 and its role in movement generation and control. Aside from reach direction and force production, M1 firing rates covary with a surprisingly large and diverse set of movement parameters, making it difficult to draw unitary conclusions about the drivers of this activity. The results of numerous studies have argued in favor of kinematics (like direction) as the primary type of movement

feature “represented” or “encoded” by M1 activity. Coding for movement velocity, which is the direction scaled by speed, appears to be particularly robust (Schwartz 1993; Moran & Schwartz 1999; Wang et al 2007), and our research group recently demonstrated skilled control of velocity through a BCI (Inoue et al 2018). However, M1 activity has also been related to distance or position (Kettner, Schwartz, & Georgopoulos 1988; Fu et al 1995), acceleration (Ashe & Georgopoulos 1994), and joint angular velocity (Reina, Moran, & Schwartz 2001).

As described above, other studies have emphasized the importance of kinetic parameters (i.e. forces) in driving M1 activity (Evarts 1968, 1969; Cheney & Fetz 1980; Georgopoulos et al 1992; Sergio, Hamel-Pâquet, & Kalaska 2005). A related perspective is that M1 is best understood as a source of “upper motoneurons” that specify the strength of muscle contraction. In this view, M1 is positioned lower in the hierarchy of motor control (closer to the final effector) in contrast to the relatively high position assumed by advocates of kinematic coding (closer to an abstract movement plan). This notion is supported functionally by studies demonstrating “post-spike effects” of M1 neurons on electromyograms (EMG) recorded from forelimb muscles (Fetz & Cheney 1980; Schieber & Rivlis 2005), and by analyses that predict EMG patterns from M1 firing patterns (Morrow & Miller 2003; Cherian, Krucoff, & Miller 2011). Yet, M1 is highly active even when muscles are not, for example while monkeys wait for a “go” cue in instructed-delay tasks (Tanji & Evarts 1976; Georgopoulos, Crutcher, & Schwartz 1989; Alexander & Crutcher 1990a). In fact, M1 PTNs exhibit directionally-tuned delay period activity even more commonly than non-PTNs (Tanji & Evarts 1976).

While these experimental findings may seem in conflict, there is general agreement that “the search for a single, dominant, encoded movement parameter [in M1] is unlikely to be fruitful” (Johnson, Mason, & Ebner 2001). Indeed, apparently strong parameter tuning can be

found to be invalid in new contexts (Caminiti, Johnson, & Urbano 1990; Aflalo & Graziano 2006), single units often display "mixed" selectivity for several parameters (Lurito, Georgakopoulos, & Georgopoulos 1991; Shen & Alexander 1997a), and tuning properties may not be stationary in time (Fu et al 1995; Georgopoulos et al 1989; Churchland & Shenoy 2007; Churchland et al 2010). Thus, clear and succinct descriptions of the "function" of M1 have been elusive.

1.2.2 Parameter dissociation and tuning changes

In part, controversies surrounding the specific information encoded by M1 neurons are complicated by the fact that many of the relevant parameters are correlated during normal behavior (Evarts 1968). For example, if a neuron fires most strongly during an upward reaching movement, we cannot readily discern if its discharge relates to a cognitive plan specifying "move up," or to the contraction of the arm muscles that counter gravity to enact this plan. Moreover, arm muscles are cosine-tuned to reach direction much like M1 neurons (though with a different distribution of PDs, see Schwartz & Moran 1999). Central to this difficulty as it pertains to neural coding is the notion of "coordinate systems" of movement parameters— that is, neurons may encode information in (for example) "extrinsic" coordinates of movement in a workspace, or "intrinsic" coordinates of an arm muscle or joint angle.

Substantial progress in resolving the "muscles vs. movements" debate of M1 coding was achieved by Kakei, Hoffman, & Strick (1999) using an elegant task that dissociated three relevant coordinate systems. A monkey was trained to make wrist movements in eight directions with the wrist oriented in different postures, and directional tuning of muscles and neurons was evaluated for each posture. As the wrist was rotated from a pronated to supinated posture, the

PDs of several wrist and finger muscles also rotated relative to extrinsic space. Although the angle of the wrist rotated a full 180°, muscle PD changes were only around 70° on average; thus, extrinsic movement direction was dissociated from both wrist-centric and muscle-centric coordinates. Leveraging this dissociation, the PDs of M1 neurons were measured for each wrist posture and compared with those of the limb muscles. Posture-related tuning changes were found for roughly 30% of the population, with PD shifts in a range comparable to that observed for limb muscles. These neurons were considered to be "muscle-like" in their tuning properties. Conversely, the PDs of about half the population were found to be relatively invariant with wrist posture and thus considered to have "extrinsic-like" tuning properties.

1.2.3 Motor adaptation

Although directional tuning is clearly related to multiple movement parameters, the study by Kakei, Hoffman, & Strick (1999) is a potent demonstration that this tuning property is useful for probing neural processes, given a properly designed dissociation paradigm. Motor adaptation is another scenario in which movement parameters can be dissociated. For example, Shadmehr & Mussa-Ivaldi (1994) studied reaches made using a manipulandum to which they could apply a mechanical force field. When the force field was applied, direction of movement was effectively dissociated from the direction of applied forces. Although subjects initially produced distorted trajectories, they gradually adapted to the perturbation and recovered their original performance. Reaches are believed to involve a kinematic plan, which tends to follow a smooth, straight path to an end point (Morasso 1981; Flash & Hogan 1985), and an "internal model" that relates this plan to musculoskeletal dynamics. Shadmehr & Mussa-Ivaldi (1994) reasoned that, to produce a straight trajectory under perturbed conditions, subjects would need to update an internal model of

the dynamics between the environment, their motor output, and their hand path. Put another way, force field adaptation dissociates the kinematic plan from the "inverse model" that generates motor output. Using a similar force field perturbation, Li, Padoa-Schioppa, & Bizzi (2001) found that M1 directional tuning in monkeys can be related to both of these aspects of movement.

1.3 TEMPORAL DYNAMICS OF M1 ACTIVITY DURING BEHAVIOR

1.3.1 Visuomotor transformations

While many studies have focused on the relation between M1 activity and movement per se, others have found that visuospatial processes related to movement constitute yet a further class of M1 drivers. These processes have been studied by dissociating vision and movement, often using a visuomotor perturbation. As with force field perturbations, abrupt application of a visuomotor perturbation initially causes subjects to produce distorted trajectories, and normal performance is gradually recovered over repeated trials (Bossom 1965; Clower et al 1996; Yin & Kitazawa 2001; Krakauer, Ghez, & Ghilardi 2005). Visuomotor perturbations can be implemented using prisms (Baizer, Kralj-Hans, & Glickstein 1999; Kurata & Hoshi 1999; Yin & Kitazawa 2001) or rotated mappings between movement and a computer cursor (i.e. "visuomotor rotations," Shen & Alexander 1997a; Wise et al 1998; Krakauer et al 2000; Krakauer et al 2005). Results of these studies have highlighted the temporally complex nature of M1 firing: beyond heterogeneous or mixed parameter coding, neurons often code for different parameters in sequence as behavior progresses. Although the original work on cosine tuning used firing rates

averaged over the whole movement (e.g. Georgopoulos et al 1982), the detailed time-course of firing can be much more complex than this relationship would suggest.

In one notable experiment, Shen & Alexander (1997a) trained monkeys to make center-out movements using a joystick that moved a computer cursor. During a session, the mapping from joystick to cursor was rotated by 90° to mimic a visuomotor rotation– moving the joystick rightward (for example) would move the cursor downward. Monkeys required about 80-120 trials to recover normal performance after the mapping change. Once vision and movement were dissociated, M1 directional encoding was related to one or the other (or both). During the reaction time, right before movement, many neurons coded for the direction of the visual instruction irrespective of the upcoming joystick movement. However, encoding properties gradually changed throughout movement to instead reflect joystick direction, or a combination of cursor and joystick direction. This result suggested that the type of information encoded by M1 neurons could change in sequence under certain task constraints. Other studies that have dissociated vision from movement have produced similar findings. For example, M1 neurons in monkeys trained to make a movement in the direction opposite to a stimulus first coded for the stimulus, then for the movement (Zhang et al 1997; Alexander & Crutcher 1990a).

These coding patterns have commonly been interpreted as a hallmark of a "sensorimotor transformation," which involves a serial sequence of steps to convert movement-relevant sensory information (i.e. visuospatial stimuli) into an appropriate motor response. In this view, cortical networks transform information from sensory coordinates to motor coordinates, though Kalaska et al (1997) have noted that "the co-ordinate transformation model may be only a metaphorical description of the underlying causal mechanisms" of cortical activity. Still, they maintain that the framework "retains heuristic value… [for characterizing] the properties of single cells and

neural circuits" and their relation to behavior. Indeed, visuomotor dissociation tasks have demonstrated that time-varying encoding patterns in M1 may carry information sequentially in a manner that is appropriate for a given task.

In line with this view, other studies have indicated the particular cognitive strategy used during a task is reflected in M1 temporal patterns. Georgopoulos et al (1989) trained a monkey to make center-out reaches based on a rule determined by the brightness of a cued reach target. For one brightness level, the monkey reached directly toward the target as in normal center-out tasks. For another brightness level, reaches were made 90° to the illuminated target. The physical movement was therefore dissociated from the visuospatial stimulus. Movement direction was decoded from M1 firing rates recorded during the task using the population vector algorithm (PVA). As with previous studies (e.g. Georgopoulos et al 1986), when reaches were made toward the cued target, PVA output grew in magnitude shortly before the onset of movement and pointed straight toward the reach direction. However, when reaches were made 90° to the cue, the population vector initially pointed in the direction of the cue before rotating and pointing toward the actual movement direction. This result was believed to reflect a "mental rotation" of an imagined vector that first pointed toward the stimulus, then toward the movement (Georgopoulos & Massey 1987).

In a related experiment (Pellizzer, Sargent, & Georgopoulos 1995), a monkey was trained on a visuomotor "context-recall" task that required a movement direction to be chosen based on the temporal order of a sequence of three visuospatial cues. When one of the cues changed color, the monkey reached toward the cue that followed that one in the sequence. This rule meant that a reach was never made to the first target, and since no additional cues were displayed after the third, only the first or second cue changed color on each trial. The correct response was

therefore fully specified by the second cue in the sequence. If the second cue changed color, a reach was made to the third cue. Otherwise, a reach was made to the second cue. This task constraint was reflected in the of responses of M1 neurons: based on the directional tuning of each neuron, it was found that responses initially coded for the direction of the second cue before abruptly switching and coding for the upcoming movement direction. This property could also be demonstrated using the PVA, as was the case with the mental rotation experiment. However, rather than a gradual rotation, the context-recall task was associated with an abrupt encoding change, possibly reflecting the different cognitive strategies employed in the two tasks (Pellizzer & Georgopoulos 1993).

1.3.2 The "dynamical systems" hypothesis of M1 activity

Although the temporal patterning of single unit firing is especially obvious when certain task parameters have been dissociated from one another, encoding patterns may also be non-stationary during normal behavior (Fu et al 1995; Churchland et al 2007; Wang et al 2007; Churchland et al 2010). Some investigators have interpreted non-stationarity as an indication that parameter coding is not a valid or useful descriptor of firing rates. In a recent controversial study, Churchland et al (2012) proposed an alternative framework that treats the activity of dozens or hundreds of neurons as a certain class of dynamical system in which the activity pattern of the population at one moment in time follows lawfully from the previous pattern. They argued the intrinsic properties of the neuronal network dictate a set of evolution rules, something akin to a central pattern generator that limits the need for input-driven changes in activity. In this way, neurons need not "represent" anything; their patterning instead serves to

directly generate movement and any observed correlation with task parameters is incidental (cf. Fetz 1992).

To substantiate their dynamical systems framework, Churchland et al (2012) developed a novel dimensionality-reduction method that identified a simple evolution rule for M1 activity during reaching. Their method, called jPCA, is a variant of principal component analysis (PCA), which maps observations along many dimensions (i.e. firing rates of many neurons) to a new set of summary dimensions (principal components). When the neuronal data were viewed from their carefully-chosen summary dimensions, the pattern of population activity appeared to rotate smoothly away from its starting point, and activity during each movement type traced out a unique rotational path. In a two-dimensional subspace (for example), a rotational pattern is achieved when the observations along the two dimensions are oscillatory and out of phase. Out-of-phase oscillations along the top two jPCA dimensions were found to explain a moderate amount of variance in M1 population activity, and these patterns were hypothesized to form a flexible set of bases from which muscle activity could be constructed. The authors proposed that rotational or oscillatory dynamics are the dominant mechanism by which M1 contributes to movement. They further suggested this framework may be a more appropriate account of M1 activity patterns than classical "representational" coding.

Although the work of Churchland et al (2012) marked a valuable shift in analytical and visualization methods for neuronal data, there are several important criticisms of their approach, the treatment of inputs to the system, and the focus on rotational dynamics. One study, for example, showed that robust rotational patterns could be found by applying jPCA to a simulated population of velocity-tuned neurons, with variable response onsets similar to real data (Michaels, Dann, & Scherberger 2016). The authors concluded that "jPCA alone is not

sufficient to distinguish between a representational model with lags and the dynamical model proposed by Churchland et al." It is important to note, however, that this study also identified differences between a simple velocity-tuned population and real neuronal data during a reaching task. Indeed, neuronal responses during reaching often do not simply rise and fall with movement velocity but seem to have additional response features not explained by a single tuning model (e.g. Churchland & Shenoy 2007).

1.3.3 Feedback-related responses in M1

The dynamical systems hypothesis has emphasized the dominance of intrinsic "evolution rules" that dictate temporal patterning. In this view, the motor cortex is a "dynamical machine" (Churchland et al 2010) that relies on local recurrent circuits to operate without the influence of external input (Churchland et al 2012; Russo et al 2018). However, it is well known that M1 neurons respond to somatic input when the limb is moved passively, and responses are modulated by muscle stretch and tactile stimulation (Lemon & Porter 1976; Wolpaw 1980). Moreover, small perturbations applied to the arm during movement modulate M1 neurons at a short-latency (Evarts & Fromm 1977), suggesting somatic feedback can alter response dynamics.

More recently, it has been shown that M1 activity may also carry higher-order information related to movement feedback. Inoue, Uchimura, & Kitazawa (2016) recorded activity from M1 neurons while monkeys made reaching movements to visual targets observed through wedge prisms. The prisms were adjusted on each trial to randomly change the amount of visual error, which caused the monkeys to make end-point errors of varying magnitude. To prevent online corrective movements that would have countered the prism perturbation, the movement was hidden from view using a liquid crystal shutter. The shutter was opened

immediately after the movement to reveal the error of the end-point relative to the visual target. Using an information-theoretic analysis, it was shown that the firing rates of many neurons, measured shortly after the shutter was opened, shared mutual information with trial-to-trial end-point errors. This finding suggested that visual feedback could modulate M1 encoding.

M1 activity related to visual feedback has also been reported during ongoing behavior (Stavisky et al 2017). Firing rates of single units were recorded during a reaching task in which the hand was not directly observed, but was instead represented by a cursor on a computer display. During some trials, cursor position was displaced from its trajectory mid-reach, which required monkeys to make a corrective movement that countered the displacement and moved the cursor back toward the target. Firing rates were analyzed using dimensionality-reduction methods, and population activity was found to reflect the perturbation in two separate phases. The first phase seemed to correspond to short-latency visuomotor feedback related to the unexpected cursor displacement, while the second phase was related to the subsequent corrective movement. These results suggested that visual feedback could alter ongoing dynamics in M1 during behavior, at odds with recent work suggesting M1 operates independently of external inputs over extended periods of time (Russo et al 2018).

1.3.4 Discrete state changes in M1

Though M1 activity undoubtedly evolves in time during movement, there is considerable evidence that firing rates do not necessarily change smoothly as predicted by the rotational dynamics proposed by Churchland et al (2012). For example, Abeles et al (1995) measured the rates of small populations of M1 neurons while monkeys performed different reaching tasks and found that firing rates transitioned abruptly between "quasi-stationary states." Hidden Markov

models (HMM), which are used to model processes that transition between discrete states, were found to accurately describe the observed activity patterns. State transitions identified by the HMMs corresponded to transitions in the monkeys' behavior. Further, it was shown that the cross-correlation patterns of pairs of responses varied considerably when computed using data recorded during different states, suggesting rapid network reorganization (cf. Elsayed et al 2016). The authors concluded that "neural networks in the brain dwell most of the time in stable configurations of activity ('attractors' or 'states'), each having distinct firing rates and neuronal interactions."

The results of Abeles et al (1995) suggested that distinct network states are associated with distinct behavioral states; transitions in behavior could be recognized in the transitions of activity patterning. This type of rapid activity change was also observed during the "context-recall" task described above (Pellizzer, Sargent, & Georgopoulos 1995). In fact, rapid change in M1 patterning associated with behavioral transitions has been observed in several different contexts: when monkeys transition from planning to executing a movement (Elsayed et al 2016; Lara et al 2018); from actively moving to idly sitting (Velliste et al 2014); from reaching to grasping (Rouse & Schieber 2016b); and from accelerating to decelerating the arm (Harpaz et al 2018). A smooth rotational model of M1 rates (e.g. Churchland et al 2012) may therefore obscure underlying single-unit dynamics.

1.4 BEHAVIORAL STRUCTURE OF REACHING MOVEMENTS

1.4.1 Two-component model of reaching

Although M1 firing is temporally complex, transitions in neuronal firing patterns are correlated with transitions in behavioral state (Abeles et al 1995; Vaadia et al 1995; Shmiel et al 2006; Velliste et al 2014; Elsayed et al 2016; Rouse & Schieber 2016b; Lara et al 2018; Harpaz et al 2018). Therefore, in establishing a relation between M1 firing and reaching, it is pertinent to examine the temporal structure of reaching movements. Classical work by Woodworth (1899) established that point-to-point movements consist mainly of two components. Subjects made accurate drawing movements with a pencil on a roll of paper attached to a revolving drum. Woodworth was then able to evaluate detailed records of the time-course of movements. In some experiments, subjects accurately terminated their movements at a small dot or line. The initial portion of trajectories brought the hand close to the end point, and small corrective movements were made to accurately acquire the target. While the initial displacement was rapid and stereotyped across trials, the target homing phase was slower and more variable. He therefore speculated that these movements engaged two separate processes in sequence. Woodworth specified the movement time with a metronome and found that target homing was only effective for movements of sufficient duration. He concluded that "rapid movement does not allow time enough for the later [trajectory] adjustments. The later adjustments are reactions to stimuli set up by the movement, and a rapid movement does not allow for the reaction-time." In support of his two-component model of limb control, Woodworth found that very fast movements were about as inaccurate as those made with the eyes closed.

Woodworth's observations have led to an enduring perspective that the initial part of the movement is ballistic– a "pre-programmed" displacement unfolds for some time before sensory information can be processed meaningfully. Neurophysiological support for ballistic movement components has been described by Evarts & Fromm (1977), who recorded M1 firing rates in monkeys and applied small mechanical perturbations during movement. The perturbations modulated firing rates only for small precise movements but not for large rapid movements, indicating that proprioceptive feedback is also probably limited during "initial displacements" of the arm. These findings imply that the structure of reaching movements might be recognizable in the central motor system.

1.4.2 Reaching without vision

In Woodworth's (1899) original experiments, the contribution of vision to reaching was studied by instructing subjects to close their eyes before movement. On eyes-closed trials, the end-point target could not be acquired very accurately: without visual feedback, "we have no exact knowledge of where the goal is, and so cannot use our finer adjustments" to acquire the target (Woodworth 1899). Keele & Posner (1968) instead used electronically-controlled lighting to randomly eliminate visual feedback without the subject's prior knowledge. They attempted to refine Woodworth's estimate of the amount of time needed to process feedback and use it to improve the accuracy of their movements. They found that movements more rapid than 190 ms were just as accurate in the light and dark, whereas movements at least 260 ms long were more accurate with the lights on. Therefore, visual processing time was likely somewhere between 190 and 260 ms. Subsequent experiments demonstrated that prior knowledge of whether feedback would be available can influence visual processing time (Zelaznik, Hawkins, &

Kisselburgh 1983), though ~200 ms seems a reasonable estimate in the case that subjects are uncertain.

Meyer et al (1988) studied the composition of point-to-point movements with and without feedback using a manipulandum-controlled cursor that could be hidden trial-to-trial. The initial ("primary") sub-movements were comparable regardless of feedback, in line with the notion of a "pre-programmed" displacement. Interestingly, they observed secondary sub-movements even when the cursor was disabled, and those sub-movements were of similar frequency and duration as those made with cursor feedback. This suggested that multi-component reaches were the result of a core strategy adopted during precise movement. However, the accuracy of secondary sub-movements was clearly diminished without feedback, suggesting they occurred as part of a process that was distinct from the "pre-programmed" primary sub-movement, and instead "stemmed from active processing of visual feedback." It is therefore likely that these two processes are associated with distinct patterns of neural activity appropriate for each phase of the movement.

1.5 SUMMARY

The primary motor cortex is anatomically and functionally implicated in motor control, though its specific role is debated. M1 activity correlates with a very large number of behavioral parameters, and firing patterns are temporally complex. Given the diverse and seemingly contradictory nature of M1 coding, what conclusions can be drawn about the function of this cortical region? How should we interpret patterned firing rates during behavior? Fetz (1992) has likened the effort to "reading tea leaves." Yet, there are clear principles that may help parse

activity patterns. For example, cosine-like directional tuning is a robust feature of M1 firing, even if this tuning relates to diverse behavioral parameters. Numerous studies have leveraged this tuning property, along with dissociation tasks, to determine information content carried in M1 during behavior. Firing patterns are often complex, but encoding changes seem to align with distinct components of behavior. Johnson & Ebner (2000) comment that, while the "encoding of multiple parameters in a motor cortical neuron is ambiguous… these apparent ambiguities are dealt with by a temporal parcellation scheme in which different parameters are signaled at different times in the task." Their "temporal parcellation scheme" posits that "the various signals and coordinate systems [recognized in M1] vary in time and evolve throughout a behavioral sequence, consistent with the demands of the task and the required sensorimotor transformations" (Johnson, Mason, & Ebner 2001). This framework may be good compromise between the various descriptions and accounts of M1 activity and function.

2.0 TEMPORALLY SEGMENTED DIRECTIONALITY IN THE MOTOR CORTEX[1]

2.1 SUMMARY

Developing models of the dynamic and complex patterns of information processing that take place during behavior is a major thrust of systems neuroscience. An underlying assumption of many models is that the same set of rules applies across different conditions. This has been the case for directional tuning during volitional movement; a single cosine function has been remarkably robust for describing the encoding of movement direction in different types of neurons, in many locations of the nervous system, and even across species. However, detailed examination of the tuning time course in motor cortex suggests that direction coding may be labile. Here we show that there are discrete time epochs within single reaches, between which individual neurons change their tuning. Our findings suggest that motor cortical activity patterns may reflect consistent changes in the state of the control system during center-out reaching. These transitions are likely linked to different behavioral components, suggesting that the task defines changes in the operational structure of the control system.

[1] This chapter was published as: S.B. Suway, J. Orellana, A.J.C. McMorland, G.W. Fraser, Z. Liu, M. Velliste, S.M. Chase, R.E. Kass, A.B. Schwartz. Temporally Segmented Directionality in the Motor Cortex. Cerebral Cortex 28: 2326-2339, 2018 (Advance access June 2017). Reproduced with permission of Oxford University Press.

2.2 INTRODUCTION

Understanding and modeling the dynamic activation of neural ensembles is a major goal of systems neurophysiology. The pioneering experiments of Georgopoulos et al. (1982) in which monkeys reached to radial targets from a central start position showed that the tuning relation between single-unit neural firing rates in the motor cortex and movement direction could be fit with a broad cosine tuning function that spanned all movement directions. Direction was shown to be "encoded" in a way that was determined by a neuron's "preferred direction" (direction of peak firing). These features were used in the population vector algorithm to extract movement direction from neural activity (Georgopoulos et al., 1983). Cosine tuning has since been shown to be robust for many paradigms and characterizes unit activity recorded in structures throughout the neural axis during movement (Schwartz, 1994; van Hemmen and Schwartz, 2008). This type of broad tuning has been expanded to multiple parameters describing the motion of the arm, wrist and fingers, and, as a general descriptor of upper limb dynamics, forms the basis for brain-computer interfaces in which recorded neural activity drives the movement of external devices (Wessberg et al., 2000; Serruya et al., 2002; Taylor, Helms Tillery, and Schwartz, 2002; Hochberg et al., 2006; Wodlinger et al., 2015). However, evidence is accumulating that a cell's preferred direction may be unstable (Churchland and Shenoy, 2007; Sergio and Kalaska, 1998; Suminski et al., 2015; Hatsopoulos, Xu, and Amit, 2007). Directional tuning may even change within the course of a single reaching movement. Because of these rapid changes in a neuron's directional sensitivity, the canonical aspect of a cell's preferred direction and the meaning of directional tuning have been called into question.

We carefully examined changes of direction representation in motor cortical discharge as reaching takes place. In simulation, we determined that complex firing rate patterns and apparent

tuning lability can result from neuronal tuning to multiple kinematic variables. However, even though adding more kinematic parameters in the simulation could lead to tuning lability, it did not account for the within-trial changes observed in the actual data. Instead, we found discrete, consistent changes between three well-defined epochs across the population of neurons. As we will show, many neurons display stable preferred directions within each of these epochs. During the sharp transitions between tuning epochs, modulation dropped to a minimum. These transitions corresponded to different phases of the reach, suggesting that the tuning changes were driven by the behavioral structure of the task.

The tuning functions generated in the original center-out experiments compared the *mean* discharge rate, a single number calculated across the duration of movement, to a single movement direction from the center position to the target. When expanding these findings to account for task dynamics, the simplest assumption would be that the model was stationary—the same model would be valid at each time point through the trial. If the hand moves in a straight line from the center position to the radial target, direction-induced neural modulation should be constant during the trial. According to the cosine tuning model, if direction is the only factor governing the firing rate of a neuron, a profile of firing rate versus time, throughout a movement, should be flat (Figure 2.1A).

The problem can be formalized with a simple planar model used in the original 2D center-out task:

$$\lambda = \beta_0 + \beta_x D_x + \beta_y D_y \tag{2.1}$$

where λ is the firing rate of a cell, D_x, D_y are coordinates of a unit vector pointing in the movement direction, β_0 is baseline firing rate, β_x and β_y are the coordinates of a vector in the

cell's preferred direction with modulation depth $m = \sqrt{\beta_x^2 + \beta_y^2}$. If D_x and D_y are constant (straight movement), then the neuron's firing rate should remain unchanged throughout the trial.

Although the basic finding of cosine tuning has been confirmed repeatedly, motor cortical firing during reaching movements is rarely constant. This suggests that either a neuron's discharge rate is governed by additional, non-directional factors (requiring an extension of the original cosine-tuning model) or that the directional specificity of each neuron changes during the reach (with the possibility that direction coding in general may be invalid). Trial-averaged firing rate profiles of motor cortical firing rates are almost always phasic, with peaks and valleys at different points within the reach, resembling the examples shown in Figure 2.1B or C. This shows that simple direction encoding by these neurons is not, by itself, an adequate description of motor cortical function, motivating investigators to consider the effect of movement variables other than direction on firing rate. The simplest of these are position and its successive derivatives, velocity and acceleration (Ashe and Georgopoulos, 1994). In single neurons, interactions between direction and these additional parameters can produce complex changes in firing rate with temporal instability in both the amplitudes and preferred directions of directional tuning. In this paper, we first confirm the temporal instability of directional tuning during single reaches (Figure 2.3), then show that extending the model with non-directional parameters does not account for this observation (Supplemental Figure 2.11). Finally, we describe distinct phases of the reach, defined by each neuron's directional modulation, in which tuning is robust and stable (Figures 2.4, 2.5). This indicates that direction is a major determinant of firing rate, but suggests the existence of distinct states within a single reach. The step-changes in preferred direction clearly evident in our results (Figure 2.4E) signify discrete segmentation of neural processing during reaching.

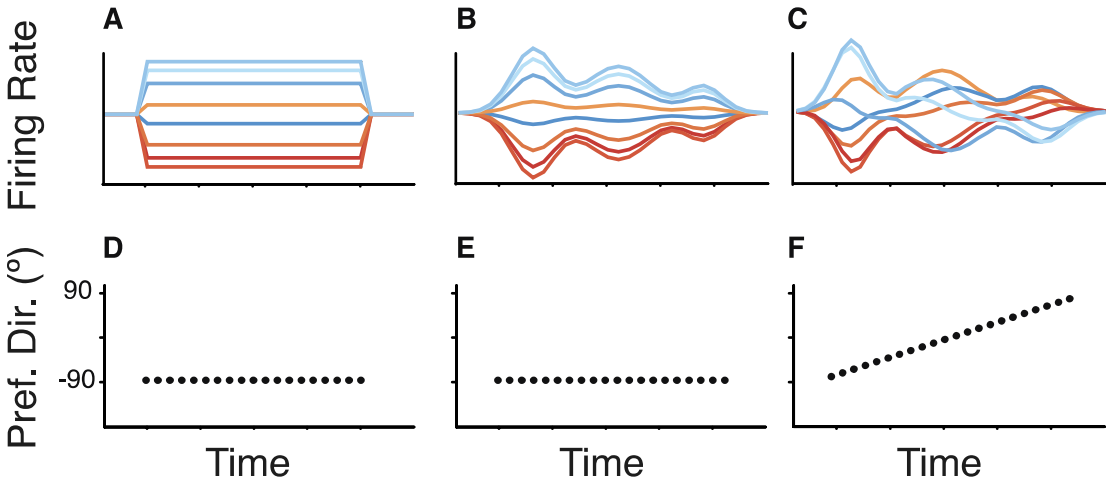

Figure 2.1 Simulated examples of tuning vs time

Each example is composed of 8 straight, center-out, point-to-point reaches. The top row is the firing rate time course of the simulated neuron, with colored traces corresponding to movement in a particular direction. The bottom row shows the preferred direction calculated in separate bins throughout the trial. A,D. A canonical neuron that behaves according to the classic cosine model, in which direction is the only determinant of firing rate. Assuming a straight movement, the firing rate would be a step-function and constant during the movement. The preferred directions of tuning functions calculated in small windows throughout the trial are constant across windows. B,E. A more typical firing rate pattern in which the neuron shows several peaks of firing during reaching. In this case, the entire length of the profile is scaled uniformly by direction and the preferred direction is stationary. C,F. A neuron with varying modulation patterns during the trial. Neural modulation again has multiple components during the reach, but firing rate is both direction- and time- dependent. At each point in time, direction has a different effect on firing rate. In this case, the preferred direction changes continually through the trial.

2.3 MATERIALS AND METHODS

2.3.1 Behavioral task, neural recording, and neural data processing

Rhesus monkeys (*Macaca mulatta*, male, 3 animals) were seated in a primate chair, with one arm restrained. An infrared marker was attached to the free hand and three-dimensional position was monitored at 60 Hz using an Optotrak 3020 motion capture device (Northern Digital Inc, Waterloo, Ontario, Canada). The monkeys could not see their moving arm, but instead observed a computer monitor showing a virtual reality (VR) environment where the position of the hand was represented by a spherical cursor. The scene was displayed in three dimensions using a depth-displaying monitor (Virtual Window, Dimension Technologies Inc, Rochester, NY, USA). The monkeys were trained to perform a center-out reaching task in which they had to move the cursor from a starting location toward radial targets equidistant from the center. Workspace radii were 7.4, 6.5, and 8.0 cm for monkeys C, F, and N, respectively. Monkeys C and N performed a version of the task with targets arranged in a 2D plane, while monkey F performed the task with a 3D arrangement of targets (Figure 2.2). Monkeys C and F were required to hold the position of the target for a few hundred milliseconds. Monkey N wasn't required to hold, but instead was allowed to return back to the home position immediately. All movement conditions were presented in a pseudo-random fashion: failed trials were returned to the queue until each set was completed. Monkey C performed 47 repetitions to each of 16 targets; monkey F performed 26 repetitions to each of 26 targets; monkey N performed 40 repetitions to each of 58 targets. Additional information about data sets is detailed in Section 2.7.2.1.

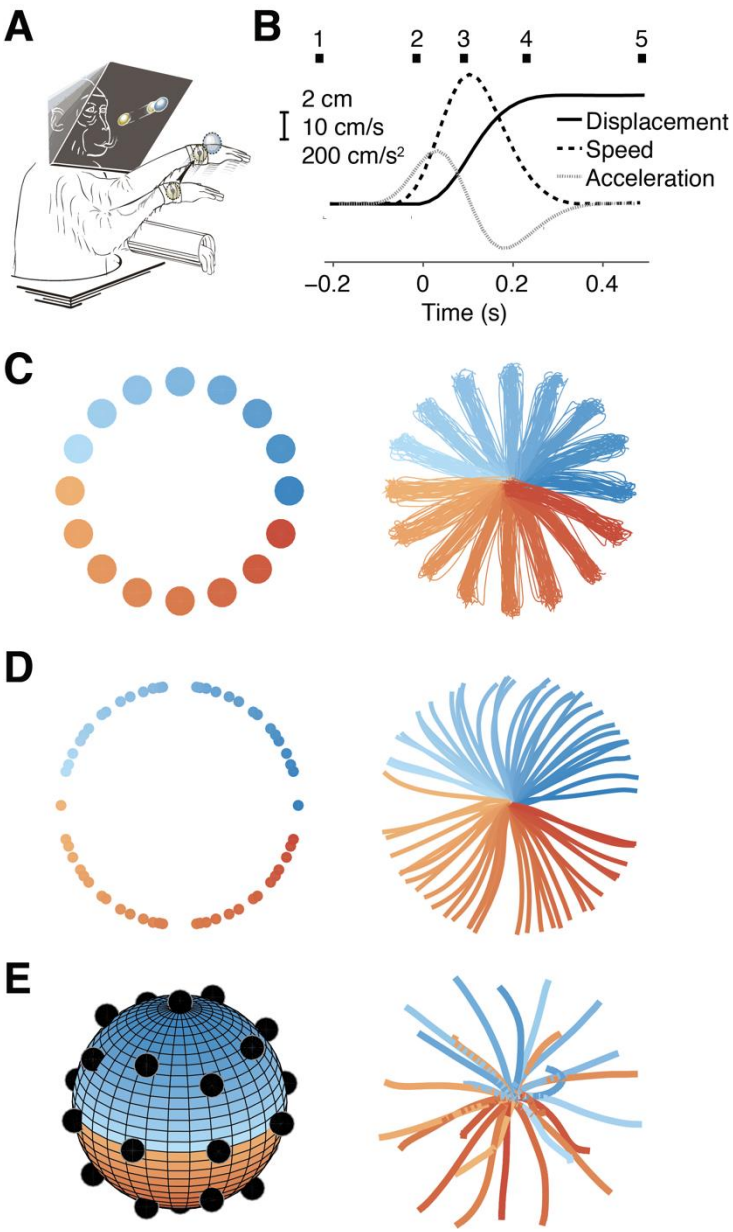

Figure 2.2 Experimental design and kinematic results

A. Monkeys reached for radial targets displayed in a 3D computer monitor. **B**. Representative kinematics from a single monkey C experiment. 1 Target display; 2 Movement onset; 3 Peak velocity; 4 End of movement; 5 Reward administration. **C**. Two-dimensional target arrangement, monkey C. Left- targets individually displayed in virtual reality. Right- Trajectories of the hand collected during one experiment, with all repeated movements shown (47 per target). **D**. Two-dimensional target arrangement, monkey N. **E**. Three-dimensional target arrangement, monkey F. Trial-averaged hand trajectories are shown for monkeys N and F. In panels **C-E**, the color of the targets corresponds to that of the trajectories. See also Supplemental Figure 2.9 for electrode array placement.

We determined five task-related and kinematic events for each trial: target show, movement onset, peak velocity, movement offset, and the end of a hold period (“hold off”). Movement onset and offset times were defined as the point when hand speed passed 15% of its maximum. Electrical activity from single neurons was recorded extracellularly using chronically-implanted 96-channel ‘Utah’ electrode arrays (UEAs; Blackrock Microsystems, Salt Lake City, UT, USA) implanted approximately in the arm area of the pre-central gyrus (for locations see Supplemental Figure 2.9).

We recorded 93, 119, 185 single units for C, F, and N. For monkeys F and N, waveform snippets were stored for each threshold crossing and neuronal spikes were sorted offline using "Offline Sorter" (Plexon Inc, Dallas, TX, USA). For monkey C, single units were sorted manually online, with sorts adjusted periodically to maintain isolation as needed. Neural data recorded from monkeys N and C were collected over one day. For monkey F, neural data were collected over five consecutive recording days. To average across days, units were identified as being the same using multi-day unit identity analysis (Fraser and Schwartz, 2012). This was accomplished by calculating a similarity score derived from four independent metrics: mean firing rate, spike waveform, auto-correlation, and cross-correlation with the rest of the population (see also Section 2.7.2.2).

Spike trains from each unit were binned in 20 ms intervals to compute instantaneous firing rate. A Gaussian kernel with standard deviation of 50 ms was used for smoothing. The hand position values were resampled and interpolated to match the time steps of the firing rates using the function “pchip” in Matlab. On each trial, we took movement onset to be time zero and then aligned across trials to additional landmarks at target show, peak velocity, movement offset, and, for monkeys C and F, hold off. This was accomplished by choosing the number of bins

between each landmark so that the average bin size was 20ms, and counting the spikes within each bin. Thus, although a single bin spans 20ms on average, it may span more or less time on a given trial. Because time was normalized independently between each alignment point, each segment of the task was associated with different trial-to-trial variation in bin length. The standard deviation of the bin length during the most variable epoch across our data sets was 6 ms.

Lastly, we normalized firing rate by scaling its value between zero and one based on minimum and maximum firing rate, and subtracting the mean over targets (cross-condition mean) at each time step. Amplitude normalization is critical only for the population vector analyses (described below), but it is also convenient when comparing tuning characteristics between neurons with a different range of firing rates. Subtracting the cross-condition mean allowed us to focus on the dynamics of tuning, rather than dynamics shared across targets. All analyses described below used the time-normalized data. All analyses except those described in Supplemental Figure 2.11 used the mean-subtracted, amplitude normalized changes in firing rates. Analyses described for that figure relate firing rates to kinematic variables, which in principle may modulate spiking irrespective of target (e.g., scalar speed or displacement). These preprocessing steps were therefore not performed on data used in those analyses.

We sometimes observed that neurons would fire at rates that were unexpected, based on their typical trial-to-trial variability. If this happened in the same trial for many of the units in the recorded sample, all data for that trial were removed to improve the robustness of non-linear model fitting (described below). The firing rate profiles from individual units were analyzed as vectors in time. Each trial vector was correlated with all the other trials to that target (using dot products) and if the mean correlation between that trial and all others was below .2, it was

deemed to be of 'high variance'. This was repeated for each individual unit in that trial. If more than 84% (the percentile corresponding to mean+1 std on a normal distribution) of the neurons had high variance for that trial, the trial was excluded (for all neurons) from further analysis. Fewer than 15% of trials were excluded due to this criterion. Interestingly, the low firing rate correlation between a given trial and the mean firing rate was not strongly related to kinematic variability, suggesting that the high variability was due to noise, for example from movement artifacts in the neural recordings. Results of individual analyses did not depend critically on this exclusion, though the robustness of non-linear model fitting ("component fitting," see below) was improved.

2.3.2 Preferred direction stability test

Tuning functions were calculated using the firing rates in 20 ms increments to determine statistical changes in directional tuning for each unit. We estimated cosine tuning functions from the recorded neural data using least squares linear regression based on the following model:

$$y_i = \beta_0 + \beta_1 \cos(\theta_{targ} - \theta_{PD}) + \varepsilon_i \quad (2.2)$$

where $y_i = \hat{\lambda}_i$ is the single-trial estimate of a neuron's firing rate, β_0 is baseline rate, $\theta_{targ} - \theta_{PD}$ is the angle between the target direction and the neuron's preferred direction (PD), and ε_i is the noise (or error) representing the deviation from cosine tuning. We define $\hat{m} = \beta_1$ as the estimated magnitude or modulation depth of the tuning function. The extension to the 3D target arrangement involves expressing the $\theta_{targ} - \theta_{PD}$ covariate as the dot product of two unit-length 3D vectors. Two versions of this model were constructed assuming either that PDs are fixed in time (Eq. 2.3) or that PDs change over time (Eq. 2.4). We reasoned that if PDs were actually constant, allowing the fitted PD to vary with time would not significantly improve the likelihood

of the model. Furthermore, the variation in firing rate between time bins would be sufficiently described by a change in the modulation depth of the stable tuning function. We thus formulated a "restricted" model with a static PD such that:

$$y_{i,t} = \beta_{0,t} + \beta_{1,t} \cos(\theta_{targ} - \theta_{PD}) + \varepsilon_{i,t} \quad (2.3)$$

where *t* denotes time. We then formulated an "unrestricted" model with time-varying PD such that:

$$y_{i,t} = \beta_{0,t} + \beta_{1,t} \cos(\theta_{targ} - \theta_{PD_t}) + \varepsilon_{i,t} \quad (2.4)$$

After fitting, we then summed the log-likelihoods of the model fits at each time bin. To compare the models' relative goodness-of-fit, we performed a likelihood-ratio test. The test statistic is given as:

D = -2ln(likelihood of restricted model) + 2ln(likelihood of unrestricted model) (2.5)

The null hypothesis of the test represents the case that the restricted (i.e. constant PD) model is as likely as the unrestricted (i.e. time-variant PD) model, given the data. The value of D under the null hypothesis is assumed to follow a chi-squared distribution with degrees of freedom equal to the difference in degrees of freedom between the two models. The restricted model had 1+T degrees of freedom in the 2D target case, or 2+T in the 3D case: one (or two) from θ_{PD}, and T from the $\beta_{1,t}$ parameters, where T is the number of time steps in the testing period. The unrestricted model had 2T degrees of freedom in the 2D target case, or 3T in the 3D case: T (or 2T) from θ_{PDt} and T from the $\beta_{1,t}$ parameters. For the likelihood ratio test, this gives T – 1 dfs for the 2D target case, or 2T – 2 dfs for the 3D target case. We failed to reject the null hypothesis of "PD stability" if the *p*-value of this test exceeded (0.05 / number of time bins considered). To reduce the influence of noise, time bins for which a neuron was not significantly modulated by target direction (one-way ANOVA, $p > .001$) were excluded from the analysis. This is important

because it allows us to be sure that a result of "stable PD" cannot be attributed to a lack of tuning; a model with changing PD does not improve likelihood over a model with static PD if there is no task-related modulation of firing.

We note that θ_{targ} in Equations 2.2 – 2.4 represents the target direction (constant throughout a trial), rather than the instantaneous direction of the hand end-point. This is in contrast to the methods utilized in Supplemental Figures 2.10 and 2.11, which used the instantaneous hand direction. We found there was consistent modulation of many units after the hand had stopped, which is not expected for a model based solely on instantaneous hand direction (see, for example, Supplemental Figure 2.10B, left subpanel). Using the true hand direction was therefore only valid for part of the trial. Instead, cosine tuning to target direction was consistently found to be robust at most time points in the trial (Figure 2.3B).

2.3.3 Tuning reliability over time

Trial-to-trial reliability of preferred direction estimates (Eq. 2.2) at each 20 ms time bin was assessed by bootstrap. We first found the PD within each time bin for each unit using the original sample of data. Firing rates for each unit were then sampled with replacement across trials 1000 times. At each time bin, a distribution of unsigned angles was formed by finding the arccosine of the dot product between the unit vectors representing the original PD estimate and each of the bootstrap estimates of PD (from the trial resampling). We then found the 95% confidence interval of these angles from the percentiles of this distribution. This interval is described for the population of neurons in Figure 2.3A.

2.3.4 Component fitting

Our 20 ms tuning analysis showed that preferred directions tended to change in discrete steps during the reach. These steps, or segments, had durations of 100-150 ms and we used a novel component-fitting procedure to identify each tuning epoch (Equation 2.6). Inspection of the sequential tuning functions along with firing rate profiles revealed that the modulation depths tended to rise and fall in roughly Gaussian-shaped peaks, or components. PDs were consistent for the duration of a component, but could change abruptly from one component to the next. We therefore developed a model of firing rates that captured the observed Gaussian-shaped modulation and step-like cosine tuning features (Figure 2.4 A, B, E). For convenience, we chose to fit the modulation depth profile with a multi-component Gaussian function. Although this general function shape fit our data well, other forms likely could be substituted.

This model (Equation 2.6) defines cosine-tuned Gaussian-shaped components, specified in part by peaks in the amplitude of the neuron's tuning function, observed at different points in time through the task. The components are further defined by a PD, which was assumed to be constant for the duration of the component (an assumption which was subsequently tested for each component). Because PDs were specified separately per component, the overall PD of the neuron was allowed to rapidly change from one component to the next. These features were simultaneously captured by the following expression, which we refer to as an epoch-specific tuning model:

$$y_i = \beta_0 + \sum_{j=1}^{J} a_j \exp\left\{\frac{-(t-\mu_j)^2}{2s_j^2}\right\} \cos\left(\theta_{targ} - \theta_{PD_j}\right) + \varepsilon_i \qquad (2.6)$$

Here, $y_i = \hat{\lambda}_i$ is the expected change in single-trial firing rate at each time *t*, given a target θ_{targ}. Baseline firing rate is denoted β_0. For each of the *J* components, the modulation

depth was fitted with a Gaussian temporal profile. Thus, the j^{th} width coefficient s_j specifies the standard deviation of the Gaussian profile, and the time of maximum firing rate (a_j) is given by μ_j. The term ε_i is the noise (or error) representing the deviation from the model. Once initialization parameters were obtained (see Section 2.7.2.3), final estimates of each parameter were fit using non-linear least squares regression for the full model given in Equation 2.6.

For neurons recorded from monkeys C and F, we often observed a period of tonic firing during the hold period. Monkey N was not required to hold at the target and neurons recorded during this task did not exhibit this activity. Although this activity was typically also cosine-tuned, we intentionally avoided fitting components here, as this activity was not believed to be related to the preceding movement but rather to the act of holding at the target (Kettner, Schwartz, & Georgopoulos, 1988). Examples of this activity can be seen at the end of the trial period for units 59a and 79a from monkey C (Figure 2.4 top row) and in Supplemental Figure 2.13.

2.3.5 Population vector algorithm (PVA) analysis

The population vector algorithm (PVA) was used to decode movement direction from motor cortical units. To assess the effect of tuning lability on our ability to decode movement, we reconstructed movement trajectories with three different formulations of the population vector.

A) PVA using initial preferred directions. We found that tuning becomes maximally reliable across the population of cells shortly after target onset (Figure 2.3 A,B). In our first formulation of the PVA, we therefore computed a single PD for each cell using the average rates in a 100ms window of initial tuning. To account for small variations in tuning onset, and to ensure robust

cosine fits, the windows were determined for each cell by finding the first 100ms of significantly modulated activity. Population vectors were then computed for each trial using the equation:

$$P_t = \sum_i^N w_{it} C_i \quad (2.7)$$

where P_t is the population vector at time t on a given trial, w_{it} is the amplitude-normalized firing rate of the i^{th} unit at time t, and C_i is the single preferred direction of the i^{th} unit. To make direct comparisons between this analysis and a components-based analysis, units were included only if at least 1 component was successfully fitted (see Figure 2.5). The time series of population vectors was integrated and averaged over task repetitions to construct a movement trajectory for each reach direction.

B) Components-based PVA. Our second formulation of the PVA addressed whether PDs determined by each component in the epoch-specific tuning model would improve trajectory reconstruction. In the epoch-specific tuning model (Eq. 2.6), the PD of a single cell is constant during a tuning component but transitions rapidly to a new value when the component switches. We computed these component-determined PDs at each time bin for each cell using the epoch-specific tuning model fits. Population vectors were then computed using:

$$P_t = \sum_i^N w_{it} C_{it} \quad (2.8)$$

Here, the preferred direction C of each unit, i, is a function of time, t.

C) Bin-by-bin PVA. Finally, to determine whether PD lability at a temporal resolution higher than that of the components would further improve decoding, we computed a third version of the PVA using Equation 2.8, but, unlike the components-based method, the PDs were computed independently in each bin.

In each of these three cases, we quantified the average length of the population vectors and their angular deviation from a straight path to the target. The quality of decoding was summarized by first obtaining the length and angular deviation of the population vectors for each trial and at each time bin. All such values were then averaged. The variability in these estimates was computed via bootstrap by drawing trials randomly with replacement and re-calculating the measures using firing rates from the resampled trials. This process was repeated 1000 times.

2.3.6 Offline component recognition

In order to test the feasibility of using the component-wise preferred directions in a real-time population decoder, we investigated the possibility that patterns of population activity could be used to identify the tuning component driving a single unit at any given time point. This would be useful for real-time tracking of each neuron's "current" preferred direction during online decoding experiments, based on the concurrent activity of each recorded neuron. We trained classifiers (Linear Discriminant Analysis, LDA), one separately for each unit, using the firing rates from every unit in the recorded population. That is, a feature vector at a given trial and time bin was constructed by gathering the rates of every neuron during that bin. For each unit's classifier, class labels were defined by the previously identified tuning segments, as delimited by the intersections of the adjacent components fit for that unit. The intersections were taken to be the point at which one component's amplitude became larger than the previous component's. These labels were used to train the linear discriminant classifier for that unit. Once trained, the classifier utilized the feature vector of population rates at each 20ms sample in each trial to

predict the "current" segment of the neuron's response (and thus its PD at that time). Classification was performed only for units with at least two components.

2.3.7 Additional data collection

Data from a fourth monkey (*Macaca mulatta*, male, monkey P) were collected to clarify the effect of target acquisition strategy on the observed pattern of component timing. Monkey P was trained to distinguish the stop-and-hold requirement by the color of the target, which was chosen at random for each trial. This monkey performed the task for several months prior to collecting these data. Behavioral setup, training, and array placement were comparable to those for the other three monkeys. Neural data were acquired and sorted using the procedures described for monkey N. Data from monkey P were pooled over three days of recording and included a total of 139 neurons (average 46.3 per day).

2.4 RESULTS

2.4.1 Expanded models do not account for preferred direction lability

Data were collected as monkeys performed center-out tasks using target arrangements in 2 and 3 dimensions (Figure 2.2). Speed profiles over each movement were bell-shaped, displacement increased monotonically and acceleration profiles were biphasic, showing that these movements are consistent with smooth point-to-point movements typical of normal reaching (Figure 2.2B). To determine whether directional tuning changed within the task, we divided each trial into 20

ms time bins, and calculated local tuning functions for individual neurons in each bin. A preferred direction stability test (see Section 2.3.2) was used to determine whether local preferred directions, calculated in each bin, differed from the mean preferred direction taken from the period 100 ms prior to movement until the end of the movement. By this criterion, only 21%, 8%, and 16% of the recorded units in monkeys C, F, and N, respectively, were found to have tuning functions that were stable throughout the task.

This temporal instability could result from a statistical interaction between direction and parameters that take on varying values throughout the task. Displacement, speed, and acceleration have characteristic temporal profiles during reaching (Figure 2.2B). To explore the possibility that adding a combination of these kinematic terms to the regression equation could account for the observed instability of preferred direction, we built simulations of firing rate using a mixture of kinematic variables. The resulting firing rate profiles were complex and resulted in tuning function instability through the trial (Supplemental Figure 2.10). This possibility was tested on actual data, by using the firing rates of each unit in 16 different generalized linear models consisting of different combinations of direction, position, speed, velocity, and acceleration (Supplemental Figure 2.11A). Although specific parameters had a slight effect on model fitness, for the most part, adding additional parameters to the basic direction-only regression had little effect on the explanatory power of the model when used on actual data, as shown by the consistent goodness-of-fit (Supplemental Figure 2.11B). When assessed for temporal lability in the preferred directions (Supplemental Figure 2.11C) none of the models led to a decrease in the magnitude of preferred direction changes. Furthermore, accounting for additional kinematic terms had little effect on our statistical assessments of tuning stability for each individual neuron (Supplemental Figure 2.11D).

The large majority of deviations from stability in our data were small (20-40 degrees), but significant. Although our simulations showed that including a mixture of kinematic parameters could lead to directional instability, this analysis, performed on the collected data, suggests that simply expanding the basic model with more terms cannot explain the observed changes in preferred direction.

2.4.2 Cosine tuning is robust throughout trials, even in the presence of preferred direction lability

We characterized tuning function stability during the task with a number of analyses. Trial-to-trial directionality in each bin was examined by selecting randomly (with replacement) firing rates from different trials and calculating preferred directions repeatedly in a bootstrap procedure (Figure 2.3A). Early in the trial, preferred direction estimates for each neuron were highly unreliable (large confidence intervals), with a rapid increase in reliability about 100 ms before movement onset. A similar trend was found for the cosine tuning regression fit in each bin, where R^2 increased steeply just before movement onset (Figure 2.3B). Instead of assessing directional tuning independently in separate bins, tuning stability across bins was assessed in 200 ms (10-bin) sliding windows with the preferred direction stability test (Figure 2.3C). The percentage of cells with constant preferred directions in the 200 ms window was relatively high (approx. 75%) early in the trial and then decreased to a minimum (approx. 25%) around movement onset, before increasing as the movement progressed. While directional tuning may not be robust early in the reach (large confidence intervals), the decrease in preferred direction stability near movement onset occurred when the data fit the cosine tuning model well and with

high reliability (Figure 2.3A,B). The instability of preferred directions by this point in the reach therefore cannot be explained by unreliable tuning or low R^2 of a cosine fit.

When testing for changes in preferred directions during the behavioral trials, it would appear that directional stability is related to separate portions of the task. For instance, if only the firing rates prior to movement onset are considered (-150 to -50 ms), 90%, 81%, and 74% of the units (monkeys C, F, and N) had stable preferred directions. However, immediately after movement onset (0 – 100 ms), only 35%, 72%, and 50% of units were stable, respectively.

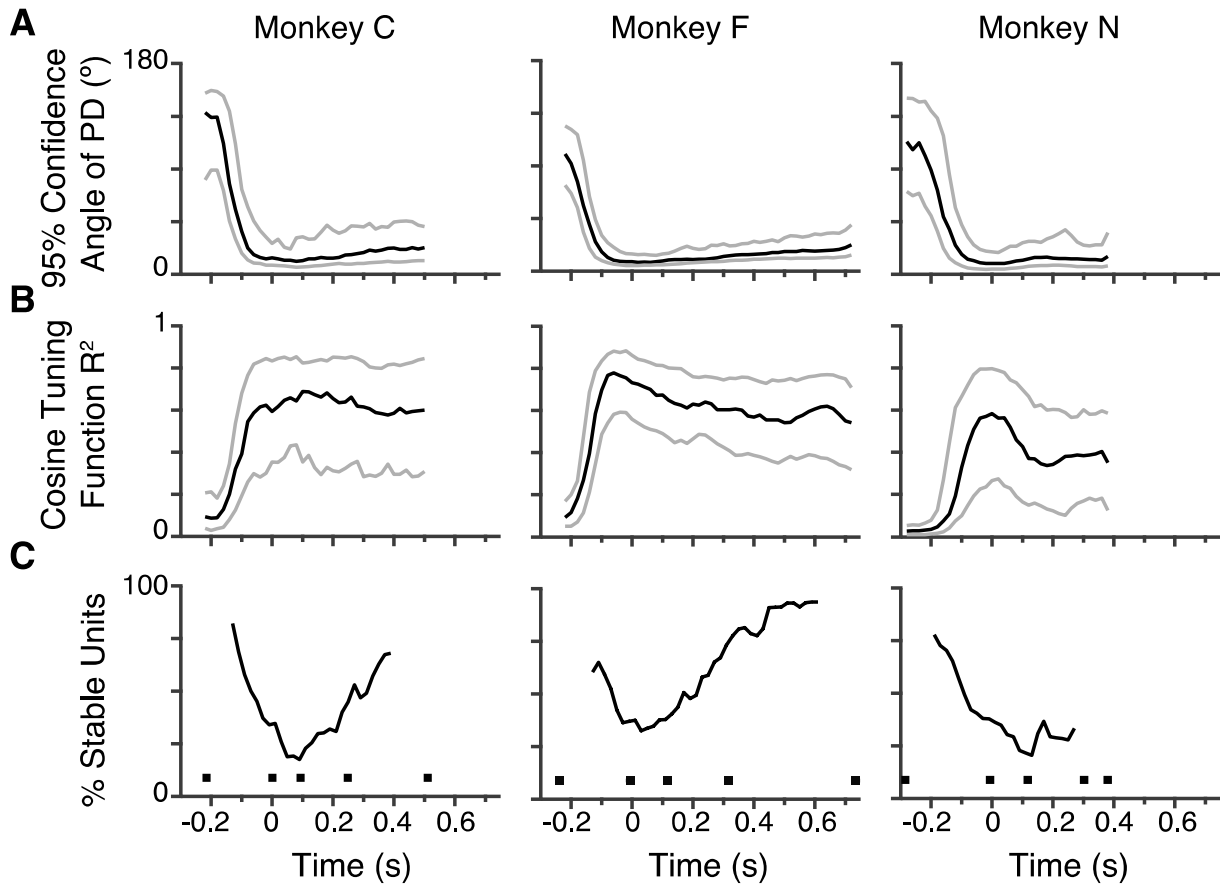

Figure 2.3 Stability vs time for the population

A. Reliability of preferred direction estimates. To determine trial-to-trial variability (width of confidence interval) for each neuron and at each time bin, PDs were repeatedly estimated using a bootstrap method (see Materials and Methods). Black lines show the population median for width of PD confidence intervals (95%) computed by the bootstrap; gray lines show quartiles (Q1, Q3) for the population. The distribution began narrowing 200 ms before movement onset, with tuning becoming maximally reliable for most neurons about 50 ms after the movement began. **B**. The same reliability trend was found when using R^2 for the cosine tuning regression. Tuning functions for each neuron were calculated in each 20 ms bin by averaging firing rates over task repetitions and fitting a cosine. Black and grey lines show the population median, Q1, and Q3 R^2 of the cosine fits to each neuronal tuning curve. **C.** Preferred direction stability was assessed within 200 ms overlapping windows that were incremented by 20 ms bins. The stability test, used previously for the whole trial, was then applied to each 200 ms window for each analyzed cell. To decrease the influence of noise, bins with poor tuning ($R^2 < 0.6$) were excluded. The percentage of units with stable preferred directions peaked 100 ms before movement onset, decreased to the beginning, and then gradually increased toward the end of the movement when the cursor was in the target. Square markers in panel **C** correspond to target show, movement onset, peak velocity, movement offset, and reward administration.

2.4.3 Episodic modulation determines changes in preferred direction

As shown in Figure 2.1, the temporal profile of firing rate can be a determinant of directional stability. Therefore, we analyzed the modulation patterns of the recorded activity for consistent features that may be related to directional tuning. We characterized direction-related modulation by using the amplitude of the tuning function (modulation depth) calculated in each 20 ms bin throughout the task. Tuning function amplitude is shown in the second row of Figure 2.4 for four different neurons. The neuron in the first column had one clear modulation peak, the neurons in the second and third columns had an early peak followed by a later peak in modulation, and the cell in the fourth column showed three distinct modulation components. For each neuron, we modeled this modulation with an epoch-specific tuning model (Equation 2.6) and the resulting Gaussian-shaped components are shown in the third row of Figure 2.4. Preferred directions calculated within the timespan of each component were found to be stable, but these directions changed rapidly between epochs (Figure 2.4, bottom row).

We next used principal components analysis (PCA) to find components of each unit's firing rate profile that covaried. This allowed us to examine prominent temporal features in the firing rate profiles without making assumptions about tuning curve shape (i.e. a cosine). Averaging across trials, we constructed a vector for each bin with components of firing rate to each target condition. We then used PCA on an T x T covariance matrix where T was the number of bins in the trial (this is sometimes referred to as "functional" PCA). The resulting eigenvectors are a function of time and represent orthogonal patterns of firing rate modulation taking place during the trials. Note that there is no consideration of directionality here, as this PCA was based solely on the variation of the observed firing rates, in contrast to the epoch-specific tuning model which was based explicitly on the typical cosine-tuning model. The two

approaches gave very similar results: the first few eigenvectors from our PCA captured variability that was described well by the epoch-specific tuning model (Figure 2.4, fourth row).

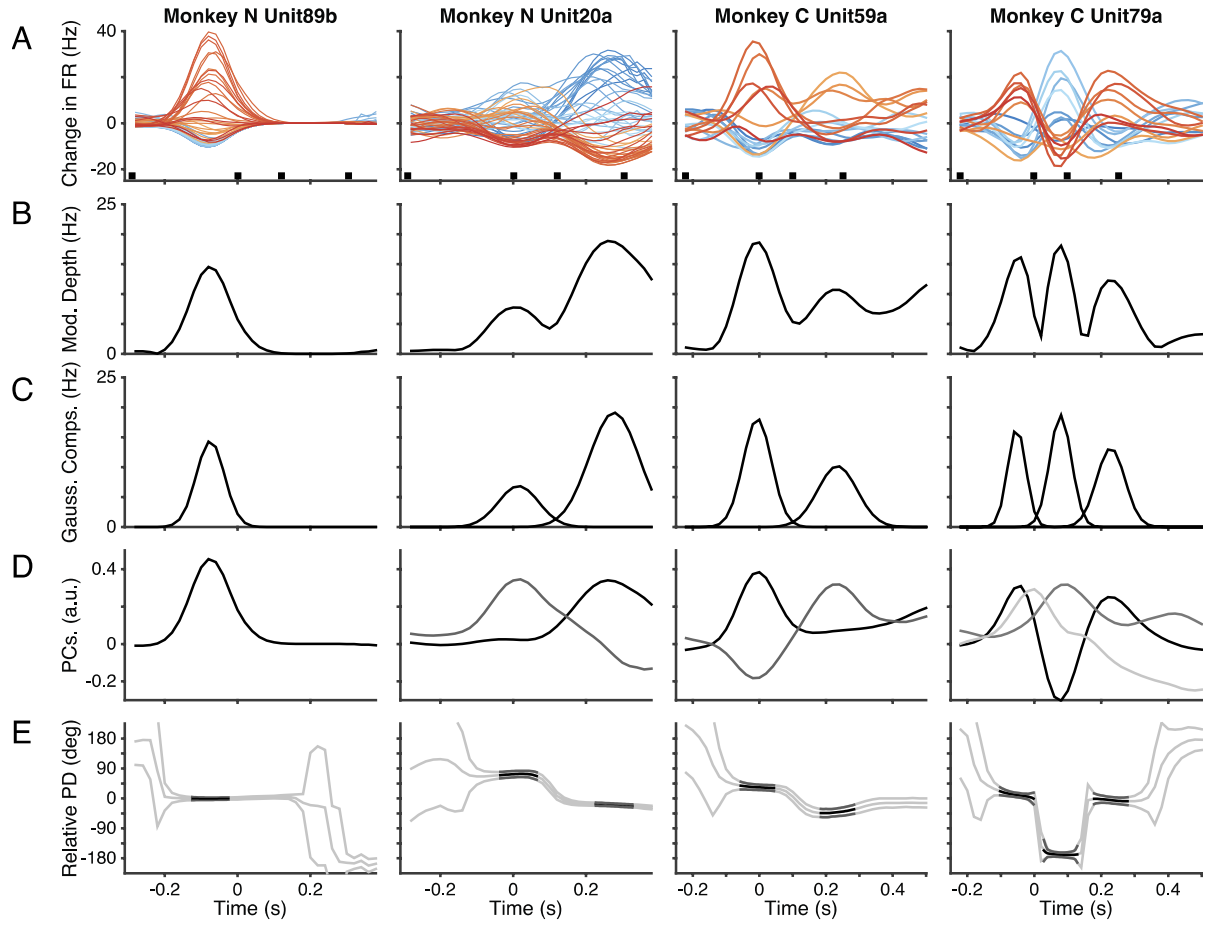

Figure 2.4 Episodic modulation patterns

Firing rate patterns of four different neuronal units are shown in columns. **A**. The first row shows trial-averaged firing rate profiles for each target direction. Square markers correspond to alignments on target show, movement onset, peak velocity, and movement offset. **B**. Tuning function amplitudes (estimated modulation depth) calculated in each bin are shown in the second row. **C**. Gaussian-shaped components fit to the tuning amplitudes are shown in the third row. **D**. Eigenvectors calculated from correlation matrices across bins are shown in the fourth row (darker line intensity for eigenvectors with more explanatory power). **E**. Preferred directions computed repeatedly over the course of the trial are shown in the fifth row. The darkened portions highlight the 100ms period centered on each Gaussian-shaped component and which was tested for PD stability. Error bars show the 95% CI (computed using bootstrap).

Dividing each reach into segments determined by the epoch-specific tuning fits showed that within those segments, preferred direction was typically constant for at least 100 ms (5 time bins) as defined by our preferred direction stability test (Figure 2.5A). Within each segment, firing rates were well described by cosine tuning (Supplemental Figure 2.12A). Further, the trend towards stability was evident even when considering only the components with very good fits to a cosine function (proportion of yellow to blue in rightmost bars in Supplemental Figure 2.12A). A majority of the neurons had a single stable component and units with three stable components were rare (Figure 2.5B). When *unstable* components were analyzed, the amount of PD change was found to be quite modest despite being statistically significant (Supplemental Figure 2.12B). Possible reasons for this finding (e.g. a conservative stability criterion) and further discussion are given in Supplemental Figure 2.13.

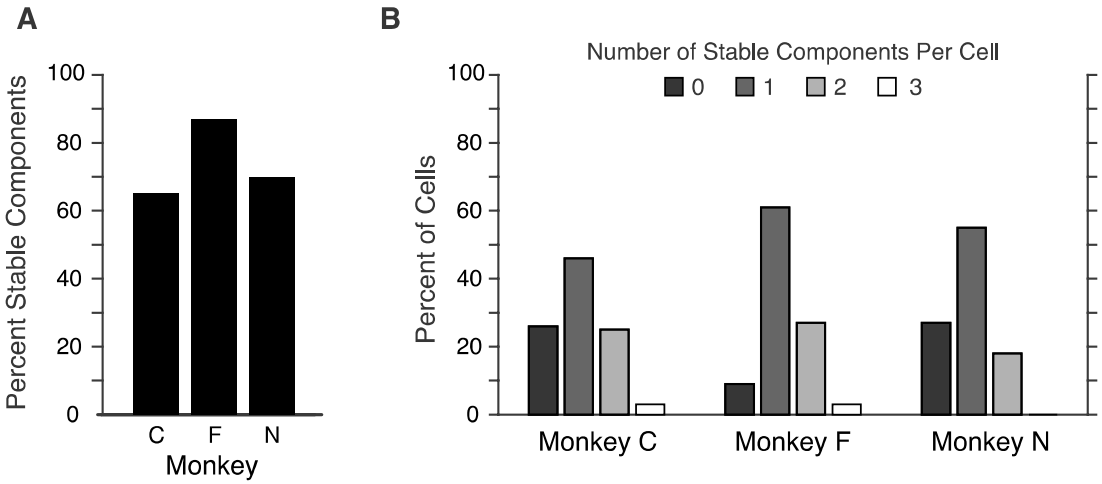

Figure 2.5 Preferred direction stability during component times

A. Component-wise directional stability. The majority of fitted components for each monkey were found to have stable PDs for the 100 ms period centered on the component peak. **B.** Number of stable components fit per neuron. The majority of components had stable preferred directions for at least 100 ms. When considering the number of components fit per neuron, regardless of stability, the breakdown was similar. From 0 to 3 components, the percentages for monkey C were 27, 33, 35, 4; for monkey F 6, 56, 34, 3; and for monkey N 39, 42, 18, 0. Total number of components fit for monkeys C, F, and N was 109, 161, and 146, respectively.

The PD differences between successive components in the example neuron displayed in the fourth column of Figure 2.4 were approximately 180º. Although the difference in PD between components was most often statistically significant (Supplemental Figure 2.12C, blue vs yellow), this extreme change in preferred direction was rare. For cells with multiple components, the angles between them (Supplemental Figure 2.12C) tended to be small, with 37% having values below 45º and only 11% in the 135-180º interval.

We next considered the effect of smoothing the recorded firing rates prior to using the preferred direction stability test. Since the test evaluates sequential time bins, smoothing the rates decreases the effective number of degrees of freedom. Given the same test statistic (Equation 2.5), the frequency of type II statistical error would then increase, inflating the reported number of "stable" components. However, smoothing also reduces the amount of noise in the firing rates, which in turn changes the value of the test statistic and has the opposite effect on the outcome of the statistical test. To directly assess these opposing effects of smoothing, we again used the PD stability test on the component data, but without applying any smoothing to the firing rates (aside from binning in 20 ms bins). We found that smoothing had very little effect on the results of the stability test. For components from monkey C, we found 65% to be stable after smoothing, and 69% stable without smoothing; from monkey F, we found 87% stable after smoothing and 80% stable without smoothing; from monkey N, we found 70% stable after smoothing and 71% stable without smoothing.

To assess the sensitivity of our stability test to different amounts of PD variation in the presence of realistic noise, we generated simulated firing rates with a preset amount of PD change (Supplemental Figure 2.14A). To approximate the characteristics of our experimental data, simulated rates were matched to each recorded neuron in terms of cosine-tuning modulation

depth and trial-to-trial firing rate variability. To directly compare our simulation to the experimental data, we restricted the test to 100ms (5 bin) windows centered on each of the Gaussian-shaped components that were fit to the real data using the epoch-specific tuning model. Therefore, the simulation results are best compared to those reported in Figure 2.5A. Modest changes in the simulated PDs were consistently identified as "unstable" by our test (Supplemental Figure 2.14B). Given that our simulated rates were carefully matched to the real data in noise and tuning depth, this finding suggests that change in the actual PDs during these components is minimal. Indeed, if the PDs within components do change, those changes are considerably smaller than those typically observed between components (Supplemental Figure 2.12C).

Although a cosine function described the firing rate data well at most time points, it is possible that our PD stability test was biased by imperfect fits. We therefore utilized another method for evaluating tuning stability based on the correlation between target-specific firing rates, which doesn't assume any parametric shape. We used this correlation metric to characterize our data in several ways, and found strong evidence for periods of stable directionality punctuated by rapid changes in tuning, without the assumption of cosine-tuning. These analyses and findings are described in detail in Supplemental Figures 2.15 – 2.18.

2.4.4 Segments of tuning occur at similar task times across the population of neurons

Next we asked whether the timing of components was consistent across the population of recorded cells. The times in the trial of each stable component (specified by the time of the Gaussian-shaped modulation peak) were used to build histograms separately for each monkey (Figure 2.6). Analysis of the histograms showed that the collective components from individual

cells tended to cluster into three epochs within the reach trial. The timing of these global epochs was consistent across monkeys, and tuning functions calculated throughout the trial tended to have maximum amplitude and consistent preferred directions within these three epochs. Interestingly, the clustering we observed for monkeys C and F was clearly more distinct than for monkey N, and for this monkey, the probability of observing components later in the task decreased. Monkeys C and F were required to stop and hold in the target before being rewarded, while monkey N was rewarded as soon as the cursor touched the target. The behavioral significance of this event therefore differed across monkeys. Because the third directional epoch is centered around the time the target was acquired, and was less prominent in monkey N, the directionality reflected at that time may be related to arresting the movement.

To test this hypothesis, we trained a fourth monkey ("P") to perform a center-out task in which the stop-and-hold requirement was cued by target color. The monkey performed this task for several months prior to collecting behavioral and neuronal data over three days. We then repeated our component-fitting procedure separately for "hold" and "no hold" reach trials, fitting a total of 265 components in the "hold" task and 184 components in the "no hold" task (231 stable in the "hold" task; 143 in the "no hold" task). We found that the number of neurons with components observed in the third epoch was markedly decreased in "no hold" trials (Figure 2.6, histograms for monkey P in panels A and B). Further, the separation in time of the three epochs was enhanced in the "hold" trials compared to "no hold" trials. This dissociation was most obvious when we combined data from the four monkeys and built histograms separately for the two task requirements (Figure 2.6 A,B bottom row). Requiring an accurate target acquisition prolonged the last phase of the reach, as we found that monkey P entered the target zone 43ms sooner in the "no hold" trials than in the "hold" trials, measured relative to movement onset (t-

test, $p = 3.3\text{x}10^{-79}$, $t = 19.4$, $df = 3391$, 95% CI = [38ms, 47ms]). These results support the idea that the third component is associated with the terminal portion of the reach.

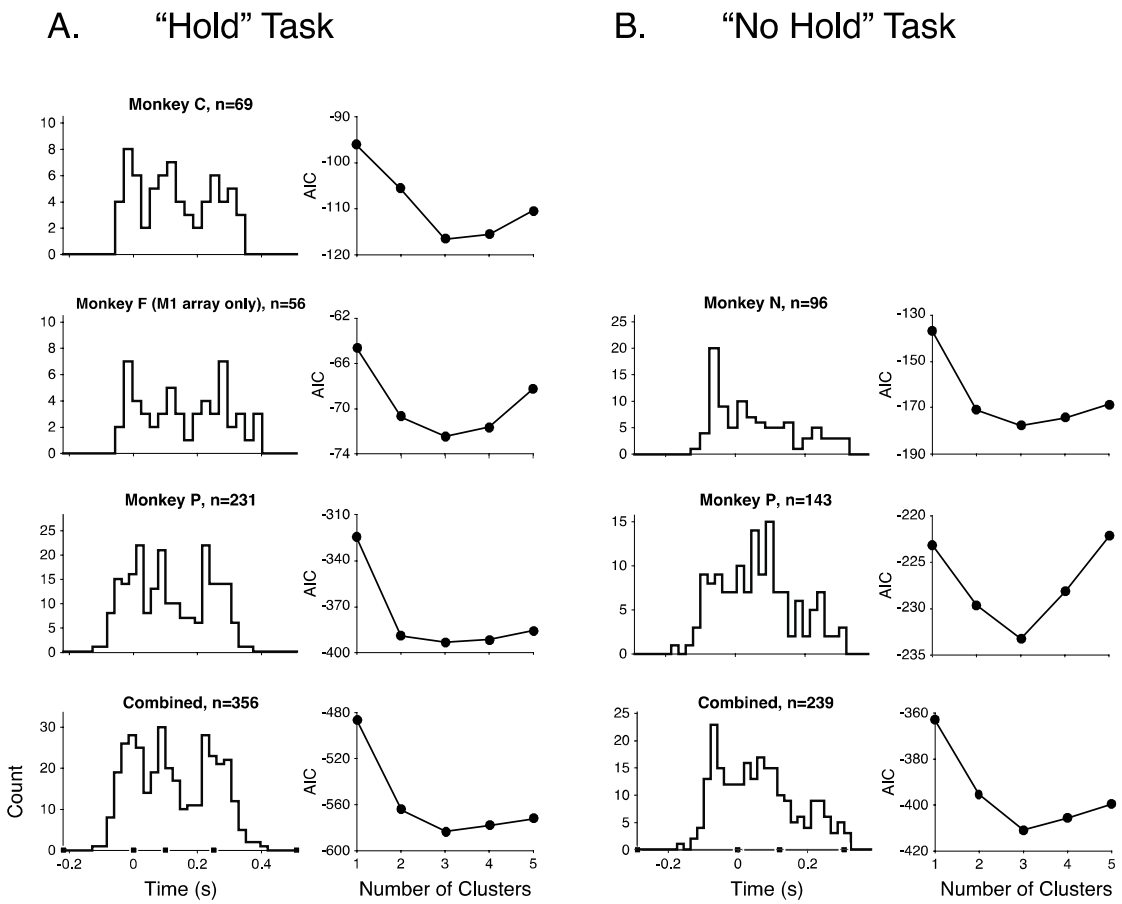

Figure 2.6 Temporal distribution of tuning components

The time of Gaussian-shaped peaks from "stable" components was used to make histograms of component occurrence for each monkey (left columns in **A** and **B**). Gaussian mixture models (GMMs) were fit to component times, and the Akaike information criterion (AIC) was computed for fits with 1 through 5 clusters (right columns in **A** and **B**). We consistently observed a minimum AIC for 3 clusters, suggesting components fall into one of three epochs, which are similar across monkeys. **A.** Stable component times observed during "hold" tasks. **B.** Stable component times observed during "no hold" tasks. Histograms in the bottom row of **A** and **B** represent data pooled from each preceding histogram. The three epochs are most distinctive in the "hold" task, but can also be observed in the "no hold" task. Movement begins at Time = 0 and the filled square symbols denote the behavioral events of target show, peak velocity, movement offset, and end of hold (if relevant). The time scales in the left column have been scaled linearly between event times to match those of monkey C; time scales in the right column have been scaled to match monkey N.

Regardless of task type, the earliest period of stable tuning began about 100 ms before movement onset, and peaked about 50 ms later. The timing of the second epoch peaked later in the reach, about 50 ms before the arm reached its peak velocity. The third epoch began around peak velocity and peaked approximately at the end of the movement. Most cells had a single modulation peak that contributed to one of the three epochs, although some cells had multiple directional components that fit into these epochs. Those cells with directional firing in an epoch tended to have stable preferred directions within that period. As shown by our tuning function analyses (Figure 2.3A,B), directional tuning is not evident early in the trial. The sharp increase in tuning function fit about 100 ms before movement onset corresponds to the beginning of the first global tuning component which peaks about 50 ms before movement onset (Figure 2.6). A second, separate directional feature peaks subsequently, overlapping the first, so that two strong and sequential directional signals in the neural population are evident during the reach. Single neurons driven by both signals could be found to have temporally-distinct tuning functions with different preferred directions (for example, the units in the middle two columns of Figure 2.4).

2.4.5 Accounting for tuning changes can be important for accurately decoding reach trajectory

The concept of directional tuning in the motor cortex has received a great deal of support since its introduction in the early 1980s (for a review- Georgopoulos and Carpenter, 2015). One of the first population vector demonstrations showed that the representation of direction evolved and pointed consistently to the target during center-out reaches (Georgopoulos et al., 1984). Following that work, neural trajectories were constructed by adding the time-series of population vectors tip-to-tail, and were found to closely match the trajectory of the arm when reaching to

targets (Georgopoulos, Kettner, and Schwartz, 1988) and when drawing a variety of figures (Schwartz, 1993; Schwartz, 1994; Moran and Schwartz, 1999; Schwartz and Moran, 1999).

How could such a large body of work, based on the idea of a fixed preferred direction, produce successful movement predictions if the tuning function is labile? Because the typical tuning model is a cosine, the firing rate average over multiple epochs of cosine tuning is also well-described by a cosine, even if PD changes occur between epochs (Supplemental Figure 2.19). As is the case with many center-out reports, using this average PD in the population vector algorithm yields accurate decoding, tending to minimize the errors due to PD changes between epochs. We found that directional tuning begins 100 ms before movement onset and extends beyond the end of movement. Using the amplitude of directional tuning, we identified three separate epochs. Within each epoch, preferred direction was constant. The first two epochs spanned the initial portion of the reach. A change in preferred direction within this span is rare; most cells had only one modulation component (occurring in the first or second global epoch) in this peri-movement period. Indeed, when we restricted our data sample to this portion of the task, trajectories decoded with population vector analysis were accurate (Supplemental Figure 2.20B). Because confining the data to this portion of the task is common practice, the issue of labile preferred direction has not been identified as a major problem in decoding experiments. Most of the directional lability found in our data samples came from the early reaction time portion of the task when directional drive was weak, or later in the task as the target was acquired and held, which may be a phase distinct from the initial portion of the movement. It should be noted that our analyses focused on directional modulation and this may differ substantially from overall firing rate modulation. Casual examination of our data suggests that that firing rate modulation tends to decrease through the reach, while directional modulation (e.g. in the third component)

may still be significant. Since the population vector algorithm is based on firing rate modulation, distortions of the neural trajectory toward the end of the reach due to changes in preferred direction may be marginalized by lower firing rates.

As a linear decoder, the population vector is optimal when neurons have cosine tuning functions and the recorded sample has a uniform distribution of preferred directions (Salinas and Abbott, 1994; Kass et al., 2005). Deviations from cosine-tuning will lead to inaccuracies in the decoded movement. Furthermore, when the preferred direction of a neuron changes during a movement, trajectories decoded from a time-series of population vectors can be distorted if these non-stationarities are not taken into account. Population vector trajectories were constructed for the reaches performed by each monkey (see Section 2.3.5). Results from monkey C are shown in Figure 2.7. When neural trajectories were constructed with a single preferred direction taken as the initial occurrence of significant tuning, the neural trajectories were distorted and shortened in the 10 and 4 o'clock directions (Figure 2.7A). However, use of component-wise preferred directions in the corresponding portion of each trial led to more accurate population vectors (Figure 2.7B, Supplemental Figure 2.20A). It could be argued that if a cell's directionality changed constantly during a reach, the most accurate reconstruction could be achieved using preferred directions with a high temporal resolution. To test this idea, we calculated preferred directions in each bin for every cell in the sample, and computed the trajectories in Figure 2.7C. These reconstructed trajectories were only slightly better than those generated with the component-wise PDs (Figure 2.7B), suggesting that the components adequately captured the changes in directionality that lead to decoding errors. It should be noted however, that these considerations apply mostly to the last portion of the trajectory. If the analysis is confined to the portion of each trial defined by the first two global components (Figure 2.6) around the time of

movement, using a single preferred direction for that epoch was almost as effective as using separate components (Supplemental Figure 2.20B).

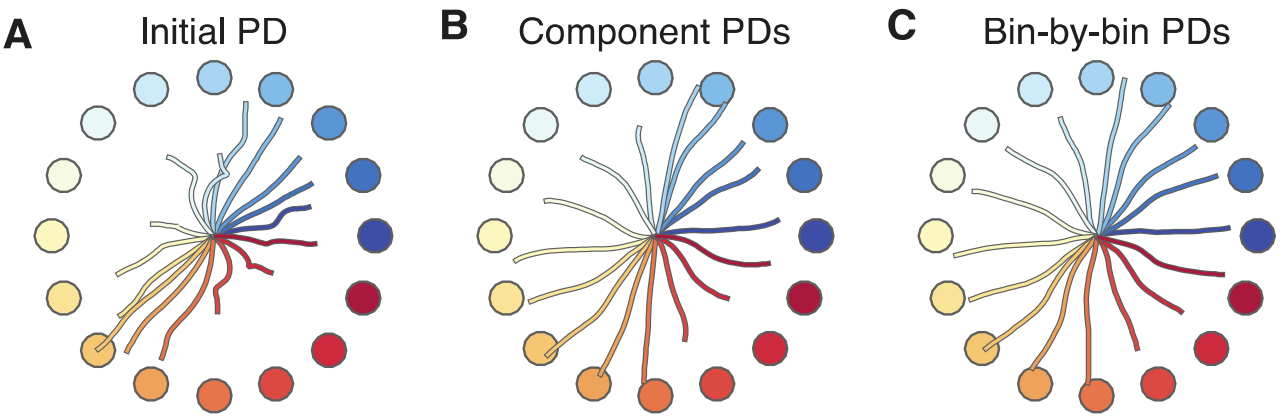

Figure 2.7 Population vector reconstructions of cursor trajectories to 16 center-out targets for monkey C

Neural trajectories were generated by calculating population vectors in each bin (from 100 ms prior to movement until 100 ms prior to end of trial) and adding the vectors tip-to-tail. **A.** The population vectors were calculated using a single, initial preferred direction for each neuron. This method resulted in highly distorted trajectories. **B.** Preferred directions from within each component were used to construct the population vectors. This method largely remedied the distortions seen in **A**. **C.** Tuning functions were calculated in each time bin and their preferred directions were used for the neural trajectories. This method yielded only modest improvements over the method used in **B**.

2.4.6 Offline identification of system state enhances accurate decoding

Assuming that the directional components effect changes in directionality simultaneously for many neurons, it should be possible to identify, at any point in time, which component is acting as the current driver on a neuron. The identity of the component could then be used as an index to the unit's component-wise preferred direction. Using each cell's segment-specified preferred direction would be expected to generate a more accurate readout of the arm's direction than a prediction produced with a single preferred direction. This real-time readout would be useful, for instance, in neural prosthetics. To demonstrate the utility of this approach, the "current" tuning segment of each neuron was detected on each trial and in each time bin by using the concurrent activity of all simultaneously recorded units in a linear-discriminant analysis. This showed that the individual components of neurons could be identified with high accuracy for the three monkeys (Figure 2.8).

To illustrate the feasibility of real-time decoding of movement trajectory using the components model, we performed an additional PVA analysis (based on Equation 2.8). In this version, the PD at every time step in each individual trial was defined by the PD of the component detected by the LDA for that sample. The identified components, when used to specify each neuron's preferred direction (per trial and time bin), produced population vectors that were comparable to those in Figure 2.7B. We used a bootstrap procedure to compare the magnitude and angular accuracy of population vectors computed this way with those computed using the true labels. The results of this method did not significantly differ from those using each cell's "ground truth" components (bootstrap test, $p > 0.05$ for each monkey).

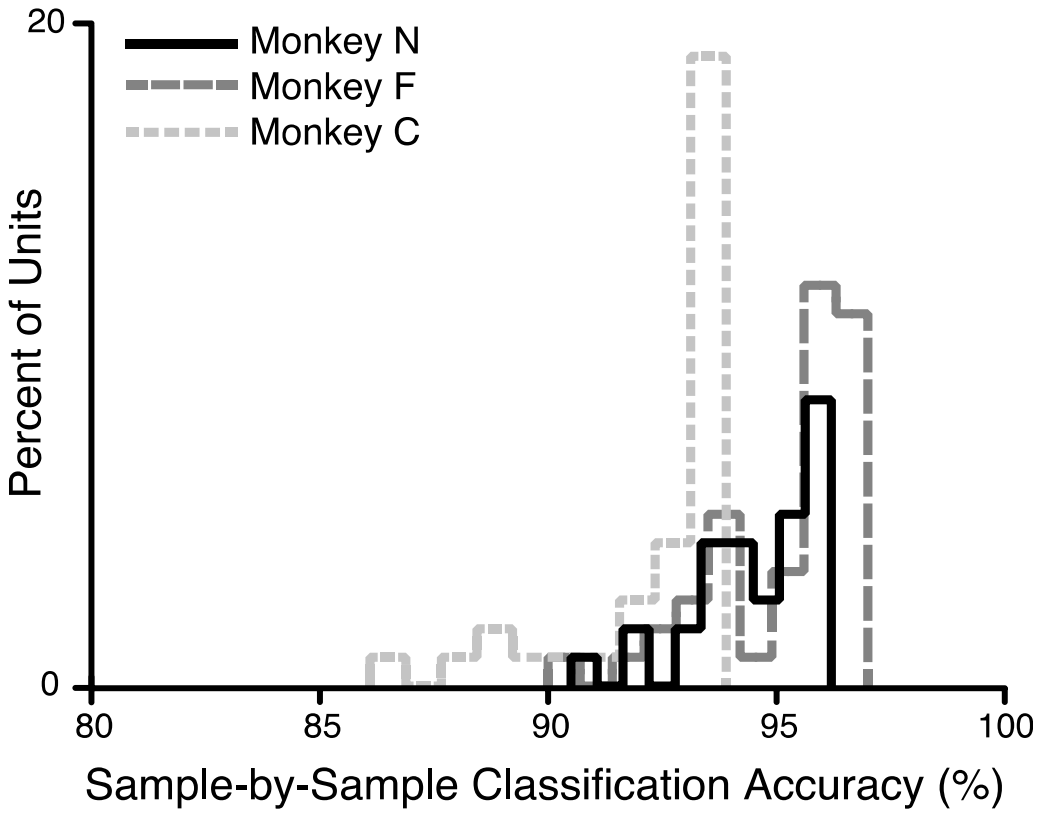

Figure 2.8 Offline component recognition

We investigated the possibility that patterns of population activity could be used to identify the tuning segment (component) of each single unit at any given moment in a trial. We trained classifiers (Linear Discriminant Analysis, LDA**),** one for each unit, using the concurrent firing rate samples from every simultaneously recorded unit. For each unit's classifier, samples in each bin were labeled based on their timing, relative to the unit's previously fitted components, with transitions delimited by the point at which one component's amplitude became larger than the previous component's. These labels were used to train the linear discriminant classifier for that unit, which was then used to predict the "current" temporal epoch for each 20 ms sample during each trial. Five-fold cross-validated success rates for units with at least two components were high: 94.7 ± 1.4%, 95.2 ± 1.6%, and 92.5 ± 2.0% for monkeys N, F, and C, respectively. Histograms of success rate per unit are displayed for each monkey. Success rate by chance was computed by randomly shuffling the group labels prior to training the classifiers, and yielded much worse classification: 50.1% ± 0.3, 48.8 ± 5.1%, and 48.2 ± 5.4% for monkeys N, F, and C.

2.5 DISCUSSION

Although the center-out task invariably shows that the mean discharge rate of neurons is tuned to movement direction, variations of this task, in which consideration of speed/amplitude (Churchland and Shenoy, 2007), shifts in workspace location (Caminiti, Johnson, and Urbano, 1990), movement fragments (Hatsopoulos, Xu, and Amit, 2007), isometric hand forces (Sergio and Kalaska, 1998, 2003), arm posture (Scott and Kalaska, 1995), limb biomechanics (Suminski et al., 2015), and covert encoding perturbations during brain-controlled tasks (Jarosiewicz et al., 2008) have shown that directional tuning functions are labile. Since some of this directional instability has been found within single reaches, concerns have been raised that the apparent representation of movement direction may be secondary to other aspects of motor control such as force generation (Sergio, Hamel-Pâquet, and Kalaska, 2005), oscillatory dynamics (Shenoy, Sahani, and Churchland, 2013), or intrinsic skeletal-muscular action (Suminski et al., 2015). One of these papers (Suminski et al., 2015), using a center-out task, reported rapid changes in motor cortical preferred directions around movement onset, a finding that is congruent with the transition between components one and two found in our work. Instead of attributing this transition to behavioral aspects of the task, these investigators suggested that the preferred direction changes were related to successive recruitment of different muscles during the task, with the argument that the primary determinant of a motor cortical neuron's preferred direction is the set of muscles to which it is connected. The primacy of motor cortical activity as a driver of muscle contraction has been controversial since the original characterization of this cortical structure (Fritsch and Hitzig, 1870). While motor cortical output undoubtedly has a role in altering muscle excitability, it is equally clear that there is no simple correspondence between a neuron's firing rate and the excitability of its "muscle field" (i.e., "upper motoneurons"), even

for the small portion of the corticofugal output that projects monosynaptically to spinal motoneurons (Griffin, Hoffman, and Strick, 2015). The historic issue of whether the motor cortex operates as a direct link to muscles or has a large role in movement planning is rooted in the assumption that a single function needs to be assigned to this anatomical structure.

Current theories of motor control link displacement, the generation of muscle activity, and behavioral output into a cohesive scheme (for a review see Diedrichsen, Shadmehr, and Ivry, 2010). In this approach, the intended state of the arm (e.g., position and velocity) is transformed by an "internal model" to a set of muscle activations, which, in turn, displace the limb as an action on the external world. Sensation from the movement is then fed back to the system. The modeled transformations are serial with the expectation that they correspond to the purported hierarchical anatomical connectivity between brain structures. Advocates of direct cortical participation in muscle contraction equate the motor cortex with the internal model, as the place where kinematics are transformed to muscle activation. Despite the attraction of these control schemes, neurophysiological results do not support the correspondence between separate brain structures and the discrete operational components used to construct these models. Multiple brain structures at different levels of the proposed hierarchy have similar directional characteristics (Suminski et al., 2015; Kalaska, Caminiti, and Georgopoulos, 1983; Johnson, Mason, and Ebner, 2001; Turner and Anderson, 1997; Kutz et al., 1997). Describing specific aspects of function is useful for framing concrete motor control issues, but localizing function to specific portions of the neural substrate may have less utility. The assignment of a single specific role to the motor cortex may be incompatible with the functional dynamics of the motor system.

There has been renewed interest in treating point-to-point reaches in the context of a dynamical system. The original idea was that an equilibrium position, defined by the joint angle

where the forces of opposing muscles balance, could be preset before the movement began (Feldman, 1966; Polit and Bizzi, 1978). When the arm was released, achieving this equilibrium would require no further input. A definitive experiment (Bizzi, Acconero, Chapple, and Hogan, 1984) disproved this idea, although modifications based on the idea of stable equilibria (Flash, 1987; Hogan, 1984) have since gained traction. More recently the 'set and forget' control concept of reaching has been revived (Churchland, et al. 2012). This recent rendition posits that trajectory specification takes place during the pre-movement period, and upon movement initiation the inherent properties of the system move the arm to a target. This concept can be described by a dynamical system formalism:

$\dot{\lambda} = C\lambda + Du;$

where λ is firing rate (the system's state variable) and u is the trajectory specification (the system's input). The term C transforms the current firing rate to a *change* in firing rate; $C\lambda$ describes the component of firing rate change that depends on the current firing rate itself. Similarly, D transforms the intended trajectory u to a change in firing rate. The integral of the first term, $\int C\lambda \ dt$, is referred to as the "natural response" and that of the second term, $\int Du \ dt$, as the "forced response". In the recent "set and forget" concept, $u = 0$ during the movement; desired velocity acts as an input before the movement begins with subsequent movement executed by the "engine of movement" in the absence of input. Indeed, we see non-step like modulation of firing rate during the reaching task (e.g. Figure 2.1A). Since this is not accounted for by different combinations of movement parameters, our results suggest that there is a system-dependent structure to motor cortical activity. While this suggests that a putatively constant direction signal can elicit a time-varying change in firing rate, our results also show that encoding of directional input is consistently maintained during discrete epochs of stable

directional tuning throughout the movement. While this study cannot rule out the possibility that these directional signals are components of the natural response, our results are consistent with control linked to behavioral drivers associated with separate states occurring during the task.

Neural operations are likely to be determined by the different control criteria needed during behavior. For instance, reaching is characterized by a set of invariants or "neural constraints" (Bernstein, 1984). A behavioral perspective can help link our findings of directional states in neural populations to global control issues. For goal-directed reaching, an initial phase, characterized by a stereotypical bell-shaped velocity profile (Morasso, 1981), during which the hand moves almost to the target, is too rapid for visual feedback to act concurrently (Hollerbach, 1982). Subsequently, a distinct terminal phase takes place, during which the hand and target are foveated together (Soechting and Lacquaniti, 1981; Meyer et al., 1988; Milner, 1992). Different cortical mechanisms may operate in these phases (Paillard, 1982; Brinkman, 1972). The first directional epoch we identified corresponds to the reaction time, the second peaks in the initial transport phase, and the last component is linked to target acquisition. Recent studies based on correlational structure between units within a recorded population suggest that this structure changes at the onset of an arm movement (Sussillo et al., 2015; Kaufman et al., 2014). Another experiment using a reach-to-grasp task has shown that there is a clear transition in motor cortical activity correlation from object location to object type at the end of the reach (Rouse and Schieber, 2016). Analysis of firing rates using Hidden Markov Models during a reaching task supports the idea of stable regimes and sharp state transitions taking place in the control system (Abeles et al., 1995). Together, these findings support the concept of discrete processes taking place sequentially through a task. By recognizing consistent non-stationarities in the time-

varying patterns of neural activity associated with motor control, we can better describe details of the processing used to generate behavior.

2.6 ACKNOWLEDGEMENTS

The authors wish to thank Valerie Ventura for her analytical advice, Jeong-Woo Sohn, Andrew Whitford, and Robert Rasmussen for collecting data used in this paper, and Scott Kennedy for helpful discussions. The authors declare no competing financial interests. This work was supported by grants from the National Institutes of Health (R01-NS047356, R01-NS050256, RC1-NS070311); and Defense Advanced Research Projects Agency (W911NF-06-1-0053, N66001-10-C-4056, N66001-12-C-4027). Author contributions: S.B.S. and J.O. made equal contributions in analyzing the data and writing drafts of the paper. A.J.C.M. collected data, performed initial data analyses, and wrote initial drafts of the manuscript. G.W.F. collected data and contributed to the analyses. Z.L. contributed to the analyses. M.V., S.M.C., and R.K. gave analytical advice. A.B.S. conceived the study, participated in the data analysis, and drafted the paper.

2.7 SUPPLEMENTAL MATERIAL

2.7.1 Supplemental Figures

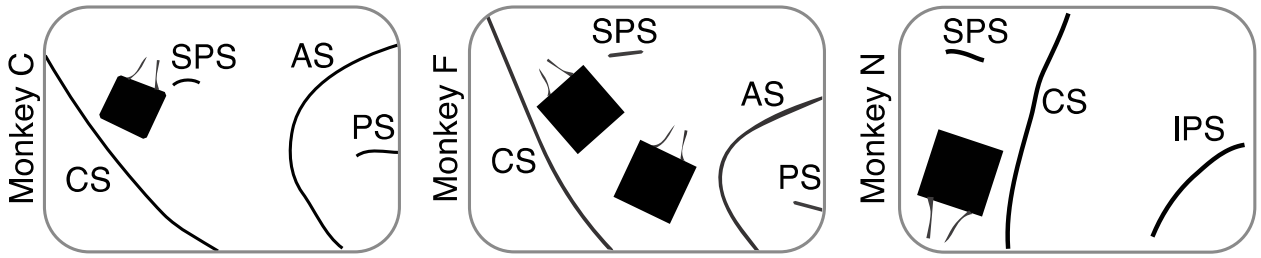

Figure 2.9 Electrode array placement

CS: central sulcus; SPS: superior pre-central sulcus; AS: arcuate sulcus; PS: principal sulcus; IPS: intraparietal sulcus. Black boxes represent Utah arrays (4 x 4 mm, drawn to scale). Small black lines connected to arrays represent wire bundles. Array placement for monkey N is approximate.

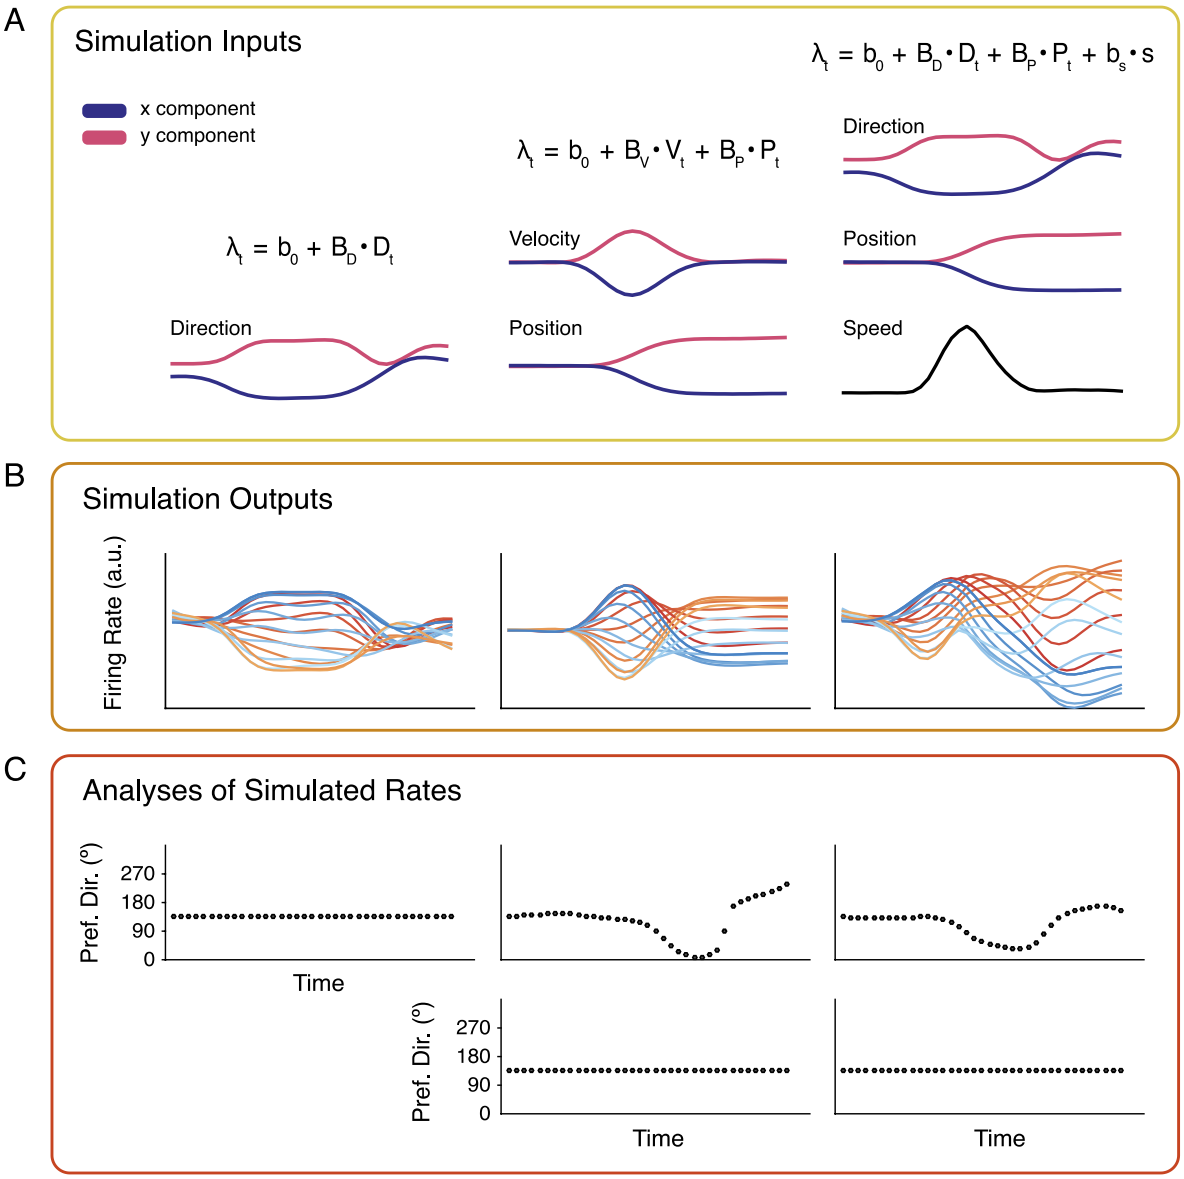

Figure 2.10 Simulation demonstrating static encoding of kinematic parameters and apparent tuning lability

For many neurons, there is a large amount of temporal variation that cannot be fit with a static directional model. Here, we used a simulation to test whether additional terms in the fitting model could account for the temporal variation in the data. The three columns in each panel describe the simulation and analysis of firing rates generated using different combinations of kinematic data collected from monkey C. From left to right, the simulation models are 1) direction-only; 2) velocity and position; 3) direction, position, and speed. **A.** Simulation models and examples of their inputs. λ_t is the neuron's firing rate at time t, b_0= baseline, D = direction, V = velocity, P = position, s = speed. The remaining terms are model coefficients that were chosen randomly and were kept constant at each time step. Lowercase letters indicate scalar quantities, and capitals indicate vector quantities. Examples of the trial-averaged x- and y-components of each kinematic term are shown for a single reach target. **B**. Outputs of simulations in A. Left, simulated rates using a direction-only generative model. The cosine tuning is expressed as step-function,

much like the idealized diagram in Figure 2.1A. Middle and right, encoding of multiple kinematic parameters produces more complicated firing rate profiles. **C**. Analysis of simulation outputs in B. The simulated firing rates were analyzed at each time step in two ways: 1) using the classical direction-only regression model (top row), or 2) using the regression model that matched the generative model (bottom row). Top left, the PD estimates of the simulated direction-only neuron are constant over time, as expected. Top middle and right, the PD estimates of these simulated neurons are temporally complex. This is because the generative model contained additional kinematic terms which were not included in the regression model. Bottom middle and right, if the regression model matches the generative model, an accurate estimate of the PD (constant over time as simulated) is recovered, despite the complex firing rate profile (panel **B**, middle and right).

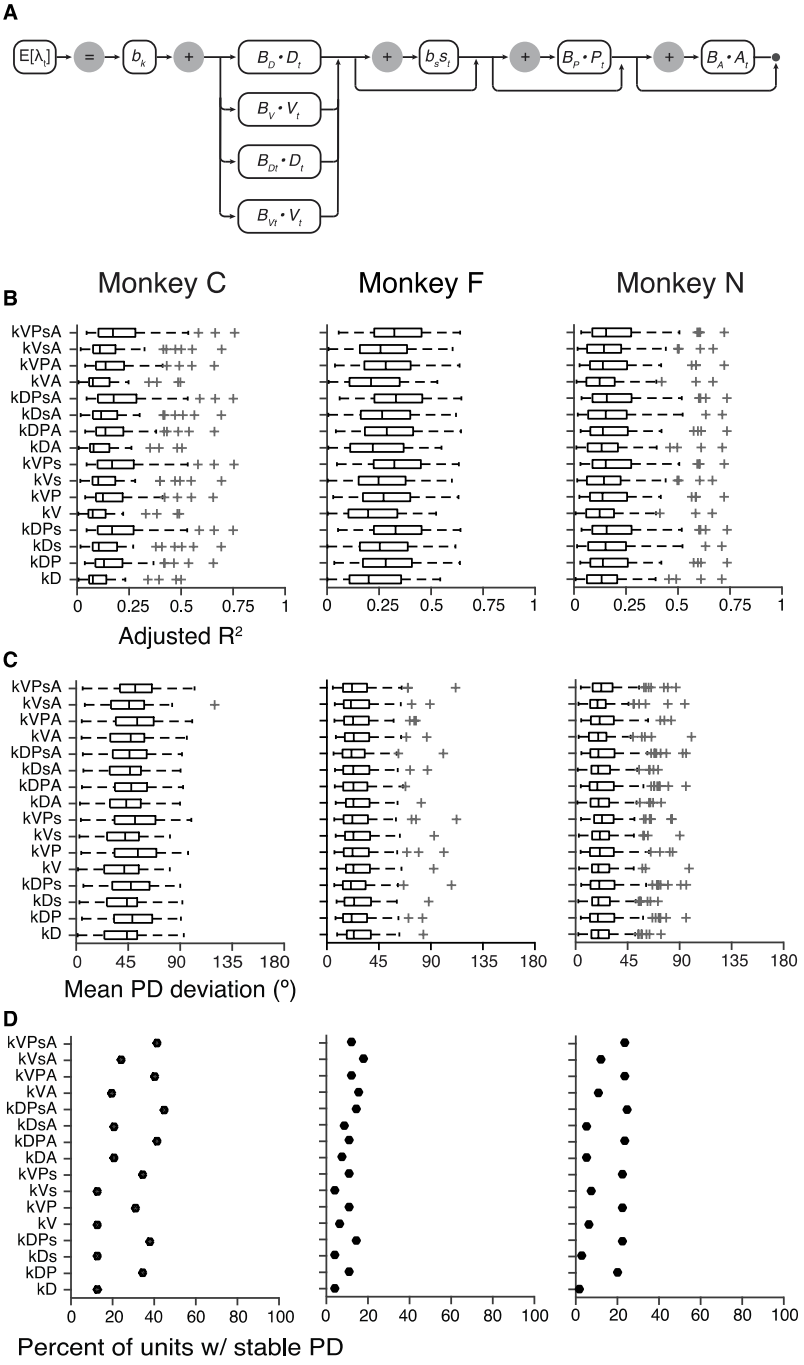

Figure 2.11 Least squares regression models of kinematic encoding

As shown in Supplemental Figure 2.10, tuning lability in simulated neurons can be observed if their kinematic drivers are not accounted for in the regression analysis. To assess the extent to which the inclusion of additional kinematic terms in the decoding model could account for tuning lability in experimental data, units' firing profiles were fit using general linear models including many combinations of kinematic parameters. Using the experimental data, we constructed 16 static models from combinations of five kinematic parameters of the hand end-point

(direction or velocity, position, speed, and acceleration) plus a baseline firing rate. Data were not trial-averaged prior to model fitting. **A**. Code syntax diagram describing model composition. Models are constructed by following the arrows; vertically stacked elements (the directional terms) represent mutually exclusive options. λ_t is the neuron's theoretical firing rate at time t, lowercase letters indicate scalar quantities, and capitals indicate vector quantities. b_k = baseline, D = direction, V = velocity, P = position, s = speed, A = acceleration. The remaining terms are coefficients fit using least squares regression analysis. Analyses included units with an average cosine tuning function $R^2 > .6$, and spanned the time period from 50 ms after target show until 100 ms after movement offset. **B**. Goodness-of-fit of models using a variety of kinematic parameters and static preferred directions (B_D and B_V in panel **A**). Box plots represent the distributions over units of adjusted R^2 values for each model. Including additional kinematic terms in the decoding model did not greatly improve model fits. One exception is the position term (P), which appears to slightly improve the fits for some neurons. This improvement can be attributed in part to the tonic discharge of neurons during the target hold period of the task; this period of tonic discharge is known to correlate strongly with position (Kettner, Schwartz, & Georgopoulos, 1988). Additionally, because position and direction are correlated in the center-out task, a change in PD for neurons in our dataset can be coarsely modelled using the kDP or kVP models, even if the directional terms are held static (see, for example, the middle panel in Supplemental Figure 2.10B). In general, however, the gains in model fit are modest. **C**. PD lability estimated using each model. To directly assess the impact of model parameters on PD excursion, we fit versions of each model in which the directional terms were allowed to vary in time (B_{Dt} and B_{Vt} in **A**). Direction and velocity were mutually exclusive in our models, and each model included one or the other. Each regression thus yielded a time series of direction or velocity coefficients that could then be compared to those from the corresponding time-invariant model. The average of these PD deviations was computed for each unit. In contrast to the direction or velocity terms, non-directional parameters (baseline, position, speed, or acceleration) were fit with single coefficients. That is, the relationship between firing rate and non-directional parameters was always assumed to be static in time. If the additional kinematic parameters explained the observed tuning lability, then the PD at each time point would remain relatively constant. Furthermore, there should be little difference between the PDs at each time point and the single PD estimated using the corresponding time-invariant model. To summarize this effect, we computed the mean of these PD deviations for each neuron and pair of models. As shown in **C**, this measure of PD deviation was largely unaffected by the additional kinematic parameters. The PD deviations of the simple direction-only model ("kD") and each of the more complex models were of similar magnitude across the population, suggesting that complex kinematic tuning cannot explain the observed PD lability. **D.** Tuning stability significance test carried out for each model. The tuning stability test described in Materials and Methods (main text) was expanded to include each kinematic parameter. This allowed us to test the hypothesis that a complex kinematic model with constant PD was as likely given the data as one with a time-varying PD. This hypothesis was tested for each combination of kinematic terms. Additional terms in the model in general did not have a large effect on the proportion of neurons judged to be stable by the hypothesis test. However, there is a clear influence of the position term on tuning stability estimates, particularly for monkey C. We attribute this effect to the same principles discussed for the similar effect of position on model fit described for panel **B**.

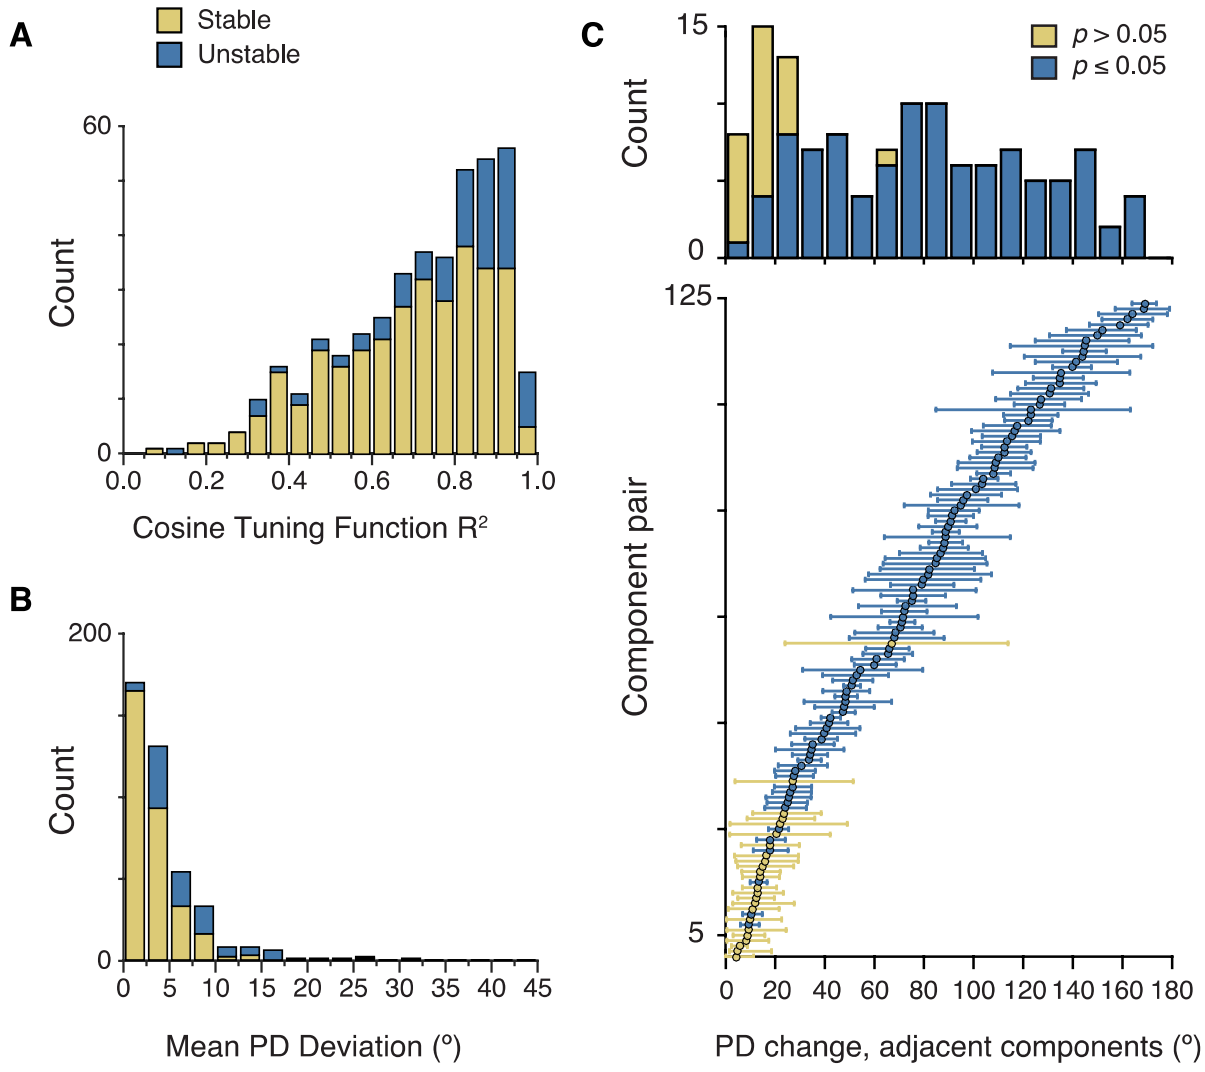

Figure 2.12 Analysis of firing rates during and between component times

To characterize firing rates during the epochs defined by our component fits, we applied regression analysis to snippets of firing rates spanning 100ms (5 bins) centered on each component peak. **A**. Quality of cosine fit to trial-averaged firing rates across targets. Rates from the snippets were averaged over task repetitions, and a single cosine function was fit using regression. Based on the large R^2 values, components' tuning functions are well-described by a cosine. **B**. A measure of within-component PD change. Two versions of the direction-only model were fit: one in which the PD was constant and one in which the PD could vary in time. We then measured the deviation between the PD fit using the constant model and each of the PDs fit for the 5 time bins. This deviation, averaged over the 5 bins, is displayed for each component. As shown by the consistently small deviation, PDs change very little for the duration of the component. Regardless of the outcome of the stability metric, mean PD deviation was below 10 degrees over 100 ms for 93% of the fitted components. In **A** and **B**, "stable" components are displayed in yellow; "unstable" components are displayed in blue. Data are pooled from all components fit to all neurons from monkeys

C, F, and N. **C.** PD changes between components. Although the *intra*-component change in PD was found to be very small, *inter*-component changes could be much larger. We thus characterized the distribution of these changes and tested each change for statistical significance. For each neuron with two or more adjacent components, we analyzed firing rates in 60 ms windows centered on each component peak. The trial- and time-averaged PD within each 60 ms window was then calculated. For each pair of sequential components, we estimated a 95% confidence interval for the difference in the two PDs. Confidence intervals were computed using a bootstrap procedure: trials were chosen randomly with replacement and firing rates from these trials were analyzed within the same 60 ms windows. These rates were used to re-estimate the average PDs for each component, and this procedure was repeated 1000 times to form PD distributions. Lastly, we computed the angular differences between the distributions of sequential components, and the 95% confidence intervals were determined by rank-ordering these angular differences and finding the 2.5^{th} and 97.5^{th} percentiles (horizontal bars). Data are shown for pairs of components adjacent in time, rank-ordered by PD change, and pooled over monkeys. The bootstrapped distributions were used to test the statistical significance of PD changes between adjacent components. The null hypothesis of a zero-degree PD difference between adjacent components was tested (yellow vs. blue in **C**). In 2D, it is possible to evaluate the signed difference between PD distributions and determine whether this difference overlaps zero. This is not possible for 3D PDs without imposing an arbitrary frame of reference. We therefore implemented a test that would generalize between both the 2D and 3D PD measurements. This was achieved by rotating the bootstrapped PD distributions of sequential component pairs to a common mean direction. Put differently, we shifted the two PD distributions to the same mean PD, while preserving the shape of each distribution. A null distribution was then created to describe the uncertainty in zero overlap due to the width of the two PD distributions. This was achieved by computing the angular differences between the rotated PDs, and a critical value for this null distribution was obtained by rank-ordering the differences and finding the 95^{th} percentile. The *p*-value of the hypothesis test was computed as the proportion of true angular deviations that exceeded this critical value. Results across monkeys were highly comparable.

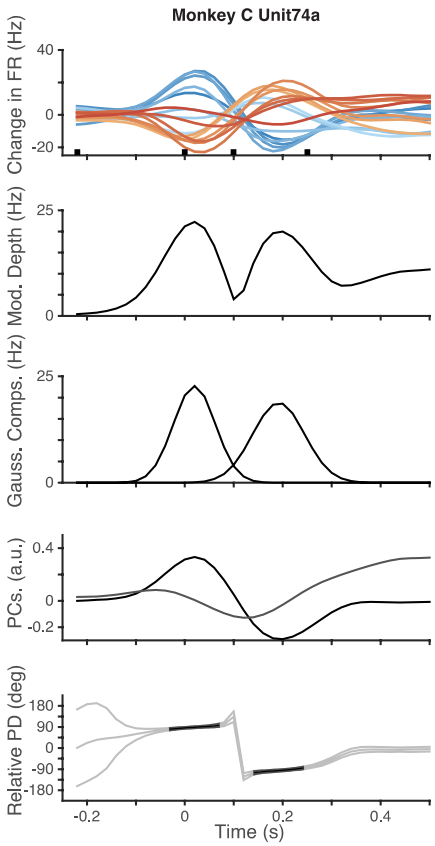

Figure 2.13 Example of firing rate "misalignment" and the effect on PD stability

Although components were most often "stable," the proportion of "unstable" PDs for components with the highest R^2 cosine fits was somewhat larger than expected (blue portions of bars in Supplemental Figure 2.12A). A possible explanation for this observation is imperfect alignment between firing rates recorded during reaches to different targets. A key assumption of the PD stability test (see Section 2.3.2) is that the time-normalized firing rates will follow a similar time course for reaches to different targets. Although time scaling based on behavioral events typically works well, it is not perfect. The neuron shown here, for example, exhibited small time-shifts across conditions that could not be remedied by alignment. For this neuron, we found patterns of multi-phasic activity that corresponded to the two peaks in the modulation depth function. The change in PD between the peaks was large and abrupt as predicted by our model. The R^2 values for a cosine fit during each component were very large, reaching above 0.96 at each peak. However, because the firing rate traces were not perfectly aligned (e.g. orange vs. red traces), the measured PDs within the components (bottom panel, darkened portion) showed a slight counterclockwise change over time. Further, because the signal-to-noise ratio for this unit is large (small confidence band around PDs, bootstrap method), the small PD change within each component was statistically significant. Both components for this neuron were found to be "unstable" by our test, even though the change in PD was very small (approximately 15 degrees over 100 ms, about 3º/bin). Details of the analyses in each panel in this figure follow the descriptions in Figure 2.4.

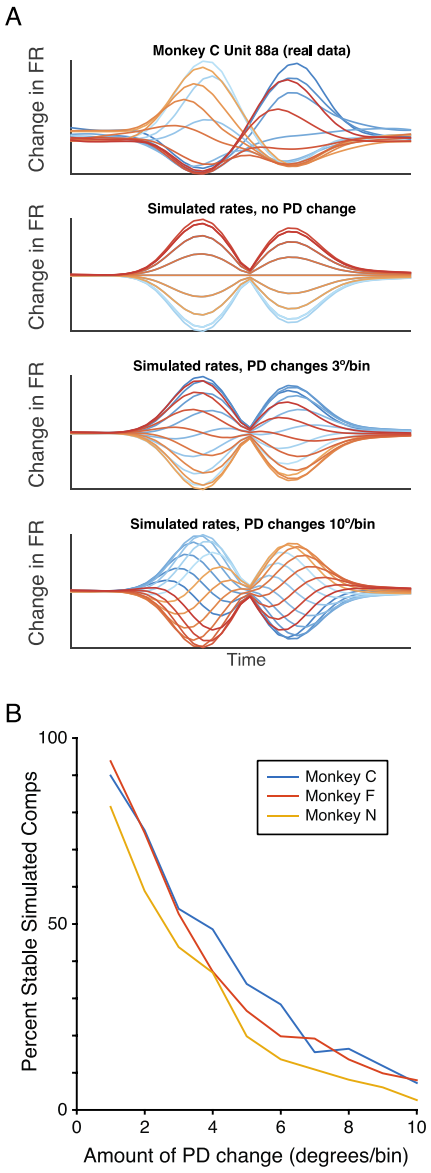

Figure 2.14 Sensitivity of stability test

To validate the tuning stability test, we simulated firing rates for each trial for every neuron from monkeys C, F, and N. Simulated rates were matched to the modulation depth and amount of noise in each time bin. For each simulated neuron, we then specified changes in preferred direction that occurred at different constant rates through the trial. The simulation was performed in two stages. 1) In the first stage, we applied the cosine-tuning regression model to each neuron at each time step, and used the resulting modulation depth function to reconstruct the expected rates at each time step. However, rather than using the estimated preferred direction at each time step, we specified a PD change over time. 2) In the second stage, we estimated the trial-to-trial variability of each neuron at each time bin by computing the standard deviation of the rates over task repetitions. We simulated this trial-to-trial variability by drawing samples, for each simulated trial at each time bin, from a normal distribution with standard deviation matched to the experimental data in that bin. The mean of the distribution on each trial and time step was

determined in the first simulation stage. The number of simulated trials was matched to the experimental data. **A.** top, real data from a neuron recorded from monkey C. The remaining sub-panels in **A** are simulated rates based on Unit 88a (prior to adding trial-to-trial noise), with modulation depth matched to the real neuron. The amount of PD change over time was different for each iteration of the simulation (shown for 0º change/bin, 3º/bin, and 10º/bin). **B**. Results of the tuning stability test applied to the simulated data, for each amount of simulated PD change. We applied the stability test to each simulated neuron, restricting the analysis to the times of the Gaussian-shaped components (fit to the real data). The test was applied to 100ms (5 bin) windows centered on each of the components.

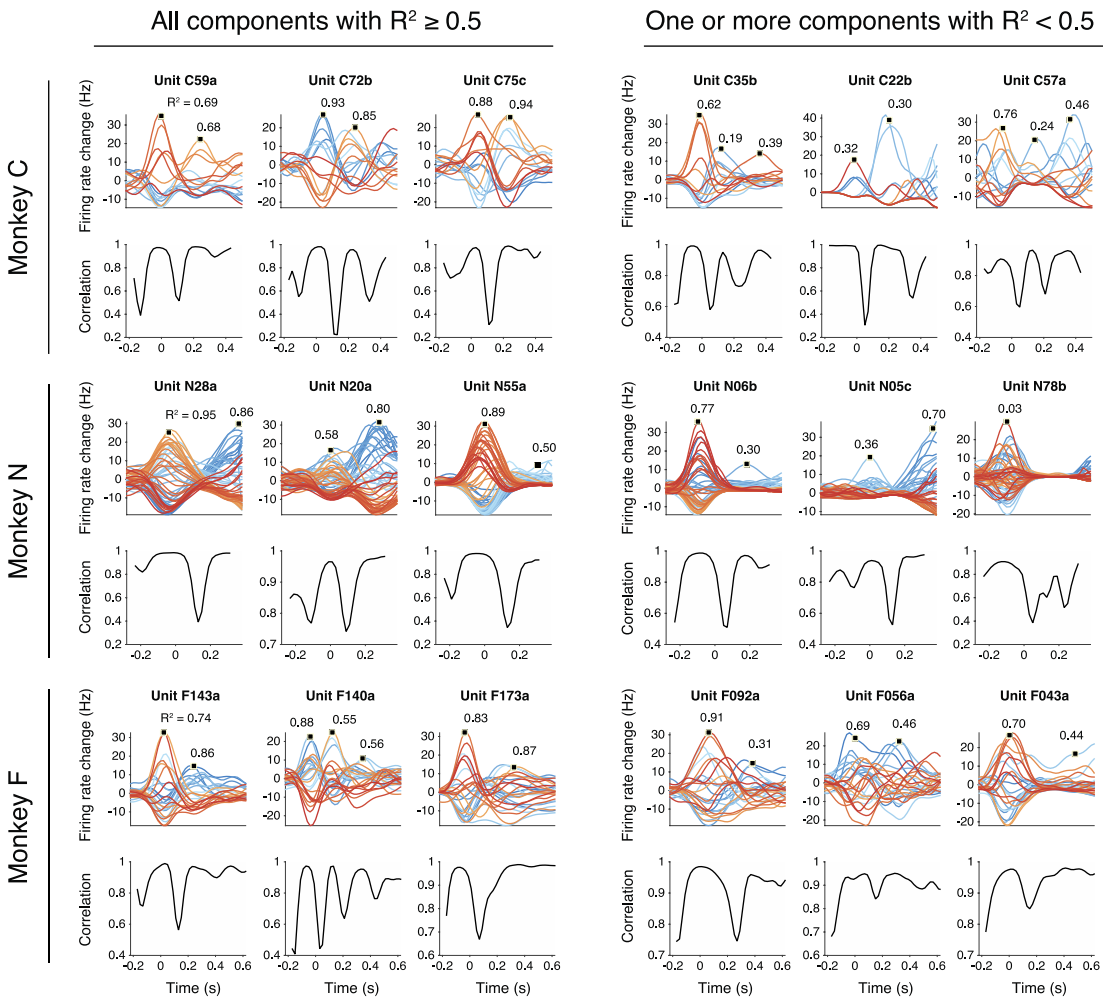

Figure 2.15 Time course of tuning correlation

An approach to evaluating tuning similarity without assuming any parametric form is to compute the correlation between pairs of tuning patterns. Each tuning pattern is considered a vector, the elements of which are the firing rates observed during each task condition. The correlation between two such vectors is a metric of tuning similarity; values near 1 indicate highly similar tuning, while values near 0 indicate dissimilar tuning. Here, we applied this metric for each neuron by correlating the tuning vectors of adjacent time points throughout the task to produce a time course of tuning similarity. This was accomplished by taking the rates in a 60 ms time window (averaged over trials and time) and correlating them with the mean rates in the subsequent 60 ms window, and repeating the procedure in steps of 20 ms. Each subpanel shows the firing rate of a single neuron (colored traces) above the corresponding tuning correlation curve (black traces). Comparing the time course of local tuning correlation with the time course of firing rate revealed plateaus of high correlation during periods of modulation, suggesting tuning was relatively stable at those times. The local correlations dropped sharply at the transitions between modulation components, indicating a period of rapid tuning change. Critically, this pattern was evident regardless of the goodness-of-fit of a cosine to the tuning pattern. Each firing rate plot has been marked in one or more places and labeled with the R^2 of a cosine fit to mean rates at that time point. Neurons grouped in the left column have high R^2

values (>= 0.5) at each modulated peak (component). Neurons grouped in the right column have at least one component with an R^2 less than 0.5. These examples demonstrate that periods of highly similar tuning (the correlation plateaus) are punctuated by rapid changes in tuning (sharp troughs), and that this phenomenon is not dependent on the quality of a cosine tuning fit. Note that each ordinate is plotted on a different scale. Since broad cosine-like tuning is a prominent feature of most components, the minimum correlation value at a transition is related to the magnitude of the PD change between components. For example, cosine tuning functions with a PD of 0º and a PD of 45º have a correlation of about 0.7.

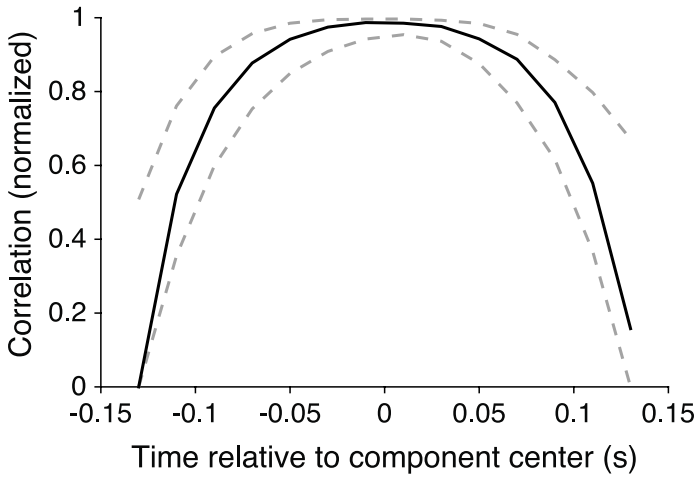

Figure 2.16 Component-centered tuning correlation

To summarize the pattern of peri-component tuning similarity, we correlated the across-target rates (averaged over trials) centered on each identified component. This was accomplished by correlating the mean rates in a 60 ms window at the center of the component with the mean rates in 60 ms windows before and after, again stepping in 20 ms increments. Time 0 represents the center of the component. Data are pooled from monkeys C, F, and N. Solid black trace, median across all components; dashed gray traces, first and third quartile. For each component-centered correlation profile, values were amplitude-normalized between 0 and 1 and time-normalized to the average component duration. Bounds of components were defined by the sharp troughs observed in Supplemental Figure 2.15 (black traces). A minimum trough value of 0.85 was chosen, although a wide range of cutoffs gave highly comparable results. If a trough was not found before or after a component center (e.g. Units C22b and F173a, respectively), a fixed period of 100 ms was taken in that direction. The exact duration was not critical, and excluding these data entirely did not qualitatively change the result. The pattern of tuning in the center bin was highly similar to the pattern observed for 50 to 100 ms before and after. At longer intervals, the correlation values tended to be much lower, dropping abruptly before the onset or after the offset of the component. For the majority of the identified components, the pattern of tuning observed at the peak remained consistent for approximately 100-200 ms. This is in good agreement with the duration of the plateaus observed in Figure S7, in which the tuning reference point was sliding rather than fixed to each component center.

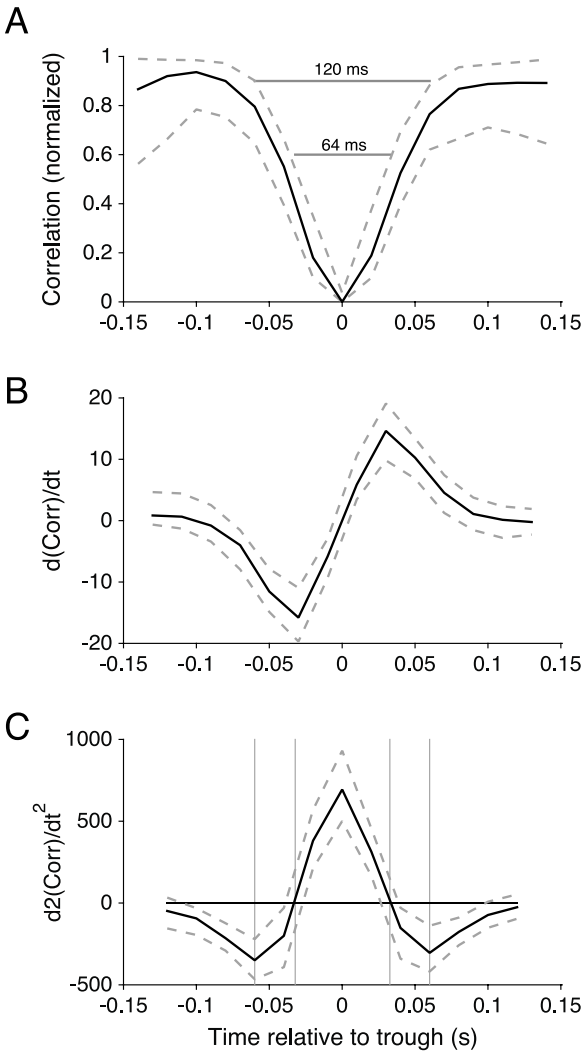

Figure 2.17 Analysis of abrupt changes in tuning correlation

To study tuning transitions in more detail, we isolated each tuning-correlation trough before and after a component (black traces in Supplemental Figure 2.15). The first and second time derivatives of each trough were then used to find the duration of the tuning transition. Troughs were included if their minimum value dropped below 0.85, although a wide range of cutoffs gave similar results. Data are pooled from monkeys C, F, and N. Solid black traces, median; dashed gray traces, first and third quartile for the population. **A.** Time-course of tuning transition. Each occurrence has been scaled between 0 and 1. Solid gray horizontal bars, median times of features described for subsequent subpanels. **B.** First time derivative of data shown in **A**. Two symmetric extrema are evident and mark the bounds of the most rapid period of tuning change. **C.** Second time derivative of data shown in **A**. The extrema seen in panel **B** were identified by the zero-crossings marked by the inner pair of gray bars (median distance, 64 ms). Before and after this rapid transition, the rate of correlation change slowed to zero. This second, slower feature was identified by the symmetric minima of the second derivative (outer pair of gray bars), which have a median distance of 120 ms. These analyses show that the transitions between components are relatively brief and of a consistent shape.

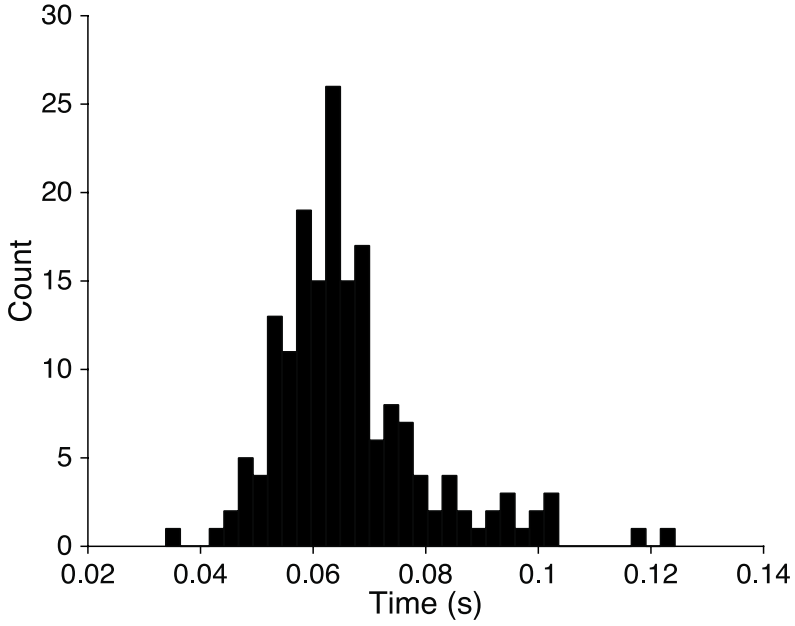

Figure 2.18 Variability in the duration of rapid phases of tuning change

To assess variability in the duration of tuning transitions, we measured the time between the two extrema of each trough's first derivative as shown in Figure S9. These times had a median value of 64 ms and were relatively consistent across occurrences. As examples, this timing for the prominent transitions of the three units in the first column of Supplemental Figure 2.15 was 58 ms, 62 ms, and 64 ms, respectively. Data in this figure are pooled from monkeys C, F, and N.

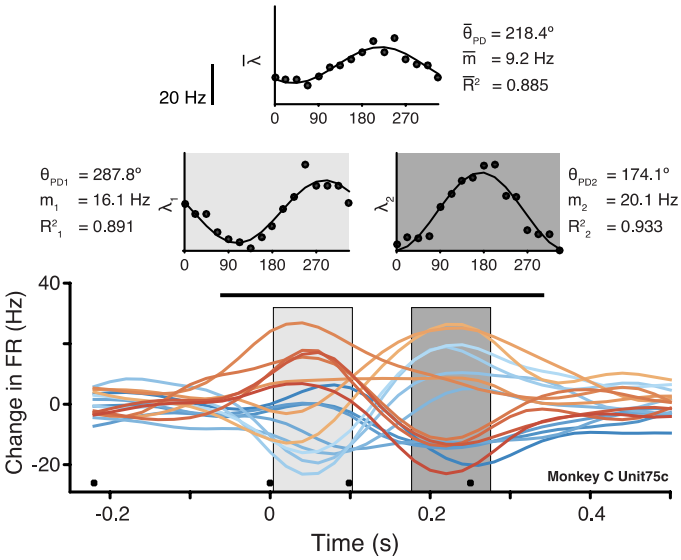

Figure 2.19 Averaging periods of cosine tuning also results in cosine tuning

Colored traces bottom panel: firing rates from a neuron that exhibits two epochs of modulation, marked by the two shaded gray boxes over the firing rates. Upper panels: Cosine tuning functions estimated from different time windows. The two tuning functions shaded in gray (middle row) correspond to the shaded gray boxes on the firing rate plot (bottom row). The tuning function in the top row corresponds to the period marked by the long black bar above the firing rates. Early work characterizing cosine tuning in motor cortex typically reported tuning curves calculated using firing rates averaged over relatively long periods of time. We found that fitting a cosine function to rates averaged in this way typically yields very high R^2 values, even for neurons with labile PDs. This is likely due to a property of cosines well known from the theory of Fourier series: summing two cosines of equal frequency but different phase (i.e. PD) results in a third cosine of the same frequency, with different amplitude and phase. Only in the event that the two tuning functions are 180º apart and of equal amplitude, complete cancellation of the tuning will be observed. Suppose the firing rate of a neuron at each time step can be described by a cosine tuning function:

$\lambda_t = m_t cos(\theta_{reach} - \theta_{PDt})$

The average of rates across time steps can be described with a cosine tuning function as well. If we represent this mean tuning function as a vector pointing towards the mean PD, $\bar{\theta}_{PD}$, the *x* and *y* Cartesian coordinates of the vector are related to the tuning functions at each time step as follows:

$$x = N_t^{-1} \sum_{t=1}^{N_t} m_t \cos\left(\theta_{PD_t}\right)$$

$$y = N_t^{-1} \sum_{t=1}^{N_t} m_t \sin\left(\theta_{PD_t}\right)$$

The resultant preferred direction and modulation depth are $\bar{\theta}_{PD} = tan^{-1}\left(\frac{y}{x}\right)$ and $\bar{m} = \sqrt{x^2 + y^2}$, respectively. As shown in this figure, when cosine-tuned epochs in real data are averaged together, their mean is well fit by a cosine.

(A)	Monkey C	Monkey F	Monkey N
Population vector angular error			
Early PD	20.7 ± 1.6°	23.7 ± 1.0°	18.7 ± 0.9°
Component PDs	11.7 ± 0.6°	19.5 ± 0.6°] n.s.	8.0 ± 0.3°
Bin-by-bin PDs	9.4 ± 0.4°	17.6 ± 0.5°	7.0 ± 0.2°
Population vector magnitude			
Early PD	5.9 ± 0.1	8.1 ± 0.1	8.7 ± 0.04
Component PDs	7.0 ± 0.1	9.6 ± 0.2] n.s.	10.1 ± 0.07
Bin-by-bin PDs	8.1 ± 0.1	10.2 ± 0.2	10.8 ± 0.08

(B)			
Population vector angular error			
Early PD	18.9 ± 1.4°] n.s.	15.3 ± 0.7°] n.s.	6.2 ± 0.4°] n.s.
Component PDs	17.5 ± 1.6°	15.3 ± 0.7°] n.s.	5.5 ± 0.4°] n.s.
Bin-by-bin PDs	11.7 ± 0.8°	14.0 ± 0.6°	4.7 ± 0.3 °
Population vector magnitude			
Early PD	6.8 ± 0.1	10.9 ± 0.2] n.s.	11.9 ± 0.1] n.s.
Component PDs	7.5 ± 0.1	11.1 ± 0.2] n.s.	11.8 ± 0.1
Bin-by-bin PDs	8.3 ± 0.1	11.6 ± 0.2	12.3 ± 0.1

Figure 2.20 Quality of PVA results using three different methods and two time periods

Values are mean over targets, trials, and time, plus/minus SD. Angular error is reported in degrees; vector magnitude is reported in arbitrary units. Methods were compared using a bootstrap test over task repetitions. As a correction for multiple comparisons, we chose an alpha level of 0.01; "n.s." indicates not significant at this alpha. **A**. Analysis spanning the period from 100ms prior to movement until 100ms prior to end of trial. Generally, accounting for tuning lability improved decoding. Using the component PDs rather than the early PD resulted in the largest improvement; smaller gains were observed by re-fitting the PDs at each time point. **B**. Analysis restricted to the period from -150ms to 100ms relative to movement onset. When just this initial portion of the population vector is considered, the difference between PVA methods is small.

2.7.2 Supplemental Experimental Procedures

2.7.2.1 Additional information about data sets

Mean reaction times for monkeys C, F, and N were 220, 236, and 290 ms, respectively. In the same order, mean movement times were 250, 318, and 304 ms. For the three monkeys, cursor radii were 10, 8-9, and 10 mm. For monkeys C and F, target diameters were 10 and 8-9 mm. Target sizes varied for monkey N. We sorted 93, 119, 185 single units for C, F, and N. Using our component fitting procedure (see Section 2.3.4), we fit 65, 98, and 119 units.

2.7.2.2 Control for multiunit activity

It is possible that multi-phasic firing could be an artifact of multi-unit, rather than single-unit activity. To examine this possibility, we analyzed the action potential waveforms of the same unit recorded in different epochs of the task. We performed an analysis using data from monkey N for each unit with two components (n = 34 units) using a binary classification problem for each unit using the sampled waveforms recorded during 2320 trials (58 targets, 40 repetitions per target). A Naive Bayes classifier was trained to discriminate spike waveforms according to two classes: occurrence before and after peak velocity. Waveforms from these two classes were sampled randomly with replacement to obtain a matching number of samples in each class. This was repeated 2500 times. Then, we computed the cross-validated decoding accuracy of the classification. The decoding accuracy was near 0.5 for each cell, suggesting that the spike waveforms did not change between the two epochs. Spike sorting errors are unlikely to account for tuning changes in our data.

2.7.2.3 Additional details on fitting the epoch-specific tuning model

Equation 2.6 describes the epoch-specific tuning model used to characterize segments of cosine tuning. This model relies on nonlinear least-squares fitting, which requires reasonable initialization parameters. We implemented a two-step procedure to obtain initialization parameters for each of the coefficients in Equation 2.6. In the first initialization step, we repeated our preliminary analysis by fitting a cosine tuning function at each time bin. As before, this yielded a time-series of preferred directions and associated modulation depths. We refer to the time-course of the estimated modulation depth as $\hat{m}_t$ (e.g. Figure 2.4B). The second step involved fitting $\hat{m}_t$ with a multi-component Gaussian function. This fit required initialization parameters as well: we obtained the local maxima of $\hat{m}_t$ using the "findpeaks" function in Matlab, and the amplitude and location of these peaks served to prime the multi-component Gaussian fit.

We never observed more than three prominent peaks in $\hat{m}_t$. Therefore, to choose the number of components to fit, we evaluated up to three versions of the fit, composed of 1, 2, and 3 Gaussian components, rank-ordered by amplitude. We then chose the optimal number of components based on the cross-validated mean-squared-error of each fit. Thus, the second step yielded the number of components to fit, as well as initial estimates of the parameters corresponding to the Gaussian-shaped modulation depth of each component. An initial estimate of each component's PD was then computed using the firing rates in the segment of time spanning $\mu_j \pm s_j$ (see Eq. 2.6, main text).

We asked whether the use of a 50 ms SD Gaussian smoothing kernel affected our ability to observe additional peaks in the modulation depth functions of each neuron. To address this question, we examined versions of firing rate profiles after smoothing with a 30ms SD kernel

and observed highly comparable results. Additionally, we estimated the firing rates from spike counts in 2 ms bins using a generalized linear model and natural regression splines, rather than using Gaussian kernel smoothing (20 ms bins). We controlled the level of smoothing by changing the number of equally spaced knots over time. The knots were spaced from 30 to 80 ms apart and the results were highly comparable to those obtained using Gaussian kernels. Spacing at smaller intervals increased the variance of the fit, but the number of high amplitude modulation depth peaks was not changed.

3.0 ACTIVITY IN PRIMARY MOTOR CORTEX RELATED TO VISUAL FEEDBACK[2]

3.1 INTRODUCTION

The primary motor cortex (M1) is thought to play a critical role in volitional movement, and the activity of M1 neurons is known to vary strongly with features of behavior. Georgopoulos et al (1982) demonstrated that most neurons in M1 are "tuned" to the direction of arm movement during reaching. These experiments showed that the relationship between a neuron's firing rate and the direction of movement is described by a cosine function. The robustness of this relationship has been taken as evidence that M1 "encodes" simple representations of movement. However, critics of representational models have argued that responses of single neurons can be complex, and simple models often do not account for the temporal details of neuronal activation (Churchland & Shenoy 2007). Further, the responses of M1 neurons covary with a wide range of behavioral parameters, beyond what would be captured by a unitary tuning model (Georgopoulos et al 1992; Kakei, Hoffman, & Strick 1999; Caminiti, Johnson, & Urbano 1990; Zhang et al 1997). Investigators have attempted to improve these simple representational models by

[2] This chapter was submitted for publication as Suway & Schwartz, Activity in primary motor cortex related to visual feedback.

including additional features of movement, but gains in explanatory power have been modest (Paninski et al 2004; Aflalo & Graziano 2006; Suway et al 2017). Although the early directional tuning work (Georgopoulos et al 1982; Schwartz, Kettner, & Georgopoulos 1988) used firing rates averaged over the entire movement with the implicit assumption of stationary preferred directions, a number of studies have revealed time-dependent processing in M1 during reaching (Georgopoulos et al 1989; Pellizzer, Sargent, & Georgopoulos 1995; Fu et al 1995; Zhang et al 1997; Churchland et al 2012; Velliste et al 2014; Rouse & Schieber 2016). Recent work (Suway et al 2017; Harpaz et al 2018) has shown that these changes in directional tuning may take place at discrete points in the behavioral task, suggesting that they are linked to changes in the state of the system. Identifying these state changes using M1 firing rates and linking them to behavioral features important for reaching can help us understand the factors that govern information processing in the motor cortex.

Recently, investigators demonstrated that M1 neurons abruptly change their firing properties when monkeys transition from planning a movement to executing the movement (Churchland et al 2010, Kaufman et al 2014, Elsayed et al 2016; Lara et al 2018). This finding was taken as evidence that the same population of neurons performs separate functions before and after movement initiation. Results from our lab, and others, indicate that similar abrupt transitions in firing pattern occur at multiple points during the movement (Suway et al 2017, Harpaz et al 2018). In contrast to a static relation to behavior, this temporal structure suggests M1 undergoes changes in functional state in conjunction with transitions between behavioral components. Some authors have referred to this concept as a "temporal parcellation scheme" (Johnson & Ebner, 2000), though the specific parameters that drive neurons remain a topic of

debate. Importantly, the sequences of neural states that can be recognized in motor areas are likely to depend on the specific behavioral components required by a task.

Reaching to an object or target is strongly dependent on vision. Subjects focus on the target they are reaching toward while the moving hand is registered in the peripheral visual field (Paillard 1982, Paillard 1996). Accurate target acquisition takes place as a series of visually guided sub-movements as the hand's image enters the macular retina (Soechting & Lacquanti 1981; Meyer et al 1988; Milner 1992; Novak, Miller, & Houk 2002). Experiments dissociating vision from movement have provided evidence for visuospatial coding in M1 and premotor cortex, particularly before the movement begins (Georgopoulos et al 1989; Pellizzer et al 1995; Zhang et al 1997; di Pellegrino & Wise 1993). However, relatively little is known about the relation between M1 activity and ongoing visual feedback during movement (although see Schwartz, Moran, & Reina 2004, and Inoue, Uchimura, & Kitazawa 2016). One possibility is that the sequence of neuronal states is determined, in part, by visual information related to the task. Here, we hypothesized that components of the M1 response during reaching might be related to visual feedback.

To test this hypothesis, we implemented several variations of the classical center-out reaching task. We first required monkeys to adapt to a perturbation that dissociated vision from movement. Although the arm movements were similar in the perturbed condition, we found clear components of M1 activity that changed after the monkey adapted to the visual perturbation. In roughly half the trials in both the normal and perturbed conditions, the monkey's reach trajectory was hidden from view before the movement began. By comparing neuronal responses across these trial types, we found that the visibility of the reach is important for driving a component of M1 activity. The relation between this activity and movement

changed after the perturbation, and the activity was much weaker when feedback was not provided. Responses following this pattern occurred toward the end of the movements, possibly consistent with a feedback signal. In fact, the timing of putatively feedback-related firing relative to reach initiation was remarkably similar to estimates of the minimum processing time for visual feedback from a movement (Keele & Posner 1968). Our findings confirm that M1 plays a highly diverse set of roles in motor control and suggest that M1 may be driven in part by visual feedback.

3.2 MATERIALS AND METHODS

3.2.1 Behavioral Task

Two adult male monkeys (*Macaca mulatta*, R and P) were trained to perform reaches with their right arms while viewing a virtual reality (VR) environment through a depth-displaying monitor (Virtual Window, Dimension Technologies Inc.). The monkeys could not see their own arms during the task; instead, they were shown a spherical cursor in VR representing their arm positions. Cursor radii were about 0.3 cm (all such measures refer to length in physical space). Position was tracked optically at 50 Hz using an Optotrak 3020 motion capture system (Northern Digital Inc.). The monkeys performed several variations of the classical center-out reaching task. In each task, monkeys began a trial by holding their arms steady for 500-700 ms in the center of the workspace volume, marked by a spherical target with a radius of 0.6 cm. A peripheral target of the same size then appeared, cueing the monkey to begin moving. Targets could appear at one of 16 evenly spaced positions around the home position, forming a circle

with a 6 cm radius. The targets were arranged in the vertical plane, though the task required 3D movement control; reaches in front of or behind a target were not rewarded. Monkeys were permitted 600-800 ms to reach the target and received a liquid reward for each successful trial.

We leveraged the VR paradigm to impose a visual perturbation during behavior. Each session was split into two blocks with a roughly equal number of trials. In the first block, no visual perturbation was applied ("standard block"). At the end of this block, we gradually altered the mapping from hand position to VR cursor position such that the movement direction of the cursor was rotated clockwise relative to the movement direction of the hand. This paradigm is commonly referred to as a "visuomotor rotation." The angle of rotation was increased over the course of 15-25 trials, and then remained fixed at its final value for the entire second block ("rotated block"). The final rotation angle varied day to day and could be one of 45°, 67.5°, or 90°. Each of these angles is an integer multiple of the target spacing (22.5°), which ensured the same physical targets were reached in both task blocks. The scene displayed in VR did not change in the rotated block, and monkeys were required to learn the new hand-to-cursor mapping through trial and error over task repetitions. Initially, monkeys made large angular errors that had to be corrected, but performance gradually recovered as they adapted to the perturbation. This adaptation period was usually completed within about 100 trials, which were excluded from analysis.

The inclusion of a visuomotor rotation was intended to provide a way of identifying "higher-order" neural responses, i.e. responses that don't covary strictly with motor output. A perturbation-related response in our task could be driven by at least two distinct processes. One process may be related to the perceptual adaptation that occurs early in the rotated block, usually completed within tens or hundreds of trials after the rotation is applied (Wise et al 1998;

Krakauer et al 2000). This adaptation is marked behaviorally by a gradual decrease in the angular error of the reach, and neurophysiologically by changes in neuronal activity relative to the unperturbed condition (Shen & Alexander 1997a). The processing of the scene displayed in VR during a reach may be another factor driving changes in M1 firing patterns. To dissociate these processes, we randomly selected 50% of the trials in both the standard and rotated blocks, and disabled the cursor display before the movement began, 140 ms after the target appeared ("invisible trials"). The monkeys were not cued before the cursor disappeared, and the cursor did not reappear until the start of the subsequent trial. The target remained visible even when the cursor was not. Center-out reaching tasks commonly require stopping and holding within a target zone, which is difficult without visual feedback. In the tasks described here, we therefore omitted the terminal hold period; monkeys were rewarded upon touching the target with the cursor.

3.2.2 Neuronal Recordings

Monkeys were chronically implanted with 96-channel microelectrode "Utah" arrays (Blackrock Microsystems), which were inserted into the arm area of the primary motor cortex. Monkey R was implanted with two arrays; units were recorded from a single array in monkey P. Extracellular voltage signals measured from the electrode arrays were amplified, filtered, and digitized using either a Plexon MAP system (Plexon Inc.) or a Blackrock Cerebus system (Blackrock Microsystems). Spike waveforms were sorted manually offline using Plexon Offline Sorter (OFS, Plexon Inc.). Unit isolation was judged based on waveform, cluster separation, inter-spike interval histogram, autocorrelation, and cross-correlation with other units on the same electrode. Only well-isolated units were saved for further analysis.

3.2.3 EMG Recordings

Intramuscular EMG activity was recorded from several arm muscles, including the anterior deltoid, medial deltoid, biceps, triceps, flexor carpi radialis, and extensor carpi ulnaris. EMG electrodes consisted of 38-gauge multi-stranded stainless-steel wires with Teflon insulation (Cooner Wire, Chatsworth, CA). A 25-gauge needle was used to insert the wires transcutaneously. Pairs of wires were used in each muscle for bipolar recordings. After each experiment, the arm was wrapped in a semi-rigid cast to preserve the electrode insertions for a few consecutive days. Raw EMG signals were differentially amplified and sampled at approximately 2 kHz using a Tucker-Davis Technologies recording system. The recorded signals were bandpass-filtered between 100 Hz and 1000 Hz, rectified, and lowpass filtered at 7 Hz to compute the envelope. Lastly, the EMG data were downsampled to 50 Hz to match our other data signals. For monkey R, neurophysiological data were collected simultaneously with EMG data. EMG data from monkey P were collected in sessions separate from those in which M1 unit activity was recorded. For monkey P, the wrist extensor was not consistently activated over trial repetitions and was excluded from analysis.

3.2.4 Data Preprocessing

We computed neuronal firing rates by counting the number of spikes in 20 ms bins and dividing by the bin width to yield spikes/s (Hz). Rates were square-root transformed (Ashe & Georgopoulos 1994; Moran & Schwartz 1999) and smoothed using a Gaussian kernel with a 50 ms SD. For each neuron, we subtracted the across-target mean at each time bin to give a target-dependent change in rate. This procedure allowed us to focus on tuning properties, rather than

on target-independent dynamics (Churchland et al 2012, Suway et al 2017). For visualization only, we applied "PC smoothing" (Churchland et al 2010; Churchland et al 2012). This was applied to neuronal responses separately within each task type to prevent artificially mixing response properties.

For each trial, we identified the time of the target's appearance, the onset of the movement, the peak velocity, and the offset of the movement. Movement onset and offset were defined as the times when the arm reached 20% of its maximum speed for each trial. Neuronal firing rates, hand kinematics, and EMG signals were normalized in time and aligned using these behavioral landmarks. This was accomplished by setting a fixed number of time bins between epochs and interpolating each trial to match this number (using the "pchip" function in Matlab, MathWorks, Natick, MA). The kinematics of each trial were inspected to ensure that reaches were straight and accurate, with bell-shaped velocity profiles. Excessively curved reaches (e.g. during the adaptation period between blocks) were discarded.

For monkey P, we analyzed four neurophysiological and behavioral datasets, as well as EMG data recorded over two days. For this monkey, we were only able to record meaningful neurophysiological data during sessions with a 45° visuomotor rotation due to degradation of the array signal for later sessions. We analyzed five neurophysiological and behavioral datasets from monkey R, including three sessions with a 90° perturbation, one session with a 67.5° perturbation, and one with a 45° perturbation. Two of the 90° sessions for this monkey also included simultaneously recorded EMG activity. Sessions from monkey R included 143 neurons per day on average (SD = 43.6 neurons, 715 total observed neurons). Sessions from monkey P included 35.5 neurons on average (SD = 5.8 neurons, 142 total observed neurons). Because so few neurons were observed daily for monkey P, we combined these datasets before analysis.

3.2.5 Tuning Curve Analysis

We computed cosine tuning functions using the recorded neuronal firing rates and hand movement directions. These were calculated using ordinary least-squares linear regression and the following cosine tuning model:

(3.1) $y = k + B \bullet \cos(\theta - \theta_{PD}) + \varepsilon$

where y is the estimate of a neuron's firing rate, k is the baseline rate, B is the amplitude of the tuning function (modulation depth), θ is the physical movement direction, θ_{PD} is the preferred direction (PD), and ε is the noise or fitting error. Since the firing rates were mean-subtracted (see above), the baseline term was always zero. The θ parameter refers to instantaneous hand direction, which was shifted back in time by 140 ms relative to neuronal responses. We also applied Equation 3.1 to EMG activations in place of firing rates. In that case, direction was shifted backwards by 100 ms. We note that preferred directions computed by this method are expressed relative to physical space, rather than VR space. This distinction is important when describing PDs before and after the visuomotor perturbation.

To compare PDs in the standard and rotated blocks, we computed the angular difference between the two tuning curves. The statistical significance of this difference was found using a bootstrap procedure. Trials were resampled with replacement 1000 times, and the angular difference was computed for each iteration. We then computed the 95% confidence interval of this difference over the 1000 iterations. If the confidence interval overlapped 0, the PD difference was not considered to be significantly different. This test is thus two-tailed, with an alpha level of 0.05. When comparing PDs, we required that the firing rates during both blocks were fit by a cosine function with an R^2 of at least 0.5. We used a significant change in PD as an

operational definition for patterns of neuronal firing potentially related to adaptation to the perturbation.

3.2.6 Trial Segmentation

For some analyses, trials were segmented into several epochs. We used factor analysis (FA) to identify independent components of firing across time for all recorded neurons. Typically, FA is applied to neuronal data by treating each neuron as a "dimension," and reducing the dimensionality to a few summary components. Here, we instead used time bins as the dimensions input to FA, with each neuron contributing observations along those dimensions. This process may be familiar as a spike-sorting method where principal components calculated from a temporal sample of voltage are used to separate different spike waveforms. Here, we use the method to find temporal features of neurons' firing rate pattern that covary across reaches to different targets. These temporal features occurred sequentially and were used to segment the trial. Since we were interested in tuning patterns regardless of magnitude, we normalized firings rates by their range at each time bin prior to applying FA. In practice, this step did not substantially alter the results.

3.2.7 Reverse Regression Analysis

We used reverse regression to find linear combinations of the neural data that could decode behavioral variables of interest. More common decoding methods such as the Population Vector Algorithm (PVA) or Optimal Linear Estimator (OLE) rely on explicit tuning model fits for each recorded unit, but here we sought to avoid these assumptions about encoding. Reverse

regression is a simple multiple regression approach that treats each neuron's firing rate as an explanatory variable, and a variable of interest as the dependent or response variable (Kass, Ventura, & Brown 2005; Inoue et al 2018). In our case, the variables to decode were the *x*- and *y*-components of the displayed direction. The following equations describe the relationship:

(3.2) $V_x = k_x + B_x R + \varepsilon$

(3.3) $V_y = k_y + B_y R + \varepsilon$

where V_x and V_y are the *x*- and *y*-components of the displayed direction, k_x and k_y are constant offsets, B_x and B_y are 1-by-N vectors of regression weights (one weight for each of the N neurons), R is a matrix of firing rates with size N-by-(time*conditions), and ε is the noise or fitting error. The regression weight vectors represent scaled axes in neural space, chosen such that orthogonal projections onto these axes best correlate with the *x*- and *y*-components of the displayed movement direction. The free parameters in Equations 3.2 and 3.3 were fit using ridge regression rather than ordinary least squares regression. Ridge regression is useful in the case that regression terms are correlated with one another, as is often true for firing patterns of different neurons. Typically, the regularization parameter for ridge regression is chosen to maximize cross-validation performance. However, we found that performance was very similar over a large range of values, with no clear maximum. We therefore arbitrarily chose a regularization strength of 1. That is, in the regression loss function, we assigned equal weighting to the sum of squared residuals and the regularization term.

We applied Equations 3.2 and 3.3 to neural data from the start of the trial until 100 ms after movement offset. The *V* variables were set to zero until 50 ms after the target appeared. These pre-target time bins were included as an additional constraint on fitting, which ensured that the decoder output was near zero when there was not yet any direction displayed. As with

the trial segmentation procedure described above, we normalized firings rates by their range at each time bin prior to fitting, though this step was not critical.

3.3 RESULTS

3.3.1 Single-unit directional encoding changes after visuomotor adaptation

Neurons in M1 are directionally tuned. Here, we used this property to examine possible effects of visuomotor adaptation on neural firing. Monkeys made reaches to the same physical targets before and after adaptation (Figure 3.1 A-B). An example neuron recorded during those trials was found to fire strongly during upward reaches (Figure 3.1C, blue rasters). If this tuning property was related to the direction of movement in physical space, one would expect to see the same response for an upward reach before and after the monkey adapted to the perturbation. However, during the rotated block, this neuron fired at a much lower rate for upward reaches compared to the standard block (Figure 3.1D, blue rasters). One possibility is that this neuron was directionally tuned in VR space, rather than physical space. In that case, one would expect the highest firing rate during trials when the target cue (and displayed reach trajectory) was upward, even if the monkey didn't physically reach in that direction. This tuning property is, in fact, what was observed for the example response in Figure 3.1. In this session, the perturbation was 90° clockwise. Therefore, in the rotated block, a target cue upward always elicited a reach to the left (Figure 3.1B, red traces and red circle). These leftward reach trials had rates that were similar to those for upward reaches in the standard block (red vs. blue rasters, Figure 3.1 C-D), suggesting that cursor-related information was encoded in this neuron's activity.

We next examined the smoothed, trial-averaged firing rate profile of this neuron across targets, and found the response had a similar shape and timing in both blocks (Figure 3.1 E-F, different colored traces). Note that the color scheme used for firing rate plots was defined relative to physical movement direction in both blocks; the encoding properties previously seen for the raster plots can therefore be recognized qualitatively in the ordering of rates across targets in panels E-F (e.g. light blue vs. dark red traces). Across-target rates averaged over the period marked in Figure 3.1E (black bar) were fit well by cosine functions in each task block (Figure 3.1G). The preferred directions (PDs) of the fits, computed relative to physical space, had an angular difference of about 105° between blocks. The PD in the standard block was subtracted from that in the rotated block, rather than the other way around, which yielded a positive difference for responses that were sensitive to the visuomotor rotation like the example in Figure 3.1. A negative change in PD between blocks was also possible, though we will show below that this pattern was less common.

Using a bootstrap procedure, we evaluated the across-trial consistency of the PD change shown in Figure 3.1G by resampling individual trials with replacement and recomputing PD change for each bootstrap iteration. The 95% confidence interval over bootstrap iterations was [96°, 113°], meaning the PD change was consistent over trials within about 17°. This consistency suggests the effect demonstrated in panels C-G is likely related to the task design. To better illustrate this relationship, we replotted the tuning curves in each block relative to VR cursor movement direction, rather than hand direction (Figure 3.1H). The similar phase of the cosine functions in panel H indicates that tuning was similar before and after adaptation when computed relative to cursor direction, in contrast to the large difference observed with respect to hand direction. This observation further supports the idea that this neuron's tuning was driven

by cursor-related information. For subsequent figures, firing rates and tuning curves will be plotted with respect to physical movement.

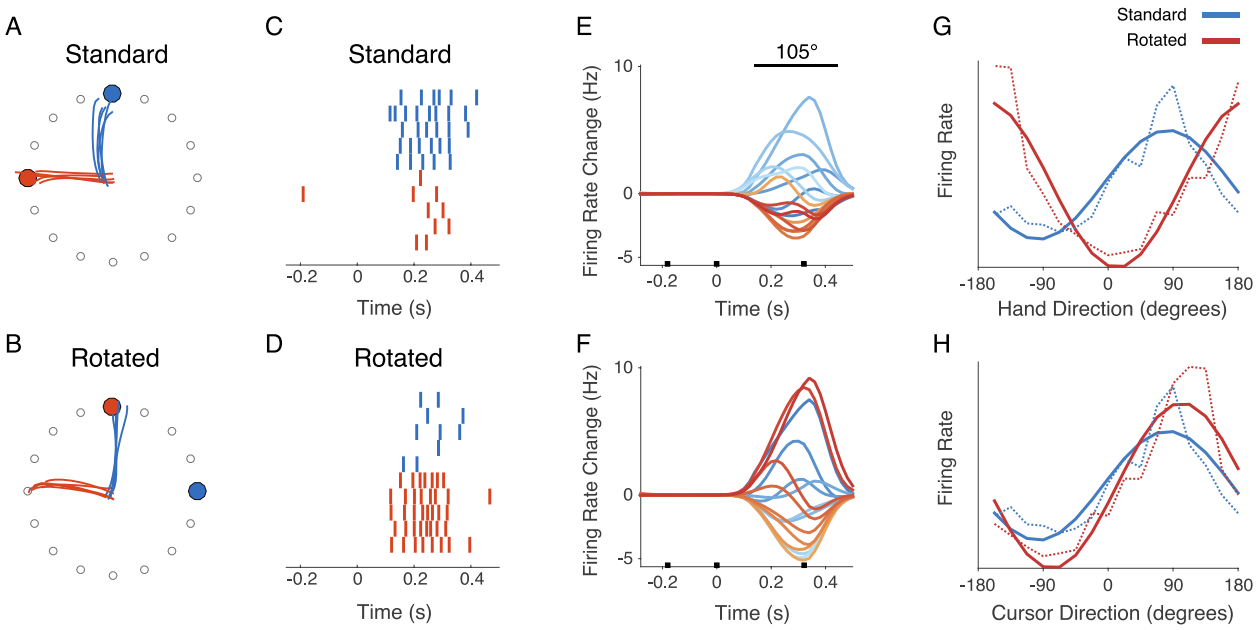

Figure 3.1 Firing rate properties change after visual adaptation.

A) Five repetitions of upward reaches (blue traces) and five of leftward reaches (red traces) in the standard block. The traces represent the physical location of the hand in space, and the red and blue circles represent the displayed position of the targets in VR, which match the hand endpoints in this block. B) Same as in A but for the rotated block (90° in this session). Note that the positions of the displayed targets are rotated clockwise relative to the hand endpoints in this block. C) Raster plot for one neuron during the same trials in A, using the same color scheme. Time 0 marks movement onset. D) Same as in C but for the trials shown in B. E) Smoothed, trial-averaged firing rates for the same neuron during the standard block, with responses plotted for all 16 reach directions (different colored traces). Black square markers along the abscissa mark the time of target show, the onset of movement, and the offset of movement, respectively. F) Same as in E but for the rotated block. The color scheme is defined with respect to hand direction, not cursor direction. G) Tuning curves relative to hand direction, computed for the responses in E and F averaged over the period marked by the black bar in E. Solid lines are model fits (Eq. 3.1), dashed lines are observed rates. Number over black bar in E reports the angular PD difference between cosine fits. H) Same as in panel G but with tuning expressed relative to cursor direction rather than hand direction. The PDs of the responses in panels E-F were very similar when expressed in cursor-centric coordinates.

The activity displayed in Figure 3.1 had only a single modulation peak. However, M1 neurons often fire with multiple modulation peaks during reaching (Churchland & Shenoy 2007; Churchland et al 2012; Suway et al 2017). We observed many neurons with multiple modulation peaks in our datasets. Interestingly, individual modulation peaks of single neurons often showed a variety of patterns following adaptation (Figure 3.2). Sometimes patterns displayed an early modulation component related to the adaptation, while a subsequent component remained unchanged (Figure 3.2 A,C). Other examples showed the opposite temporal pattern (Figure 3.2 B,D). Figure 3.2E and 3.2F show activity patterns with both components affected by adaptation, but each to a different degree. Figure 3.2G shows a unit response with an early component unaffected by adaptation, and a second component with a highly reduced amplitude after adaptation. These observations may support the hypothesis that distinct peaks of modulation can correspond to distinct encoding epochs. We also observed responses with no significant relation to the adaptation (not shown), indicating a constant encoding of the physical movement direction. Although the example response shown in Figure 3.1 had a tuning change of similar magnitude to the visuomotor perturbation, the examples in Figure 3.2 indicate the magnitude of change was variable across responses.

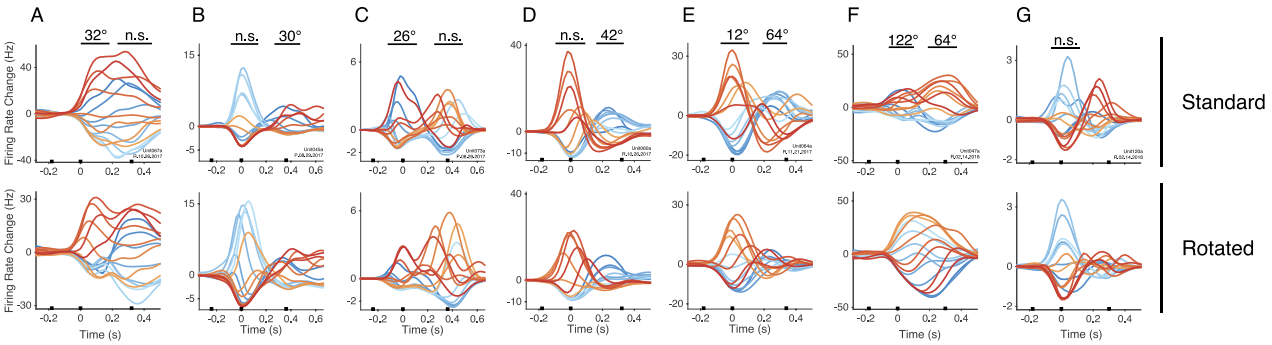

Figure 3.2 Components of single unit responses are independently affected by visuomotor adaptation.

Each column represents the response of a single neuron for all reach directions in the standard block (top row) and rotated block (bottom row). The responses in panels A-D were from 45° sessions; the response in panel E was from a 67.5° session; the responses in F and G were from 90° sessions. Black bars above response components mark time periods where we applied the same tuning comparison method as in Figure 3.1E-G (i.e. tuning compared between blocks, computed using hand direction). Numbers above bars show the angular difference in PD between blocks, and imply that this difference was statistically significant ($p < 0.05$); n.s. indicates no significant difference ($p >= 0.05$).

3.3.2 Adaptation effects are widespread and depend on rotation magnitude

We next sought to summarize the effects described in the previous section for each neuron in our dataset. Since the effect of adaptation had a strong time-dependent effect on single-neuron firing rates (Figure 3.2), we first segmented trials into several epochs using factor analysis (FA, see Methods). We chose to identify four factors from our data, representing prominent firing patterns over four epochs (Supplemental Figure 3.9). The epochs were distributed roughly evenly through the trial. Previously, our group identified similar epochs based on the timing of modulation peaks of single neurons (Suway et al 2017). Using that approach, we identified three distinct epochs roughly corresponding to the second, third, and fourth epochs shown in Supplemental Figure 3.9. The additional epoch in the current FA procedure targeted activity

during the reaction time and early movement. This factor explained the least amount of variance of the four, but we retained it for additional temporal resolution. The trial epochs corresponded to similar behavioral events in both monkeys, although the first two epochs for monkey P were slightly shifted forward in time compared to monkey R (Supplemental Figure 3.9A vs B). This could possibly be related to performance differences between monkeys; reaction time and reach duration were both longer for monkey P.

Once the epochs were defined using FA, we considered firing rates averaged over 100 ms -windows centered on each factor's peak. Within each of the four windows, we identified neurons with responses well-fit by cosine functions ($R^2 > 0.5$) in both the standard and rotated blocks. We then computed PD differences between the two blocks for each neuron and epoch and produced an angular histogram of these differences (Figure 3.3). The solid black traces in Figure 3.3 represent statistically significant changes, while the dashed gray lines show non-significant changes. We found that the PD-difference between standard and rotated blocks was related to the magnitude of the visuomotor perturbation. Differences in PD tended to be much larger for 90° sessions than for 45° sessions (Figure 3.3 A vs B). This property can also be seen in the single- unit responses in Figure 3.2 (A-D vs E-F). The proportion of neurons with a significant PD change (regardless of magnitude) after adaptation was also larger for 90° sessions (black vs. gray histograms, 3.3 A-B). This may be due to the generally larger changes in PD in those sessions; assuming similar statistical power to detect these changes, one would expect a higher detection rate for larger changes. The majority of PD differences were in the positive direction, like the responses shown in Figures 3.1 and 3.2.

For responses following this pattern, the most common observation was a PD change of roughly half the perturbation angle. This intermediate-valued change could possibly reflect an

incomplete adaptation process. To test this possibility, we re-computed PD changes using only trials from either the first or second half of the rotated block. If the intermediate-valued PD changes reflected incomplete adaptation, we would expect those changes to more closely match the perturbation angle following additional exposure, i.e. during the second half of the block. However, we found the distributions of changes were very similar to those in Figure 3.3 regardless of which half of the block trials were sampled from. This consistency suggests the intermediate-valued PD changes reflect "mixed" selectivity for visuospatial and physical movement parameters (Lurito, Georgakopoulos, & Georgopoulos 1991; Shen & Alexander 1997a), rather than an incomplete adaptation process.

To compare M1 responses to the concurrent motor output, we repeated the tuning comparison procedure for EMG activity recorded from each muscle in our dataset (Figure 3.3C). EMG data were analyzed during the epochs defined by the neuronal data. However, we also re-computed the factors from the EMG data and found three clear epochs around movement onset, peak velocity, and movement offset. A fourth factor computed from the EMG data had small loadings and contained several noisy peaks, suggesting three factors were sufficient for describing these data. When PDs of EMG data were computed within the first three EMG epochs, results were not substantially different than when computed within the epochs defined by the neuronal data. To provide a more direct comparison to the analysis of neuronal data, we retained the neuronal epochs for analysis of the EMG data. The PDs computed from EMG activity tended to be similar between the two task blocks, though some differences were observed. Interestingly, those differences were almost always in the negative direction, opposite from most changes in neuronal PDs. There was little difference in this pattern between individual muscles, with two exceptions: 1) for the rare case that a significant PD change was

observed for the anterior deltoid, the change was slightly positive; 2) a significant PD change in either direction was almost never observed for the medial deltoid. For both 45° and 90° sessions, a small subset of the neuronal responses also had negative PD differences between blocks, with a similar magnitude as those observed for EMG activity. This subpopulation may therefore be considered "muscle-like" (Kakei, Hoffman, & Strick 1999). Responses with a positive PD change between blocks may be considered related to or affected by the visuomotor perturbation.

As a control to confirm the observed PD changes were task-driven, we performed a similar PD analysis using only data from the standard block. Trials within that block were partitioned into two groups, and the PD difference between the two groups was computed for each trial epoch. We then repeated the procedure 1000 times with new randomly chosen partitions and displayed the PD differences as angular histograms (Supplemental Figure 3.10). We found that these differences were always near-zero, suggesting the larger changes seen between blocks (Figure 3.3) were related to the task.

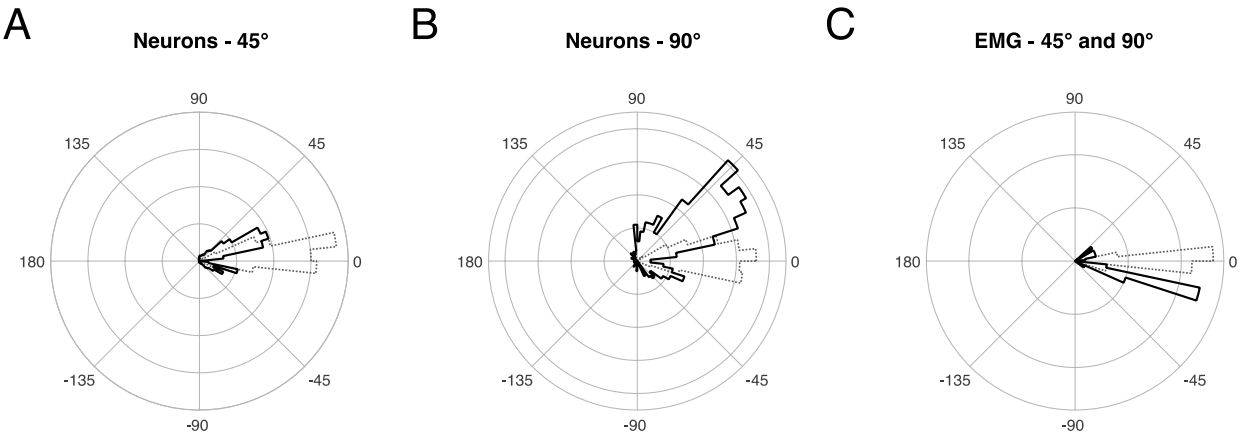

Figure 3.3 Neuronal PD changes depend on rotation magnitude, but muscle PDs are similar after adaptation

Black traces show significant PD changes, gray dashed traces show non-significant changes. Neurons and muscles could contribute up to four counts to a histogram (one for each epoch). A) PD differences for 45° sessions. Data from both monkeys were comparable and were pooled. B) Same as in A but for 90° sessions (monkey R only). C) Same as in A and B but for EMG activity instead of neuronal responses. Data from both monkeys (and rotation angles) were comparable and were pooled.

3.3.3 Features of M1 responses are related to reach visibility

In the previous sections, we described substantial effects of visuomotor adaptation on the firing patterns of M1 neurons. However, these results potentially represent a mixture of properties related either to the adaptation process, or to the mismatch between the observed movement and the physical movement. Here, we compare "visible" to "invisible" trials and show that a component of firing related to the perturbation is also related to visibility. Figure 3.4 shows the firing rates of two example neurons during the four different task conditions: the top row shows rates in the standard block, and the bottom row shows rates in the rotated block; for each subpanel, the left column shows rates in the visible trials, and the right column shows those for the invisible trials. For the response in Figure 3.4A, during visible trials, we observed one modulation peak that changed PD after the adaptation (left column). The rate of this neuron's firing was greatly decreased when reaches were not visible (right column). Another example is shown in Figure 3.4B. This response showed two modulation peaks during the visible trials (left column). Both peaks were affected by the visual perturbation, although the second peak was affected to a greater degree (46° vs 70°). Interestingly, the second modulation peak was absent when the cursor was invisible (right column). The first modulation peak was similar regardless of cursor visibility, and a similar PD change was observed following adaptation (38° vs 46°, black bars above peaks).

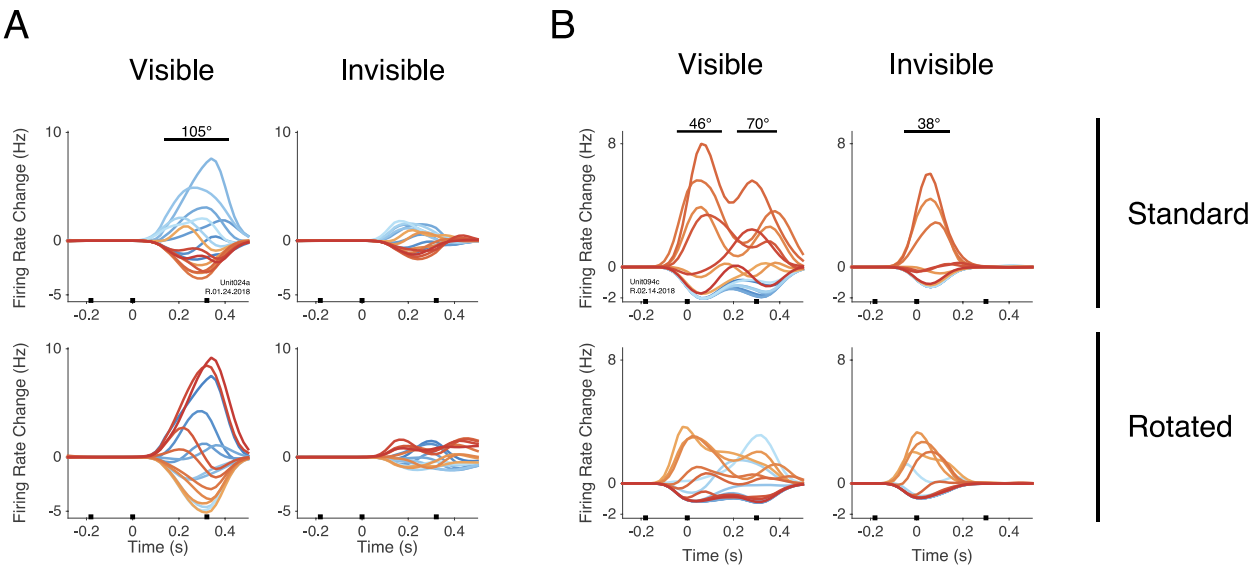

Figure 3.4 Reach visibility affects neuronal firing.

A) Firing rates of the same unit shown in Figure 3.1. The top row shows the standard condition; the bottom row shows the rotated condition. The left column shows visible trials; the right column shows invisible trials. The amplitude of firing was strongly decreased for this neuron when reaches were not visible. B) Same as in A but for a different neuron. This example response exhibited two modulation peaks when reaches were visible, but only one when reaches were not visible.

The example responses in Figure 3.4 exhibit two interesting patterns. First, for visible trials, there was a strong effect of the perturbation on their tuning functions late in the movement. Second, this tuning was weak or absent when feedback was not displayed. We asked if this phenomenon was common across the subpopulation of neurons with perturbation effects on firing rates. Because the visibility-related changes appeared to occur within a specific time span, we again utilized the trial segmentation procedure shown in Supplemental Figure 3.9. Within each of the four trial epochs, we computed the percentage of neurons with a significant PD change in the positive direction between blocks. These percentages were computed separately for visible and invisible trials and are plotted in Figure 3.5A (blue lines, solid and dashed respectively). The factor loadings marking the four trial epochs are also displayed for reference,

with scale on the right ordinate. If the visibility-related firing patterns shown in Figure 3.4 were common across the population, we would expect to find more neurons with significant PD changes between blocks for visible trials compared to invisible trials, particularly late in the movement. We found that regardless of reach visibility, the percentage of neurons with a significant PD change was large early in the trial, during the reaction time (first epoch). This percentage dropped after the movement was initiated (second epoch). The third and fourth epochs occurred around peak movement velocity and the offset of movement. During those epochs, the percentage of neurons with a PD change after adaptation increased when the monkey could see the cursor motion, but decreased when he could not. Activity related to the perturbation but not to visual feedback may be a hallmark of the adaptation process. Activity related to both task parameters, occurring towards the end of behavior, may be consistent with a feedback signal.

The example responses in Figure 3.4 showed decreased modulation amplitude around the end of the movement when reaches were not visible. If this pattern was common, that might explain why fewer PD changes were observed towards the end of invisible trials. To address this possibility, we focused on the fourth epoch of visible trials, and selected neurons with a positive PD change after adaptation. We then computed the percent change in modulation depth between those trials and invisible trials (Figure 3.5B). We found that tuning amplitude commonly decreased in the invisible trials, relative to visible trials, with the firing of many units showing a nearly complete loss of modulation. This pattern is consistent with the example responses shown in Figure 3.4.

The results in Figure 3.5 A-B were computed using a 90° session recorded from monkey R. A strikingly similar time course of effects was observed for monkey P (45° sessions, Figure

3.5C). The overall percentage of PD changes after adaptation was smaller, as expected for 45° sessions (Figure 3.3 A-B). Changes in modulation depth were also similar between monkeys: neurons with a positive PD change in epoch four of visible trials commonly showed lower tuning modulation during invisible trials (Figure 3.5D). From these results, it appears that reach visibility is important for driving M1 activity toward the end of a reach. This component is affected by the perturbation during visible trials, and is not as strongly modulated when feedback is absent. However, we also found other patterns. From Figure 3.5, it is clear that some responses maintain a relation to the perturbation with strong modulation late in the movement, even when the movement was not visible. Examples of different types of patterning are illustrated in Supplemental Figure 3.11.

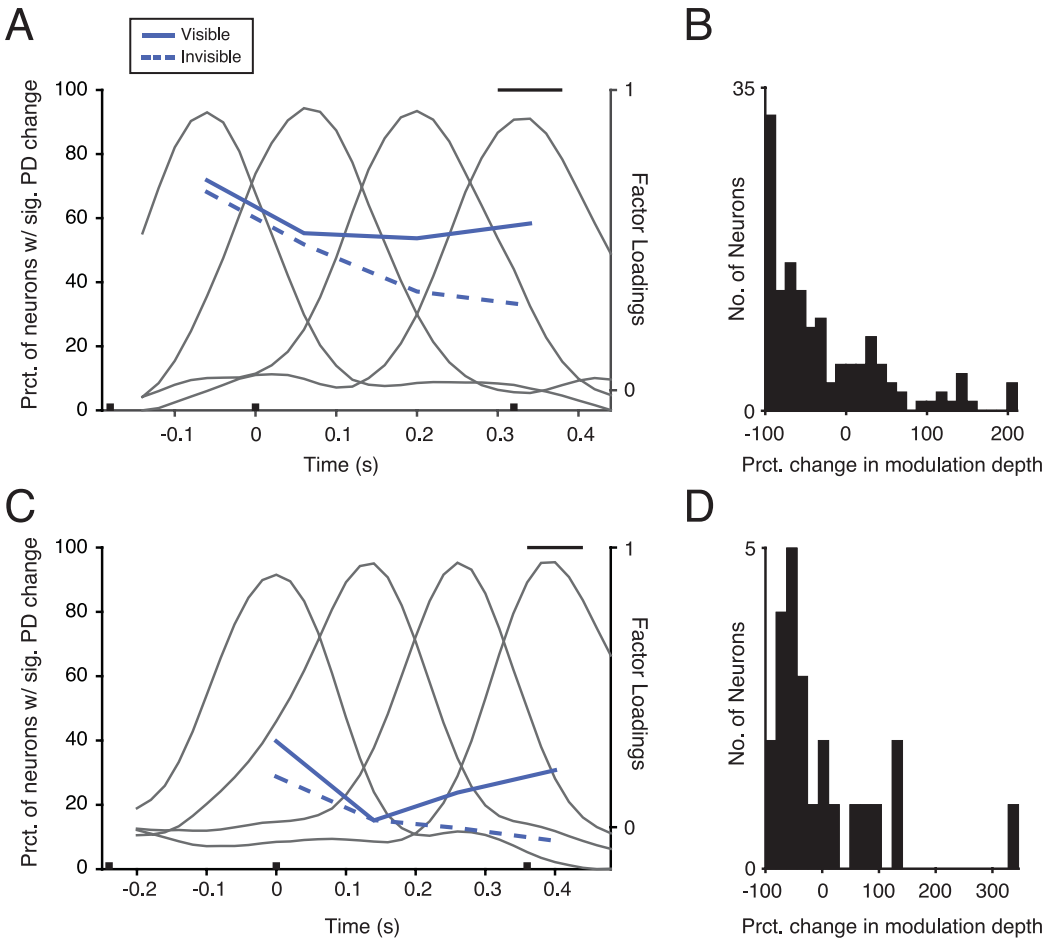

Figure 3.5 Visibility-related tuning changes occur towards the end of movement.

A) Percentage of neurons exhibiting PD change ($p < 0.05$) in the positive direction between standard and rotated blocks, computed for visible and invisible trials (blue traces, solid and dashed respectively. Scale on left ordinate). Percentage was computed within the four epochs (black Gaussian-shaped traces, scale on right ordinate). One 90° session from monkey R is shown. B) Change in modulation depth between visible and invisible trials for epoch #4 (black bar in panel A). Histogram includes responses with a positive significant PD change between blocks in the visible condition. Data in this panel were pooled over 90° sessions from monkey R. C) Same as in A, but for monkey P (45° sessions). D) Same as in B, but for monkey P.

3.3.4 Decoding cursor direction captures activity related to reach visibility

In the previous section, we measured each neuron's PD before and after adaption to identify firing patterns potentially related to reach visibility. We then determined how those patterns changed when monkeys could not see their reaches. Here, we extend this approach to a population-level analysis that does not rely on tuning models. Since many responses were affected by the perturbation, we hypothesized that the displayed movement direction could be extracted or decoded from the population's activity. We then reasoned that this procedure would perform poorly during the invisible trials if any component of the response was related to reach visibility. Reverse regression was used to build a decoder by finding linear combinations of the neural data that captured the displayed movement direction (see Methods). All units were included in the analysis regardless of their relation to the visuomotor perturbation; that is, we did not pre-select responses with significant adaptation-related tuning changes. Instead, our regression procedure naturally weights responses with this property more heavily than those without. This procedure yielded two linear combinations of rates: one for decoding the *x*-component, and one for the *y*-component of the displayed direction.

We identified these decoder weights using only the visible trials and verified that the displayed direction was decoded accurately. Figure 3.6 A-B shows one set of decoder outputs, in this case for the *y*-component of displayed direction. In the standard condition, shortly after the target appeared, decoder output was positive for upward reaches (Figure 3.6A, blue traces), and negative for downward reaches (red and orange traces). This is the expected pattern for a representation of the *y*-axis. In the rotated condition, the activation pattern changed (Figure 3.6B). Note that the color scheme of the different colored traces is defined with respect to physical space, such that blue traces (for example) always indicate upward movement, even if

the VR cursor was moving in a different direction. The example in Figure 3.6 shows data from a 90° session; in the rotated block, leftward physical reaches were associated with upward displayed movement. Since the decoder should track the displayed direction, we would then expect positive output for leftward reaches. Indeed, during the rotated block, output was positive-going for leftward reaches (Figure 3.6B, dark blue and red traces), and negative-going for rightward reaches (light orange and light blue traces). Although the same decoder (built from the visualized cursor directions) was used in both Figure 3.6A and 3.6B, the decoded components were nearly orthogonal, matching the *y*-component of displayed movement direction closely in both blocks: the five-fold cross-validated R^2 values were 0.91 and 0.94 in the standard and rotated blocks, respectively. A similar result was found for the *x*-component decoder, yielding cross-validated R^2 values of 0.90 and 0.94 in the standard and rotated blocks. The high degree of accuracy over cross-validated iterations indicates these results were consistent over trials when the movement was visible.

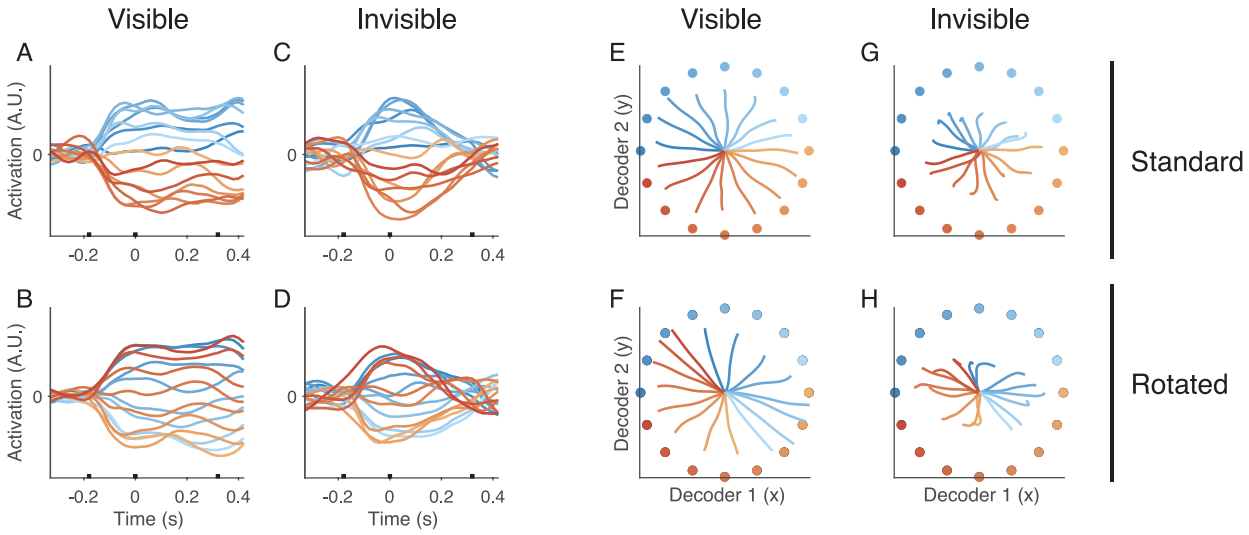

Figure 3.6 Accurate decoder readout of displayed direction depends on cursor visibility.

A) Decoder output for the y-component of displayed direction, shown in the standard block. Different colored traces show output for reaches to each target. B) Same as in A, but during the rotated block. Note the ordering across targets has changed, though the decoder weights were the same. C,D) Same as in A,B, but for invisible trials. Output amplitude became weaker and ordering of targets became less clear late in the movement. E) Integrated trajectories from the decoder shown in A-D (ordinate) and a second decoder for the *x*-component of displayed direction (abscissa). This panel shows trajectories during the standard block, for visible trials. Different colored circles around the edges mark the physical target position. F) Same as in E, but for the rotated block. Note the trajectories have rotated 90° clockwise in this block, though the physical targets are the same. G,H) Same as in E,F, but for invisible trials. Trajectories during invisible trials were shortened and distorted towards the end of movement, relative to visible trials.

Next, the rates recorded during invisible trials were used in the decoder constructed from the visible trials (Figure 3.6 C-D). For those invisible trials, the decoder result was comparable to that for visible trials shortly after the target appeared, during the early part of the movement. Over that time span, the across-target output pattern was comparable in both trial types, as was the output amplitude. However, later in the invisible trials, decoder output decreased in amplitude and the ordering across targets became less clear (Figure 3.6 C-D). This pattern matches closely with that described for the tuning analysis of single-unit responses in Figure 3.5: effects related to the perturbation were strong early in the trial regardless of reach visibility (possibly an effect of adaptation), but weaker late in the trial when reaches were not visible (possibly related to lack of feedback).

Thus far, we have focused on only one of the decoded directions, in this case, the *y*-component. To better visualize the output of both decoders (one for the *x*- and one for the *y*-component), we integrated their values over time and plotted the resulting trajectories for each condition (Figure 3.6 E-H). This integration procedure is very similar to that used in the classical population vector algorithm (Schwartz 1994; Schwartz & Moran 1999). In the standard block, the trajectories were relatively straight and accurate (Figure 3.6E). A similar result was found during the rotated block, though some skewing of the trajectories is apparent (Figure 3.6F). In this block, the trajectories were rotated nearly 90° relative to those in the standard block, closely matching the displayed direction. For example, the lightest orange trace points rightward in Figure 3.6E, and downward in Figure 3.6F. Following the same procedure for invisible trials initially resulted in trajectories similar to those in the visible trials (Figure 3.6 G-H). However, in both the standard and rotated blocks, the trajectories did not extend as far, and became distorted toward the end of movement.

We next evaluated the consistency of the results for invisible trials and utilized the trial segmentation from previous analyses to examine differences across epochs (Figure 3.7). Because the targets were arranged in a circle, decoder output for each axis should be sinusoidal over targets. The example shown in Figure 3.6C decoded the *y*-axis in the standard block, so the peak of this sinusoid should be centered on 90° (upward). Indeed, averaged over the period spanning epochs one and two (see Supplemental Figure 3.9), the across-target activation closely fit a sine function, peaking near 90° in the standard block (Figure 3.7A, blue traces). The phase of this sinusoid was shifted during the rotated block, and was instead centered near 180° (leftward, Figure 3.7A, red traces). The difference in phase between these two sinusoidal patterns was found to be consistent over trials. We used a bootstrap procedure to repeatedly resample invisible trials with replacement before repeating the procedure and re-computing the phase difference during epochs one and two (Figure 3.7B). This difference was about 85° (95% CI = 78°, 92°). In contrast, when we analyzed the period of time during epoch four, across-target activation was poorly fit by sinusoidal functions, and was very inconsistent across trials (Figure 3.7 C-D).

The results presented in Figure 3.6 and Figure 3.7 A-D were computed using neural data from a 90° session performed by monkey R. For monkey P, we recorded similar data for 45° sessions, but were unable to record meaningful neurophysiological data during 90° sessions. In principle, the same analysis can be applied to the 45° sessions. In that case, however, physical and displayed movement are not perfectly orthogonal. The correlation between them ($\rho \approx 0.7$) likely causes the decoder output to represent more of a mixture of activity related to displayed and executed movement. Nonetheless, the same analysis is presented in Figure 3.7 E-H for monkey P, and a similar trend emerged. These panels show decoder output for the *x*-component

of displayed direction. During the period spanning epochs one and two, across-target activation was well-fit by a sinusoid centered near 0° (rightward) in the standard block (Figure 3.7E, blue traces). The phase of this sinusoid shifted during the rotated block (Figure 3.7E, red traces). The difference in phase between blocks was consistent, and centered on 25°, undershooting the 45° rotation (Figure 3.7F, 95% CI = 12°, 37°). This undershoot could be due to the smaller number of rotation-sensitive responses in 45° sessions (Figure 3.3, A vs. B). When the analysis was applied during the fourth epoch, this phase change between blocks was not observed, instead centered on 0° (Figure 3.7 G-H, 95% CI = -19°, 21°).

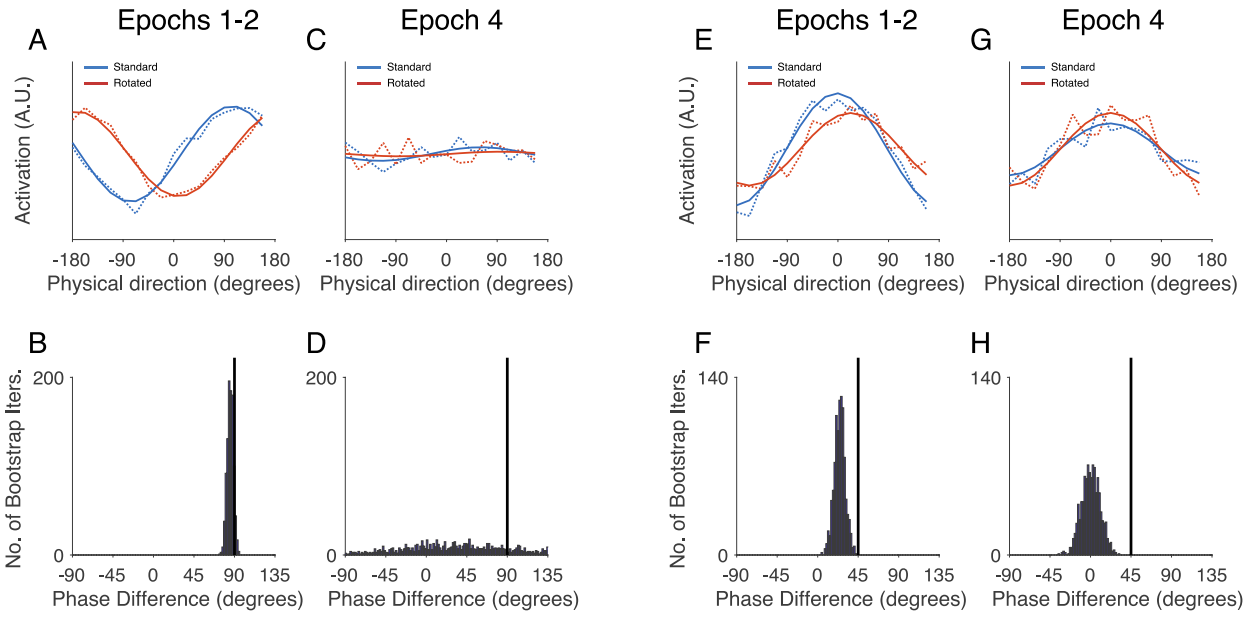

Figure 3.7 Decoder output is inconsistent late in the movement for invisible trials.

A) Same decoder output shown in Figure 3.6C and 3.6D, averaged over epochs 1 and 2, and plotted against physical reach direction. Blue traces are from the standard block (i.e. 3.6C), red traces from the rotated block (i.e. 3.6D). Dashed traces show observed data, solid traces show model fits (Eq. 3.1). B) Phase difference between the two model fits in A, shown over bootstrap iterations. The difference was consistently near 90°. C,D) Same as in A,B, but averaged over epoch 4. Late in invisible trials, the readout was noisy and inaccurate. E-H) Same as in A-D, but for monkey P (45° visuomotor rotation). These panels show readout for the *x*-component of displayed direction.

During the fourth epoch for monkey P, across-target decoder output was slightly noisier, lower in amplitude, and less consistent across trials than during epochs one and two. However, compared to the results for the 90° session, these effects were more moderate (compare Figure 3.7 C-D to G-H). It may be that these decoders do not as fully dissociate displayed and executed movement when using data from 45° sessions. Similar results were observed for a 45° session from monkey R, including a slight undershoot of the rotation angle between blocks for the early epochs. Broadly speaking, however, the results from both monkeys and rotation angles are in line with the single-neuron analyses presented in Figure 3.5: late in the movement, a component of the response related to the perturbation is less prominent if the movement is hidden from view.

3.3.5 Physical reach direction is well-represented in firing rates regardless of task condition

Although many neurons responded in different ways before and after adaptation, and in visible and invisible trials, we wondered if signals related to the physical movement could still be extracted accurately in each case. We therefore repeated the analysis in the previous section but used the reverse regression procedure to decode the physical reach direction, rather than the displayed reach direction (Figure 3.8). As before, all units were included regardless of whether their responses had adaptation-related tuning changes. Regression was again performed using only visible trials, and we again used five-fold cross-validation to confirm the accuracy of this readout. Accuracy was relatively high: the cross-validated R^2 values for the *y*-component were 0.89 and 0.93 in the standard and rotated blocks, respectively. For the *x*-component decoder, cross-validated R^2 values were 0.87 and 0.93 in the standard and rotated blocks.

Figure 3.8A shows decoder output for the *y*-component of physical reach direction. During the standard block, physical and displayed direction were identical; decoder output should then be comparable to that shown in Figure 3.6A. Indeed, the ordering of outputs across targets was similar (blue traces were positive-going, red and orange were negative-going). However, unlike the decoder for displayed direction, the output pattern shown here did not change after adaptation (Figure 3.8B, compare to 3.6B). Further, output in both blocks was relatively consistent even during invisible trials (Figure 3.8 C-D). One exception is that activation was weaker for invisible trials in the standard block, but only slightly (Figure 3.8C). When we integrated these activations, as well as those for the decoded x-component, the trajectories were relatively straight and accurate in each condition (Figure 3.8E-H). For invisible trials, the trajectories were slightly shorter for some targets, but not to the extent observed for readout of displayed direction (Compare 3.8G-H to 3.6G-H). These results suggest that the physical movement is consistently represented in M1 activity regardless of visual context, in line with the finding that a considerable portion of responses do not change PD in the rotated block (Figure 3.3 A-B).

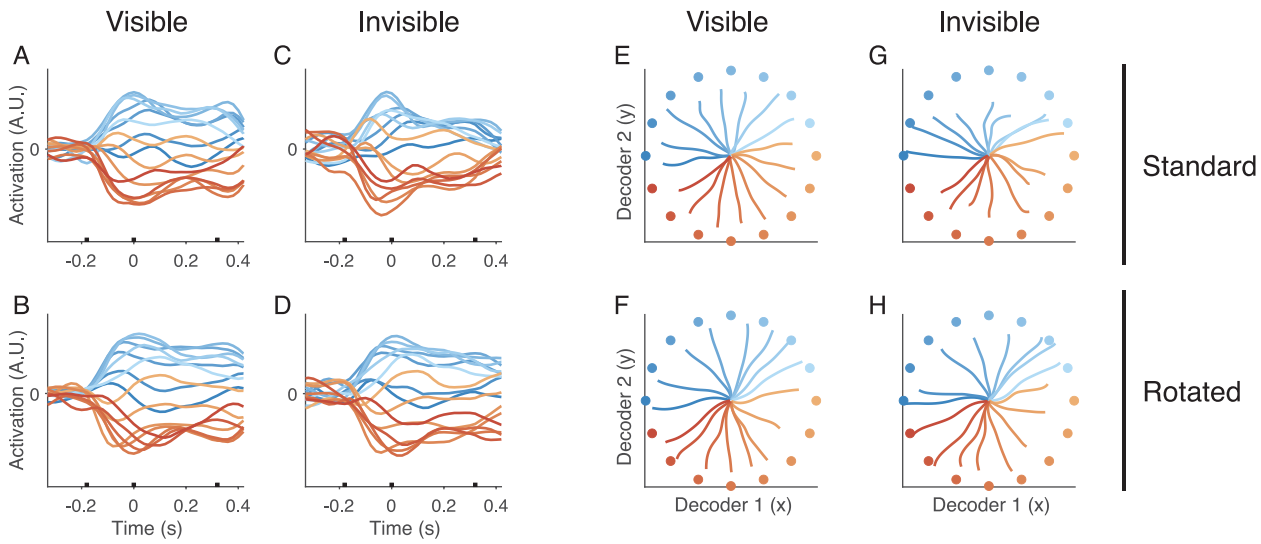

Figure 3.8 Neural readout of physical direction is accurate regardless of cursor visibility.

A-D) Same as in Figure 3.6A-D, but for the decoded *y*-component of physical direction. Decoder output was comparable in each trial type. One exception is for invisible trials in the standard block, where the output amplitude was slightly smaller towards the end of the reach, relative to visible trials. E-H) Same as in Figure 3.6E-H, but for decoded physical reach direction. Trajectories were reasonably accurate in each trial type.

3.3.6 Kinematics and muscle activity were generally comparable across task conditions

If movement kinematics were similar across task conditions, the associated EMG patterns should also be similar. However, state-related effects in M1 might impact EMG patterns directly or indirectly via projections to the spinal cord. Investigators have identified a variety of relationships between M1 activity and muscle activity (Cheney & Fetz 1980; Rathelot & Strick 2009; Kakei, Hoffman, & Strick 1999, Griffin, Hoffman, & Strick 2015; Schieber & Rivlis 2005). We found that the movement kinematics were highly comparable across task conditions. We quantified this using vector field correlation to measure the similarity between the time series of velocity vectors during movement for each pair of conditions (Shadmehr & Mussa-Ivaldi, 1994). Since there were four conditions, there were six possible pairs considered (*n* choose *k*:

$\binom{n}{k} = \binom{4}{2} = \frac{4!}{2!(4-2)!} = 6$). These correlations were very high: across all sessions, pairwise correlations between the six pairs of trial types ranged from 0.96 to 0.99. In general, the patterns of muscle activation in our datasets were also comparable across conditions. Several examples are shown in Supplemental Figure 3.12. As with earlier figures, the top and bottom rows show the EMG recorded during the standard and rotated conditions, respectively. The left and right columns of each subpanel show activity during visible and invisible trials, respectively. From Figure 3.3C, we would expect the EMG data to have relatively similar tuning properties, relative to kinematic direction, before and after adaptation. Indeed, the across-target ordering of activity was comparable in both blocks (Supplemental Figure 3.12, top vs. bottom rows). An exception can be seen in the antagonist burst of the medial deltoid.

Differences in activity between visible and invisible trials were also generally subtle (left and right columns in each panel). One important exception can be seen in the second agonist burst for some muscles. For example, the activity of the triceps around movement offset was weaker during invisible trials than during visible trials (Supplemental Figure 3.12C). This can be seen to a lesser degree in the medial deltoid (Supplemental Figure 3.12A) and wrist flexor (Supplemental Figure 3.12B). One possibility is that the larger EMG activity during visible trials is related to online error-correction or trajectory stabilization. Reaching movements are composed of a primary displacement of the arm followed by one or more corrective sub-movements. These secondary sub-movements are thought to be modulated by "active processing of visual feedback" (Meyer et al 1988), a view supported by our EMG data.

The decrease in late-movement EMG modulation during invisible trials is also noteworthy because we observed weaker modulation in the neuronal population around movement offset for those trials. However, those decreases in modulation were associated with

PD changes in the positive direction following adaptation when reaches were visible, which was not a prominent feature of EMG patterns (Figure 3.3C). Further, the magnitude of modulation change between trial types was not as large for EMG activity. To quantify this, we computed the tuning functions of each muscle during epochs 3 and 4 (see Supplemental Figure 3.9 for epoch timing) and compared their amplitudes in visible and invisible trials (Supplemental Figure 3.13). Although many muscles were less strongly activated when reaches were not visible, this decrease tended to be small: about 14% smaller on average across all observations. When we restricted the analysis to just the fourth epoch, the decrease was slightly larger, about 23% on average. This relatively small decrease may not be surprising given that the physical movement direction was well-represented in firing rates throughout both visible and invisible trials (Figure 3.8). Nonetheless, it is possible that the decreased neuronal modulation we observed around movement offset for invisible trials was in part related to the decreased EMG modulation around the same time.

Previously, our group found that the tuning patterns of single neurons changed sequentially and discretely during normal center-out reach trials (Suway et al 2017). Here, we found that these tuning changes were commonly associated with a transition in encoding properties. For example, Figure 3.2 shows several responses with multiple modulation epochs, and these could be differently impacted following visuomotor adaptation. Figure 3.4 and Supplemental Figure 3.11 show multi-peaked responses that were differently impacted by the visibility of the reach. Although EMG activity patterns were also multi-phasic during our reaching tasks, their patterns did not change as abruptly. To illustrate the nature of neuronal encoding changes, and to contrast them with EMG patterns, we compared activity of each between task types using a sliding correlation analysis (Supplemental Figure 3.14). We found

that abrupt changes in the correlation of neuronal responses between tasks were common, while this was not true for the recorded EMG activity.

3.4 DISCUSSION

Numerous lines of experimentation show that M1 modulation is correlated with surprisingly diverse aspects of motor control. This activity is related to muscle excitability (Fetz & Cheney 1980; Schieber & Rivlis 2005; Griffin, Hoffman, & Strick 2015), force production (Cheney & Fetz 1980; Kalaska et al 1989; Georgopoulos et al 1992), limb kinematics and geometry (Georgopoulos et al 1982; Moran & Schwartz 1999; van Hemmen & Schwartz 2008), visuospatial processing and visuomotor transformation (Zhang et al 1997; Alexander & Crutcher 1990a, Shen & Alexander 1997a), and cognitive processing (Georgopoulos et al 1989; Pellizzer, Sargent & Georgopoulos 1995). Several investigators have examined firing rates in M1 during movements for which the stimulus or visuospatial properties of the task were dissociated from physical movements. For example, Georgopoulos et al (1989) used a "mental rotation" task in which a monkey reached 90° counterclockwise to a stimulus/cue direction. They found that the population vector calculated from M1 firing rates initially pointed in the stimulus direction before rotating toward that of the movement. A subsequent study utilized a "context recall" version of the task in which the monkey relied on a sequence of stimuli to determine the correct movement direction (Pellizzer, Sargent, & Georgopoulos 1995). In this task, M1 neurons first coded for the direction of the stimulus in the sequence that signaled the correct response, and then abruptly changed 100-150 ms later to code the direction of the upcoming movement. A similar type of encoding change was commonly observed during the tasks described here, though

we found these changes occurred at multiple points during a movement (Figures 3.2, 3.4, and Supplemental Figures 3.11 and 3.14).

Other related dissociation paradigms have involved comparing neuronal responses during two types of movements: those made towards a stimulus/cue ("compatible" or "congruent" conditions), and those made in the direction opposite from the stimulus ("incompatible" or "incongruent" conditions). For flexion and extension of the wrist (Zhang et al 1997) or elbow (Alexander & Crutcher 1990a), many M1 neurons initially coded the stimulus direction and later coded the movement direction, a finding comparable to those from the "mental rotation" and "context recall" tasks. A similar trend was observed for M1 activity during the tasks presented here. For example, after monkeys adapted to a visuomotor rotation, the largest percentage of neurons with significantly altered tuning was observed during the reaction time and very early in the movement (Figure 3.5 A,C). This percentage declined about 150 ms later, consistent with the previous studies. In that work, the dissociation between the stimuli and physical movement ended before the movement began, and the movement was essentially unaltered by any of the preceding events. Correspondingly, effects of those events on neuronal firing were not observed throughout the movement. During the rotated block in the tasks presented here, the dissociation between vision and movement was constant, and tuning changes related to this dissociation were common late in the movement, particularly when the cursor motion was visible (Figure 3.5 A,C, blue solid traces). A similar result was found in a four-target joystick-controlled movement task with a similar type of visuomotor rotation (Shen & Alexander 1997a), suggesting this type of encoding in M1 is not related only to movement initiation.

Regardless of the time course of perturbation-related PD changes, we found the magnitude of those changes tended to be larger for larger perturbations (Figure 3.3 A-B). Most

of the observed tuning changes were in the positive direction, corresponding to adaptation to the perturbation. For example, the neuronal response in Figure 3.1 had an upward PD during the standard block, but a leftward PD during the rotated block (a counterclockwise rotation is positive by convention). Cursor movement displayed in VR was upward during the standard block for upward physical movements and upward during the rotated block for leftward physical movement– this property was apparently reflected in the firing of the neuron. Though most of the observed PD changes followed this pattern, those changes were often smaller than the angle of the perturbation. This may be in line with results of previous studies showing "mixed" selectivity for visuospatial and physical movement parameters, rather than "pure" selectivity for one or the other (Lurito, Georgakopoulos, & Georgopoulos 1991; Shen & Alexander 1997a).

Activity in M1 related to visuospatial task features has commonly been interpreted as a hallmark of a "sensorimotor transformation" (Kalaska et al 1997). This process is hypothesized to involve a serial sequence of steps that converts movement-relevant sensory information (i.e. visuospatial stimuli) into an appropriate motor response. In the experiments described here, we used a VR system to dissociate vision and movement rather than simply cueing an "incongruent" response, raising the possibility that the observed tuning changes were related to some adaptation process following the dissociation. It may be that the observed tuning changes represent a mixture of both visuospatial-related coding and adaptation-related effects on firing properties, though our experiments did not dissociate these possibilities. We have presented evidence here that tuning changes late in the movement are related to the visibility of the cursor trajectory. Although much of the empirical evidence for "sensorimotor transformation" in M1 has focused on the neural events around movement initiation, some work has also examined feedback-related firing during or following behavior. In one study, wedge prisms were used to induce visual error

prior to movement, and the firing of many M1 neurons was found to share mutual information with this error following the movement (Inoue, Uchimura, & Kitazawa 2016). Another study introduced a "cursor jump" perturbation mid-movement, in which the random displacement of a feedback cursor necessitated a corrective movement (Stavisky et al 2017). Shortly after that perturbation, putative feedback-related responses in M1 firing rates were found. In the present study, a component of M1 activity, normally observed late in the movement, was weak or absent if the cursor trajectory was not displayed, consistent with a feedback-driven neural response (Figures 3.4 through 3.7). This tuning of this response component was also related to the visuomotor rotation when the cursor trajectory was visible, as would be expected for a signal related to monitoring the VR scene. When responses during visible and invisible trials were compared, this component appeared to constitute a discrete phase or epoch of single unit firing (Figure 3.4, and Supplemental Figures 3.11 B-C and 3.14C).

The observation that many neurons changed their encoding properties discretely during movements is of interest in light of recent efforts to elucidate the "dynamics" of M1 activity during behavior (Churchland et al 2010; Churchland et al 2012; Shenoy, Sahani, & Churchland 2013; Russo et al 2018). These efforts have routinely been implemented using dimensionality-reduction techniques to visualize the temporal evolution of firing rates from a large population of neurons. A common result from these methods is that population activity appears to trace out a smooth trajectory as it evolves in time, leading to the hypothesis that firing rates are subject to a set of dynamical rules that govern the way the pattern at one moment follows lawfully from the previous pattern. We used related methodology here to find dimensions in our data that were useful for decoding the displayed cursor movement (Figure 3.6) or the physical movement (Figure 3.8). During invisible trials, activation along the displayed movement dimensions began

decreasing gradually 100-150 ms after movement onset, reaching near-zero around the time the movement ended (Figure 3.6 C-D). Although this visualization conveyed a relatively gradual evolution toward zero, the underlying changes in single-unit firing were often abrupt (see especially Figures 3.4B, and Supplemental Figures 3.11 B-C and 3.14C). This rapid change found in individual units was likely obscured in population-level analyses because of slight changes in the timing of single-unit responses (Georgopoulos et al 1982; Moran & Schwartz 1999; Cisek & Scott 1999).

The concept of smoothly evolving, lawful neuronal dynamics is appealing for the simple manner in which it captures temporally changing responses, though this simplicity could be associated with a tradeoff in flexibility. Traditional "representational" tuning models of firing tend to generalize poorly over extended time periods, although tuning is stable and robust within discrete epochs. A "temporal parcellation scheme" may constitute a useful compromise between these two descriptions of M1 activity (Johnson, Mason, & Ebner 2001). Much like the "dynamical systems" perspective, this scheme predicts that firing rates covary with task-related parameters in a time-varying sequence (Shenoy, Sahani, & Churchland 2013). However, information encoded in this sequence is likely determined by the particular set of kinematic parameters (Fu et al 1995), visuospatial features (Georgopoulos et al 1989; Pellizzer, Sargent, & Georgopoulos 1995; Zhang et al 1997; Alexander & Crutcher 1990a), reach-to-grasp requirements (Rouse & Schieber 2016), or feedback constraints (Inoue, Uchimura, & Kitazawa 2016; Stavisky et al 2017) of a given behavior, rather than by evolution rules. The common interpretation of M1 as a dynamical system focuses on the locally-driven network state, but excludes the influence of input (Churchland et al 2012; Russo et al 2018). In contrast, the abrupt changes in state signified by changes in tuning are likely input-driven, suggesting this

perspective is incomplete. In principle, neurons need not transition abruptly from the encoding of one parameter to the next, though examples of abrupt patterning in M1 have been reported in several contexts: at the transition from planning a movement to starting the movement (Elsayed et al 2016, Lara et al 2018); upon selecting the correct movement response from a sequence (Pellizzer, Sargent, & Georgopoulos 1995); when monkeys transition between distinct modes of behavior (Abeles et al 1995; Velliste et al 2014); and during single reaching movements (Suway et al 2017; Harpaz et al 2018). Discrete, step-like transitions in firing have also been reported in other brain regions (Latimer et al 2015) and motor systems (Hahnloser, Kozhevnikov, & Fee 2002). This type of patterning has been suggested as a general "neural syntax" for brain operations (Buzsáki 2010).

We found that firing patterns of many single units in M1 reflected the behavioral parameters of each task. Given that M1 is a major source of corticospinal efferents (Porter & Lemon 1993), we might expect that some aspects of these observed changes would also be found in muscle activity. This was generally not the case for tuning changes following adaptation, which were largely dissimilar for firing rates and EMG (Figure 3.3 A-B vs C). Specifically, neuronal tuning changes were most often in the positive direction and could be very large, while tuning changes for EMG activations were almost always in the negative direction and tended to be relatively small. However, for a small subset of neuronal responses, we did observe PD changes following adaptation that could be considered "muscle-like" (Kakei, Hoffman, & Strick 1999). We also observed changes in modulation strength from "visible" to "invisible" trials for both firing rates (Figure 3.5 B,D) and EMG activity (Supplemental Figure 3.13), although those decreases were typically much larger for neurons. Classically, reaching movements have been viewed as composed of two phases: an initial "transport" phase that largely does not depend on

visual feedback, and a target-homing phase that depends heavily on visual feedback (Woodworth 1899). Meyer et al (1988) found that the initial "transport" phase of reaching was invariant in the absence of visual feedback, while this was not true of secondary corrective sub-movements. It is therefore likely that the enhanced EMG activation observed during visible trials was related to visually guided sub-movements. It could be argued that changes in neuronal activation in the absence of feedback were also related to the absence of visually guided sub-movements, and not related directly to vision. However, this hypothesis would require that M1-specified corrective signals during visible trials take place in "visual" coordinates, given that the associated tuning functions changed following adaptation. This property would be at odds with recent findings suggesting feedback-related and correction-related signals in M1 are distinct (Stavisky et al 2017). If these adaptation-sensitive M1 response components are converted to muscle commands by other neural structures, this additional processing step is not likely to be linear: after adaptation, the relationship between the tuning functions of many neuronal responses and muscle activations was found to be nearly 90° out of phase, indicating that they were linearly independent (Figure 3.3 B-C). This observation poses a problem for the common assumption of linear mapping from M1 activation to EMG patterns (Kaufman et al 2014; Russo et al 2018). Indeed, although both types of signals tended to have multiple phases or modulation peaks over a trial, encoding differences of neuronal activity between tasks were large and their relationship often changed abruptly, while EMG activations tended to be similar between tasks (Supplemental Figures 3.12 and 3.14).

These findings are compatible with the general concept that the motor system is driven by latent sources related to behavioral events (Johnson et al 2001). From a dynamical systems view of M1 activity, this could be interpreted as input-generated changes in neural state. Our results

suggest that at least one input is derived from visual feedback associated with the moving cursor. The sequences of discrete encoding patterns imply the system is governed by "attractor states" (Abeles et al 1995) in addition to or instead of smooth rotational dynamics (Churchland et al 2012). Future work could be focused on identifying other drivers and characterizing how they act as input to alter system state.

3.5 SUPPLEMENTAL FIGURES

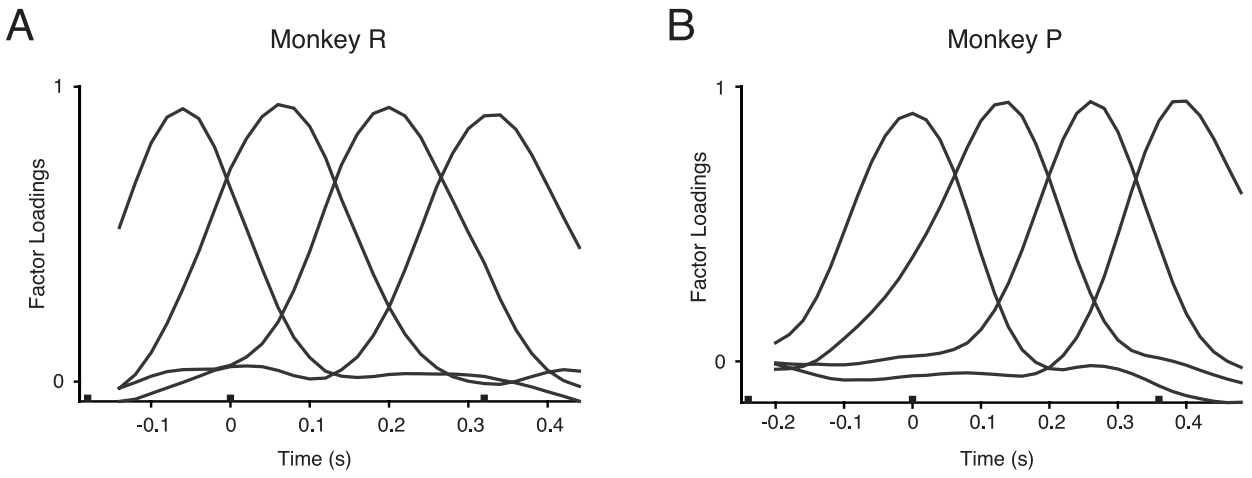

Figure 3.9 Trial segmentation using FA defined four sequential task epochs.

Black Gaussian-shaped traces show the factor loadings over time. The factors correspond roughly to 1. reaction time and movement onset, 2. early movement, 3. late movement, and 4. movement offset. A) Data from monkey R. B) Data from monkey P.

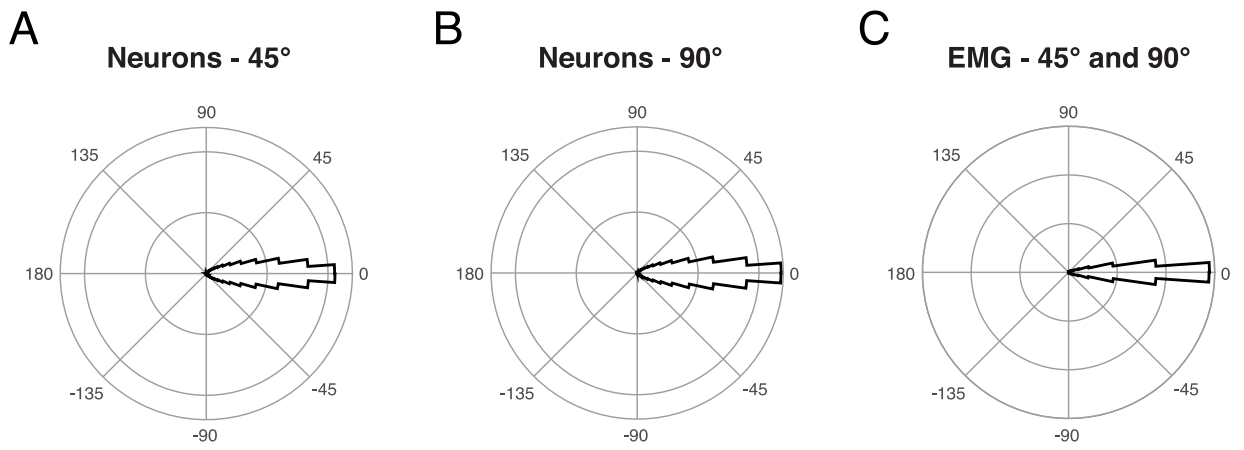

Figure 3.10 PDs of neurons and muscles are consistent during the standard block

Trials during the standard block were partitioned into two groups, and PD changes were computed between the two partitions for each trial epoch. The procedure was repeated 1000 times with new randomly chosen partitions. Changes between partitions were always very close to zero, suggesting that PDs were consistent within the standard block. See Figure 3.3 (main text) for between-block PD changes. A) Neuronal data from 45° sessions. B) Neuronal data from 90° sessions. C) EMG data from 45° and 90° sessions.

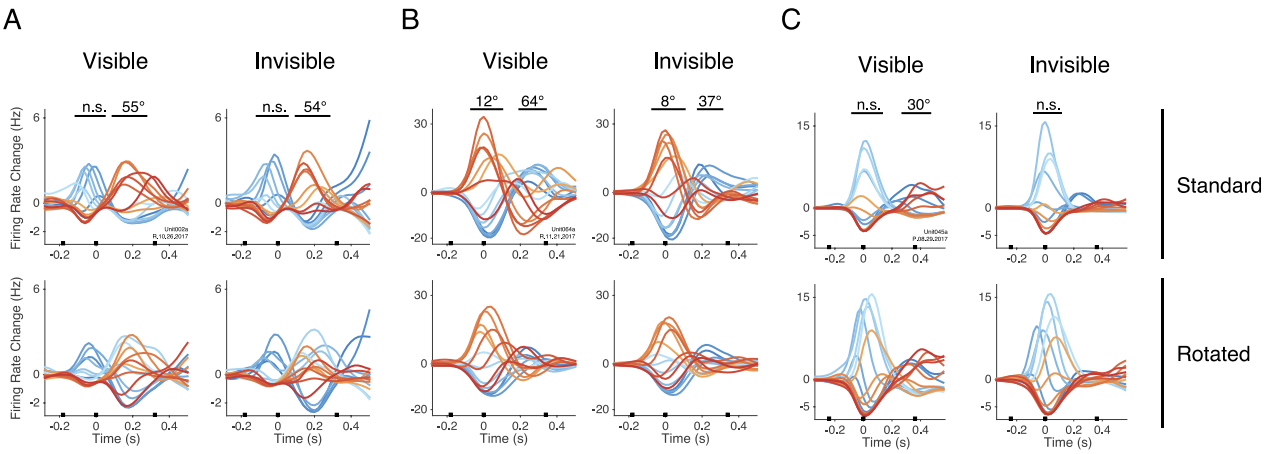

Figure 3.11 Single unit responses had varied relationships to reach visibility.

A) Unit response with two modulation components, and the second component changed PD following adaptation. This component maintained the PD change regardless of reach visibility (left vs. right columns, 55° during visible trials, 54° during invisible trials). B) Another response with two components, and again the second one changed PD following adaptation. However, the magnitude of change depended on reach visibility (64° during visible trials, 37° during invisible trials). C) Response following the same feedback-related pattern as the examples in Figure 3.4 (main text). The examples in panels B and C of this figure also appear in Figure 3.2 (main text).

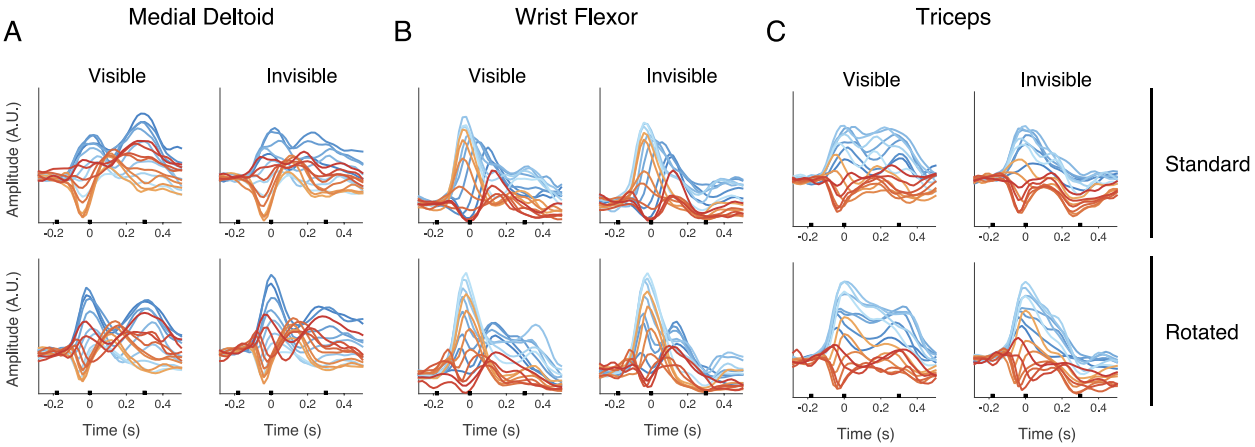

Figure 3.12 EMG activity was similar during each task condition.

Top and bottom rows show activity in standard and rotated blocks, respectively. In each subpanel, left and right columns show activity for visible and invisible trials, respectively. A) Medial deltoid. B) Flexor carpi radialis. C) Triceps. Activity was generally comparable across conditions, though some muscles had a weaker burst of activity late in the movement.

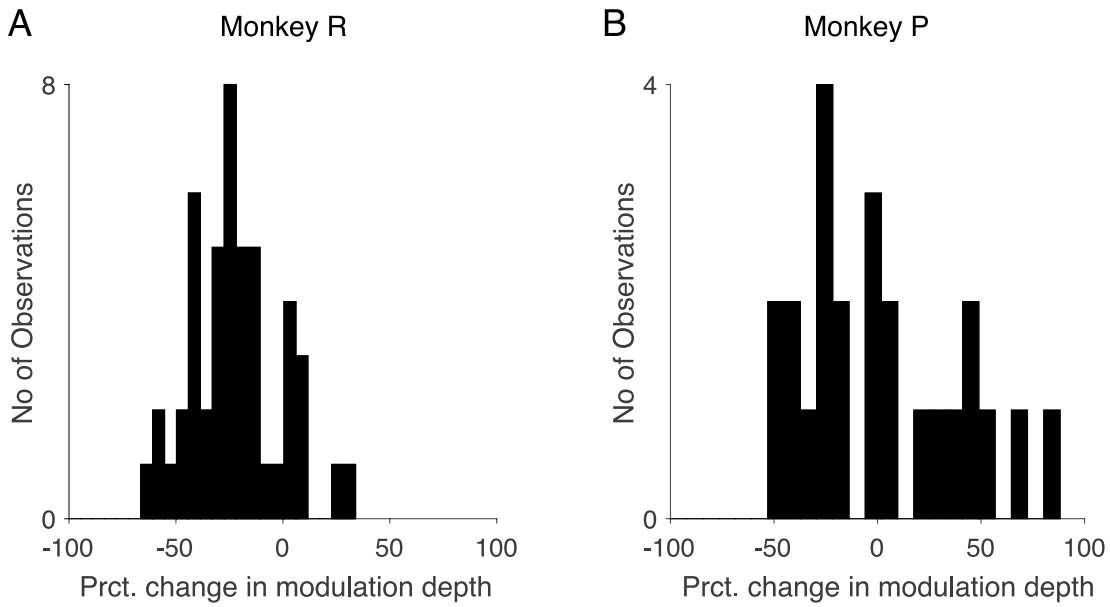

Figure 3.13 EMG modulation depth decrease for invisible trials tended to be small.

Modulation depth was computed in the 3rd and 4th trial epochs (see Supplemental Figure 3.9). We then computed the percent change in modulation depth from visible to invisible trials. This procedure was done for the standard and rotated blocks in each of two epochs; each muscle therefore contributed 4 counts to a histogram. A) Data from monkey R. B) Data from monkey P.

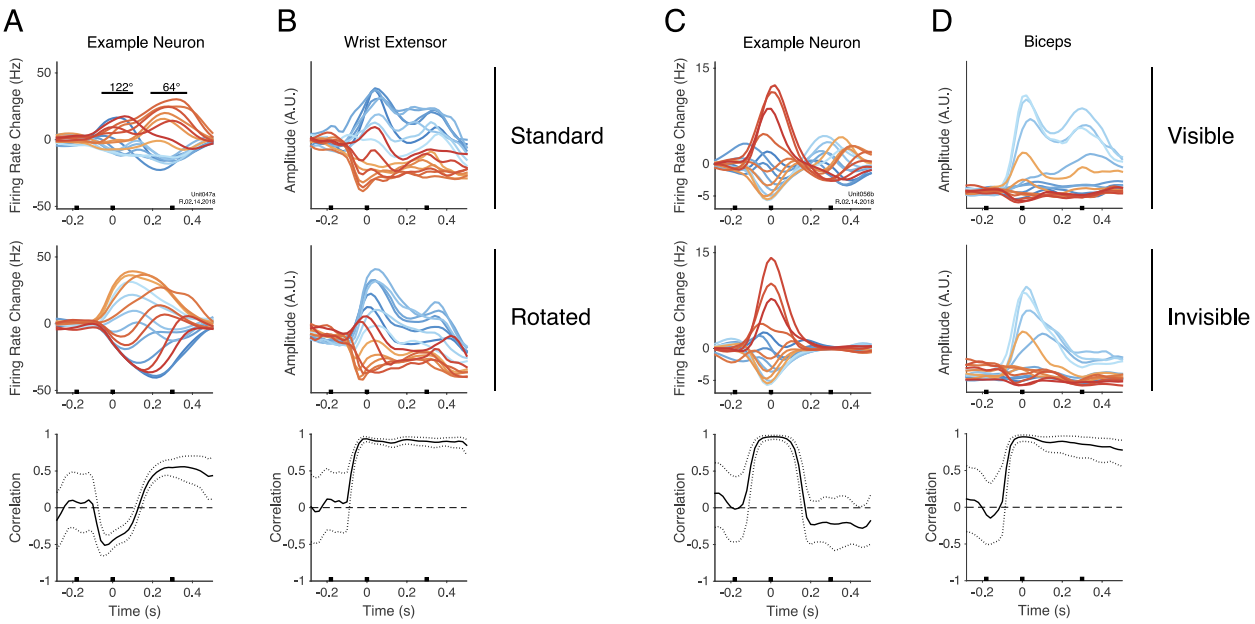

Figure 3.14 Neuronal patterns changed discretely between tasks, while EMG was consistent.

For each time bin, two vectors of rates or EMG were assembled, one for each trial type of interest. Each vector was composed of 16 values, one for each target. The correlation between the two vectors was then computed for each bin. A) Single unit response in the standard block (top row) and rotated block (middle row). Sliding correlation was computed between these responses (bottom row). Solid black trace represents the median correlation over 1000 bootstrap iterations. Dashed traces represent the 95% CI. These responses (which also appeared in Figure 3.2) had an early modulation component with a very large PD change between blocks (122°), and a later component with a smaller change (64°). During the first component, the correlation between blocks was negative. Midway through the reach, however, the amount of PD change was smaller, and the correlation abruptly became positive. B) Similar to A but for the activation of the extensor carpi ulnaris. Across-target activation was highly correlated between blocks throughout the entire trial (bottom row). C) An example unit response that had two modulation components during visible trials (top row), but only one component during invisible trials (middle row). The across-target correlation was very high (about 0.96) around movement onset before dropping in a step-like manner towards zero. D) Similar to C but for activity of the biceps. As with the wrist extensor, biceps activity was highly correlated between conditions throughout the reach, though activation appeared less consistent over trials toward the end of the reach (wider confidence band). The second agonist burst of this muscle was substantially weaker during invisible trials compared to visible trials, but the across-target activation patters remained comparable. This is in contrast to the unit response in panel C.

4.0 ACTIVITY IN PRIMARY MOTOR CORTEX DURING REACHES WITH DIFFERENT TARGET ACQUISITION REQUIREMENTS

4.1 INTRODUCTION

Classical perspectives on limb control posit that reaches consist of multiple behavioral components (Woodworth 1899; Keele & Posner 1968; Soechting & Lacquaniti 1981; Meyer et al 1988; Novak, Miller, & Houk 2002). First, there is a relatively ballistic "initial transport" phase that brings the limb toward the general location of a target. Subsequently, there is a "target homing" phase involving one or more corrective sub-movements made with the aid of visual feedback. In Chapter 2, we reported that tuning of single M1 units changed discretely during movement, and we suggested these changes were related to transitions between the phases of reaching. In support of this hypothesis, we found that unit activity could be weaker toward the end of movement if monkeys were not required to stop accurately, indicating late-movement encoding was related to halting the arm or acquiring the target. In Chapter 3, we showed that late-movement encoding could also be related to visual feedback, further linking M1 activity to classical notions of reach structure.

Here, we repeated and expanded on our previous work by recording M1 activity while monkeys made center-out reaches that differed in the way they were terminated. For one reach type, movements were accurately stopped and the arm was held within a target zone, as is typical

for center-out reaching tasks. For a second reach type, monkeys only needed to touch the target and were rewarded immediately. This design imposed a clear difference in the control requirements of target acquisition. We found that single unit encoding was initially comparable between reach types but was distinct later in the movement depending on the need to stop and hold. When there was no hold period, we often observed weaker modulation late in the movement. However, responses were diverse and it was not uncommon to observe late-movement encoding unique to either reach type. This pattern suggests M1 engages in separate processes depending on the way targets are acquired.

Previous work has shown that EMG activity can be markedly different when movements are not terminated accurately (Waters & Strick 1981). Although the stopping requirements of reaches differed in our tasks, we found EMG activity recorded from several arm muscles during the two types was surprisingly comparable, even late in the movement. We therefore wondered what factors might drive the encoding differences seen for single units. Since visual feedback is critical for stopping accurately, we hypothesized that late-movement encoding would covary with visuospatial information more strongly for accurately stopped reaches. To examine this possibility, we dissociated vision from movement using the same visuomotor rotation described in Chapter 3 and identified adaptation-related responses as previously described. Contrary to our hypothesis, late-movement encoding related to the adaptation was somewhat more common when reaches were *not* stopped accurately. Our findings suggest that processes related to target acquisition drive M1 activity differently depending on their specific control requirements.

4.2 MATERIALS AND METHODS

Details of the experimental setup, basic task parameters, neuronal and EMG data collection and pre-processing have been described in Chapter 3.

4.2.1 Behavioral task

Monkeys performed center-out reaching movements in the same virtual reality (VR) environment as described previously. Typically, center-out reaching tasks require a subject to accurately terminate movements inside a target zone and hold steady for a short period of time. Here, we sought to study M1 encoding related to this control requirement. We therefore included two trial types in each session: 1) a "hold" trial type in which reaches were stopped accurately and held within a target zone for 600-800 ms, and 2) a "no hold" trial type in which a reward was delivered upon touching the target with the cursor. The two trial types were randomly interleaved, and monkeys were trained to distinguish them based on the color of the home position and peripheral target.

As in Chapter 3, we leveraged the VR paradigm to impose a visuomotor rotation during behavior. Each session was split into two blocks ("standard" and "rotated"), and the two reach types ("hold" and "no hold") were performed in each. We analyzed five neurophysiological and behavioral datasets from each monkey. For monkey P, we analyzed EMG data recorded over two additional sessions. For monkey R, two of the five datasets included simultaneously recorded EMG activity. Sessions from monkey R included 143 neurons per day on average (SD = 43.6 neurons, 715 total observed neurons). Sessions from monkey P included 31.8 neurons on average (SD = 11.6 neurons, 159 total observed neurons).

4.2.2 Kinematic tuning analysis

For some analyses, we applied the standard cosine tuning model to neuronal firing rates as described in Chapter 3. For reference, the same tuning model is defined here:

(4.1) $y = k + B \bullet \cos(\theta - \theta_{PD}) + \varepsilon$

where y is the estimate of a neuron's firing rate, k is the baseline rate, B is the amplitude of the tuning function (modulation depth), θ is the physical movement direction, θ_{PD} is the preferred direction (PD), and ε is the noise or fitting error. Although cosine tuning is a robust descriptor of M1 activity during reaching, single unit firing is also known to covary with movement velocity and speed (Moran & Schwartz 1999; Inoue et al 2018). The simpler cosine tuning model can be extended to account for these kinematic parameters. To show this relationship, we first give an alternative expression for the cosine tuning model in Equation 4.1:

(4.2) $y = k + B_x D_x + B_y D_y + \varepsilon$

where y, k, and ε are as defined previously, B_x and B_y are regression coefficients, and D_x and D_y are the x and y components of movement direction, respectively. The PD of the cosine fit (θ_{PD} in Eq. 4.1) is then $\arctan(B_y / B_x)$, and the modulation depth (B in Eq. 4.1) is $\sqrt{B_x^2 + B_y^2}$. Movement direction is defined as a vector with unit length, whereas velocity is defined as the direction vector scaled by movement speed:

$V_x = SD_x$

$V_y = SD_y$

and,

$S = \sqrt{V_x^2 + V_y^2}$

Where V_x and V_y are the x and y components of velocity, and S is scalar speed. M1 firing rates covary with both vectorial velocity and scalar speed (Moran & Schwartz 1999), as follows:

(4.3) $y = k + B_x V_x + B_y V_y + B_s S + \varepsilon$

where B_s is a regression coefficient and all other terms have been defined previously. Equations 4.2 and 4.3 are thus related: speed acts as both a gain and an offset to the standard cosine tuning function. Equations 4.2 and 4.3 will be referred to as the "direction-only" model and the "velocity-speed" model, respectively. Because speed is a non-directional value, we did not subtract the across-target firing rate mean from unit responses prior to fitting kinematic models as was done elsewhere.

In some analyses, each kinematic model was fit to firing rates and a choice was made as to the "better" fit. The velocity-speed model is more complex than the direction-only model, so a concern is that greater explanatory value could simply be attributed to greater model flexibility. To address this concern, we utilized the Akaike Information Criterion (AIC) to select a model. The AIC reflects a balance between a model's explanatory value and its complexity, such that a more complex model would be preferred if the gains in explanatory value justified the additional degrees of freedom. We defined the "preferred" model as the one that minimized the AIC.

4.2.3 Vector field correlation

To compare sets of scalar quantities, a typical approach is to compute correlation or covariance values. In several analyses, we wished to assess the similarity of two sets of vectors, rather than scalars. For this purpose, we relied on a "vector field correlation" method described by Shadmehr & Mussa-Ivaldi (1994). The following notation is modified from the original citation. Consider two sets of vectors, arranged into matrices:

$A \in \mathbb{R}^{NxT}$

$B \in \mathbb{R}^{NxT}$

where N is the number of elements in each of the T vectors. For example, if A and B are data matrices representing time series (sets) of velocity vectors, N would represent the number of spatial dimensions (two or three for movement in a plane or volume) and T would represent the number of time points in each set. We seek an expression for the correlation between these sets:

$$(4.4) \quad \rho(A,B) = \frac{Cov(A,B)}{\sigma(A)\sigma(B)}$$

Note that this expression differs from standard correlation in that A and B are data matrices. The operation denoted by *Cov* is nonetheless meant to indicate computation of scalar covariance, rather than a covariance matrix. To yield this scalar, we defined the operation as follows:

$$(4.5) \quad Cov(A,B) = \frac{1}{T}\sum_{j=1}^{T} a_j \cdot b_j$$

where a_j and b_j are the mean-subtracted j^{th} columns of A and B, respectively, and the symbol $\cdot$ is the dot product operator. The σ operator was defined as:

$$(4.6) \quad \sigma(A) = \sqrt{\frac{1}{T}\sum_{j=1}^{T} a_j \cdot a_j}$$

For some analyses, we relied on the expression in Eq. 4.5 as a measure of both the similarity and magnitude of A and B ("vector field covariance"). For other analyses, we used the scaled version of this metric described by Eq. 4.4 to consider only the similarity between sets without regard to magnitude ("vector field correlation").

4.3 RESULTS

4.3.1 Movements were similar across task conditions

Monkeys made straight and accurate reaches to each of the 16 targets (Figure 4.1). Aside from slight differences in curvature or extent of movement, these trajectories were highly comparable in "hold" and "no hold" trials, and before and after the visuomotor adaptation. We quantified similarity using vector field correlation, computed for the trial-averaged velocity vectors during movement for each pair of conditions (Shadmehr & Mussa-Ivaldi, 1994, see methods). Since there were four conditions, there were six possible pairs of conditions (*n* choose *k*: $\binom{n}{k} = \binom{4}{2} = \frac{4!}{2!(4-2)!} = 6$). Correlations between each pair were very high: over all datasets and condition-pairs for both monkeys, values ranged from 0.95 to 0.99. Although the velocity measurements were well-correlated, a notable difference between certain conditions was the peak velocity, which tended to be higher during "no hold" trials. On average over all trials in our datasets, reaches were 33.56% and 22.14% faster for monkeys R and P respectively during "no hold" trials, relative to "hold" trials. This difference is likely due to a tradeoff between speed and accuracy depending on the stop-and-hold requirement. Note that although peak velocity was different between these two reach types, data from each trial were aligned based on the kinematic time-course (see Section 3.2.4) to ensure comparability across conditions.

We next compared EMG activity across each task condition. Average activity of several example muscles is shown in Figure 4.2. Although activity was generally comparable between task types, there were also some notable differences. One difference can be seen in the initial burst of activity around the onset of movement, which was often stronger during "no hold"

reaches due to the faster movement velocity for those trials. We also found that the second phase of the typical tri-phasic EMG pattern (i.e. the antagonist burst) was sometimes stronger for "hold" trials (e.g. Figure 4.2 A-B, left vs. right columns). This increase could be related to the need to more accurately "brake" the arm's movement in those trials (Waters & Strick 1981). Another difference can be seen for the third phase of the tri-phasic EMG pattern, occurring around the offset of movement, and which was often more tonically sustained during "hold" trials than "no hold" trials. A clear example of this pattern can be seen in Figure 4.2B (left vs. right columns). This pattern is not surprising given the need to keep the arm stationary and outstretched at the end of "hold" trials. Despite these differences, activity patterns of the muscles in our datasets were well-correlated across task types. We computed correlations between the trial-averaged EMG activity of each muscle, for each of the six pairs of conditions. Correlations computed for activity during movement were high: the median correlation values over muscles and conditions were 0.86 and 0.92 for monkeys R and P, respectively, indicating a reasonably high degree of similarity between task types.

These summary statistics were computed over all possible task pairs, but we might expect greater differences in EMG activity between "hold" and "no hold" trials than between "standard" and "rotated" trials. However, when we restricted the analysis only to "hold" and "no hold" comparisons, the median correlation did not change for monkey R and decreased only slightly for monkey P (0.90, down from 0.92). Because these reach types might be expected to be most different toward the end of movement, we wondered if EMG activity was less consistent across tasks specifically during that period. We therefore repeated the correlation analysis using activity from the time of peak velocity until movement offset. Over this period, the median correlation increased slightly to 0.88 for monkey R and decreased slightly for monkey P to 0.89.

As an additional comparison, we also considered a 100 ms window centered on movement offset, and found the median correlations were 0.90 and 0.89 for monkeys R and P, respectively. Although these correlation measures do not capture the moderate task-related differences in EMG amplitude we have described, our data broadly suggest the muscles in our datasets had comparable activity patterns during the terminal phase of both reach types.

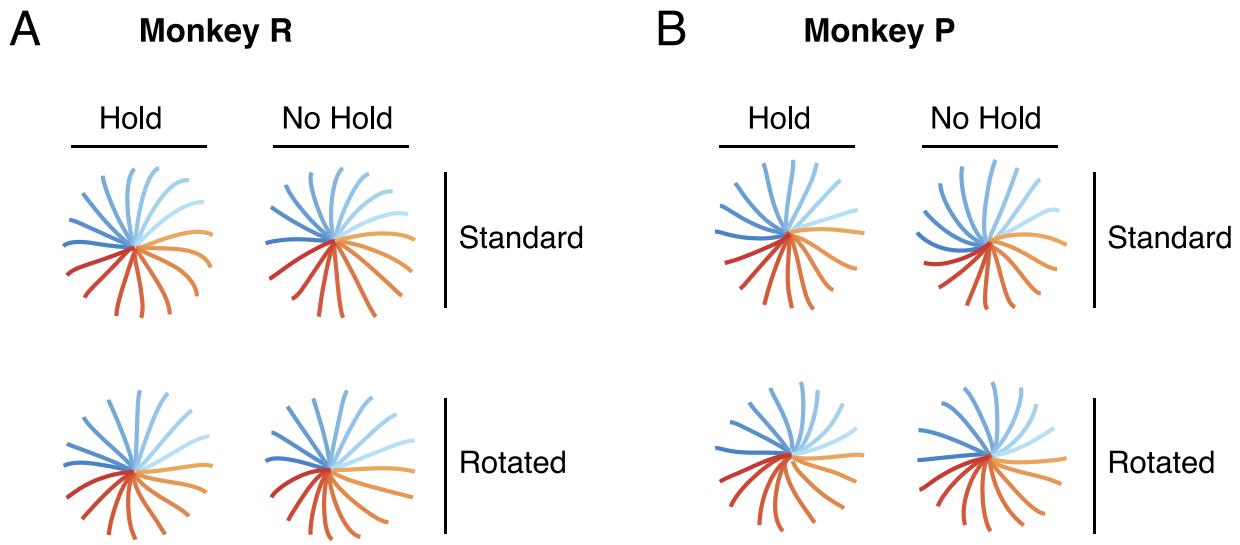

Figure 4.1 Average reach trajectories during each task type

Average reach trajectories are shown for each of the 16 targets (different colored traces) in the "standard" and "rotated" blocks (top and bottom rows), and "hold" and "no hold" trials (left and right columns). Trajectories start at movement onset and stop at movement offset. A) Data from one monkey R session with a 45° rotation. B) Data from one monkey P session with a 90° rotation.

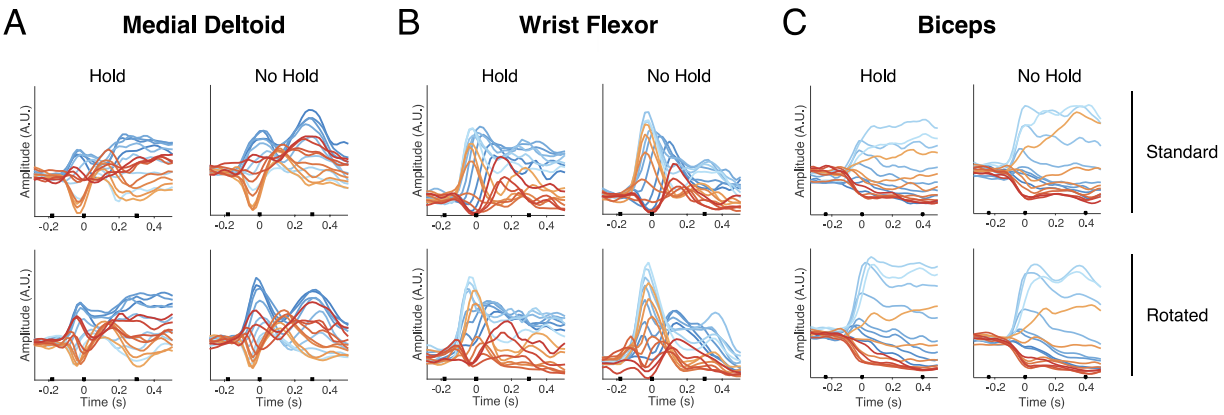

Figure 4.2 Example EMG activity during each task type

Trial-averaged EMG activity for each reach direction (different colored traces) in each task type. Top and bottom rows show the "standard" and "rotated" blocks, left and right columns in each subpanel show "hold" and "no hold" trials. The black square markers along the abscissa indicate the time of target show, movement onset, and movement offset. A) Medial deltoid activity recorded from monkey R. B) Wrist flexor activity recorded from monkey R. C) Biceps activity recorded from monkey P. EMG activity was generally comparable between task types, though some differences were apparent (see main text).

4.3.2 Firing rate patterns could be very different during "hold" and "no hold" reaches

Although movement kinematics and EMG activity were similar across task conditions, neuronal firing patterns during each could vary drastically. These differences were often most obvious when comparing "hold" and "no hold" trials. Previously, we reported that a firing component late in the movement was less common when monkeys did not need to stop and hold in the target zone (see Chapter 2). Several examples of this patterning difference are illustrated in Figure 4.3 A-C. As seen for those examples, late-trial firing often appeared to constitute a discrete phase or epoch that depended on reach type, while earlier activity was very similar regardless of the hold requirement. This discreteness contrasts with the differences in EMG activity observed between the two reach types (e.g. Figure 4.2B, left vs. right columns). The firing pattern difference is not likely related to the tonic activity sometimes seen during hold periods (Kettner, Schwartz, &

Georgopoulos 1988; Wang et al 2007), since encoding patterns diverged much earlier in the movement (for comparison see Chapter 2, Figure 2.4 third and fourth columns, and Supplemental Figure 2.13). Figure 4.3D shows the response of another unit with a late-movement firing component only during "hold" trials, although interestingly this unit was not modulated at all during the other reach type. These data, along with those presented in Chapter 2, suggest that a component of M1 activity is associated with the need to stop accurately in a target zone. This component appeared to constitute a discrete response feature that was absent during "no hold" trials.

Interestingly, the opposite pattern was also sometimes observed: a response component of some units was specific to "no hold" trials. This pattern can be seen for the examples illustrated in Figure 4.4 A-B. Like the examples in Figure 4.3 A-C, the early response components of these neurons did not vary much between reach types, while the later components were reach-specific. Another interesting pattern is illustrated in Figure 4.4C, again with a firing component late in the movement only during "no hold" trials. For both reach types, this unit was modulated around movement onset, though the directional tuning at that time changed slightly from down and left during "hold" trials (top row, red and dark orange traces) to down and right during "no hold" trials (bottom row, light orange and light blue traces). The unit response shown in Figure 4.4D was strongly modulated around both movement onset and offset during "no hold" trials, and only weakly during the entire time span of "hold" trials. These data suggest that a component of M1 activity is associated with a terminal reach phase that lacks an accurate stopping event.

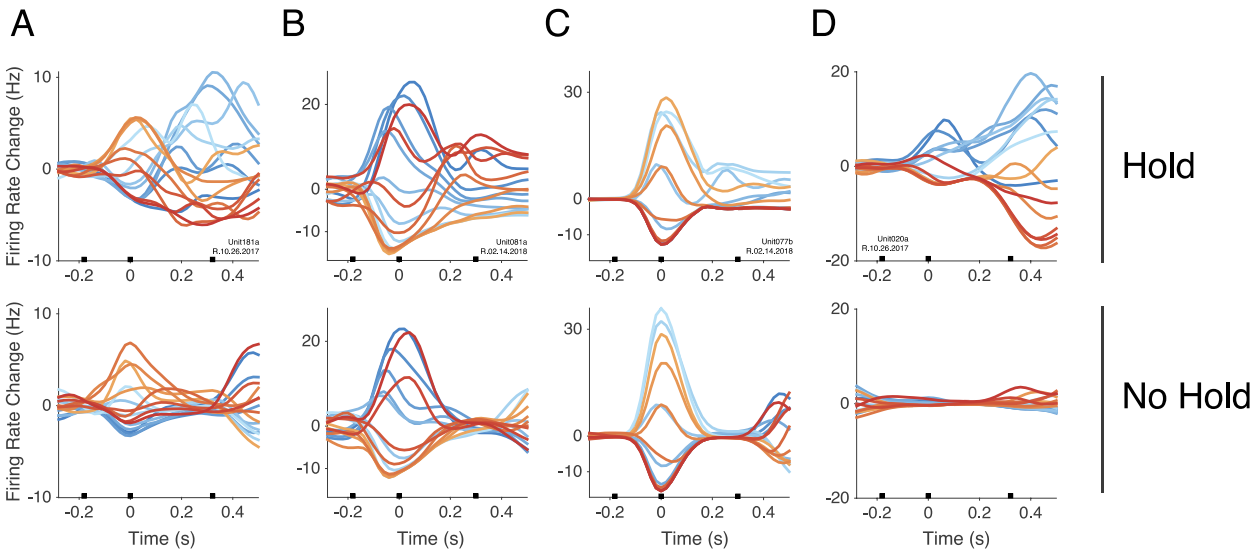

Figure 4.3 Example neuronal responses with a firing component specific to "hold" trials

Average firing rates recorded from example single units (columns) during “hold” and “no hold” trials (top and bottom rows). Color scheme matches previous figures. A-C) Unit responses that had two firing components during “hold” trials but only one during “no hold” trials. D) Unit response that had two firing components during “hold” trials but was not strongly modulated during “no hold” trials.

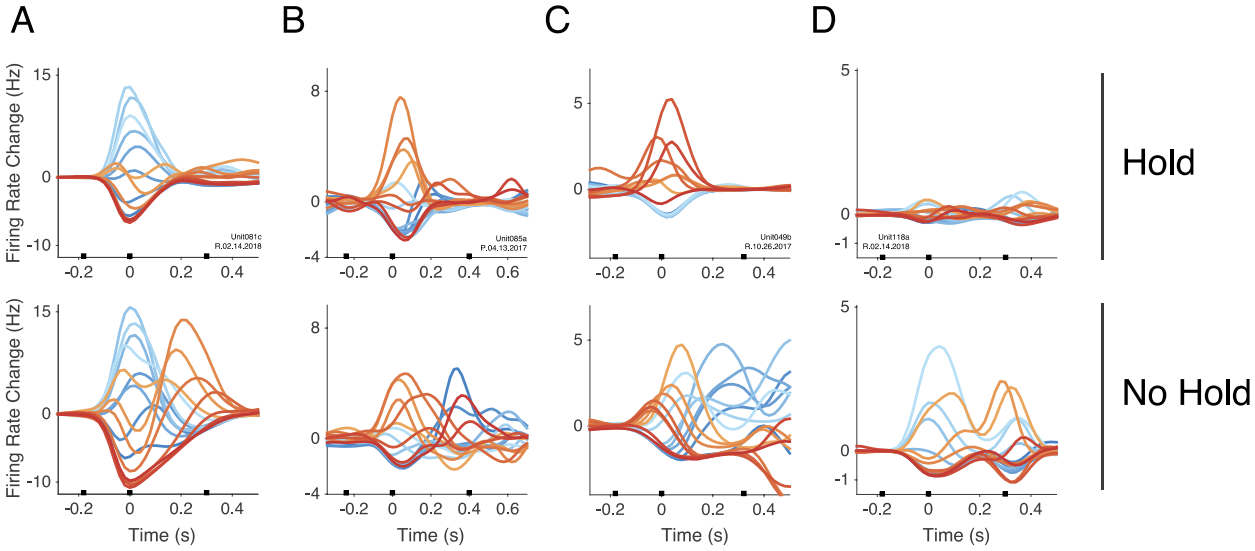

Figure 4.4 Example neuronal responses with a firing component specific to "no hold" trials

Additional example single unit responses (columns) during “hold” and “no hold” trials (top and bottom rows). A-C) Unit responses that had one firing component during “hold” trials and two during “no hold” trials. D) Unit response that was not strongly modulated during “hold” trials and had two firing components during “no hold” trials.

4.3.3 Correlation of encoding patterns between reach types is epoch-specific

In the previous section, we found that firing patterns were often similar between reach types early in the movement, before abruptly diverging mid-reach based on the hold requirement. To visualize and quantify this pattern, we applied a sliding correlation analysis to the rates of single units. For each time bin, we assembled two vectors of rates, one for each reach type. Each vector was composed of 16 firing rate values, one for each target. We then computed the correlation between the two vectors. In other words, we correlated the units' tuning at each time bin. As expected from casual inspection of response patterns, we commonly observed that activity was initially highly correlated between reach types, but poorly correlated late in the movement (Figure 4.5, bottom row). The change from high to low correlation often occurred abruptly in a step-like manner, in line with the hypothesis that there is a relatively discrete encoding phase that varies depending on the stopping requirement.

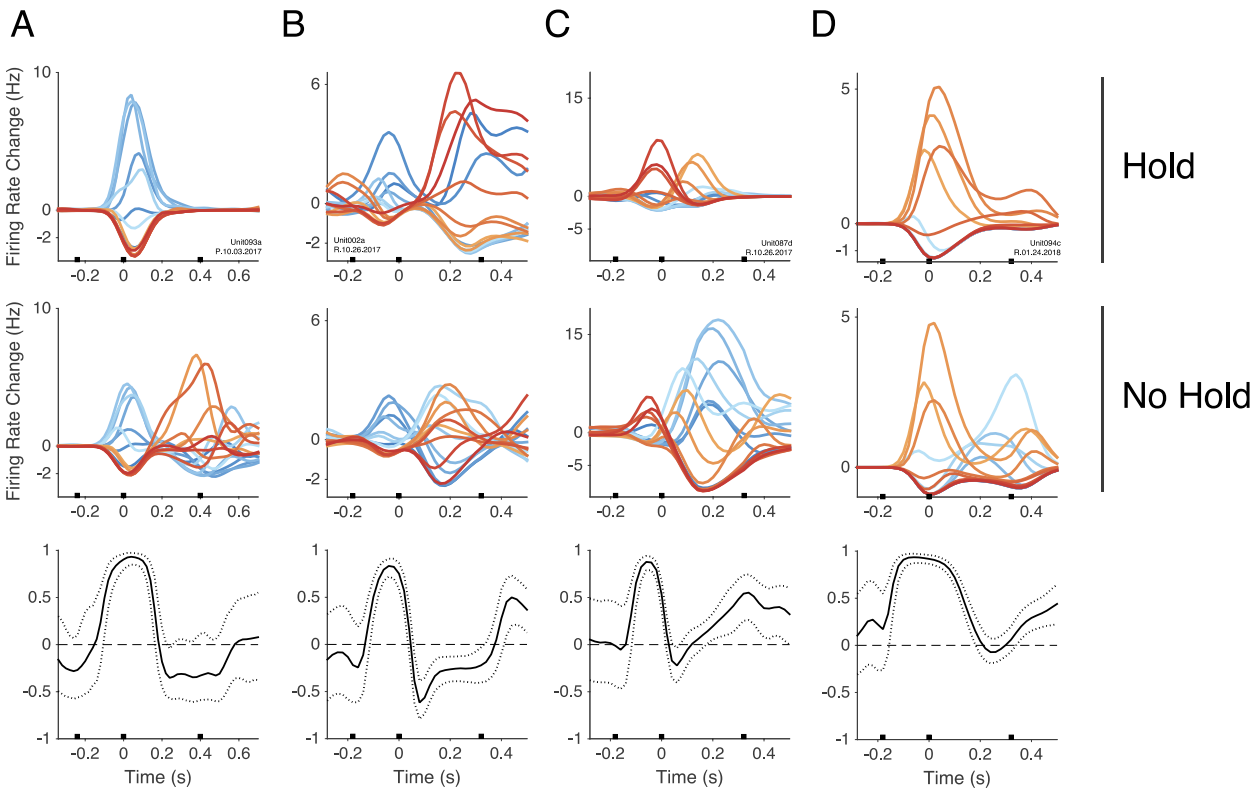

Figure 4.5 Sliding correlation analysis of encoding patterns between reach types

Additional single unit responses (columns) recorded during "hold" and "no hold" trials (top and middle rows). Bottom row shows a sliding correlation analysis in which the pattern of rates across targets was correlated between reach types, for each time bin. Trials were resampled with replacement 1000 times to assess variability. The median (solid black traces) and 95% confidence interval (dashed traces), computed over the 1000 bootstrap repetitions, are shown in each subpanel in the bottom row. For these example responses, across-target firing was highly correlated early in the trial and abruptly became poorly correlated later in the trial.

To summarize the time-varying correlation patterns for all units, we first segmented trials into a few epochs using factor analysis (FA, see Methods), and correlations were computed within these epochs. As in Chapter 3, we chose to identify four factors from our data, representing prominent firing patterns over the four epochs (Figure 4.6). The timing of factors was similar to that found previously (see Chapter 3). Once the epochs were defined using FA, we considered firing rates averaged over 100 ms windows centered on each factor's peak. Within each of the four windows, we correlated the 16-element vectors (one value per reach target) of each unit's firing rates in "hold" and "no hold" trials (Figure 4.7). Because unit responses are not always consistently modulated during behavior, a concern is that low correlation values might not reflect a genuine encoding change but may be attributed to noise. We therefore defined an inclusion criterion for each comparison. For a given epoch, the response of a unit during at least one of the reach types was required to be sufficiently consistent over trials. Consistency was assessed by dividing trials within a single reach type ("hold" or "no hold") into two equal parts and computing the trial-averaged tuning correlations between the two partitions. This procedure was repeated 1000 times with new randomly chosen partitions. Responses with a median correlation over these repetitions of less than 0.7 for both reach types were excluded from the analysis.

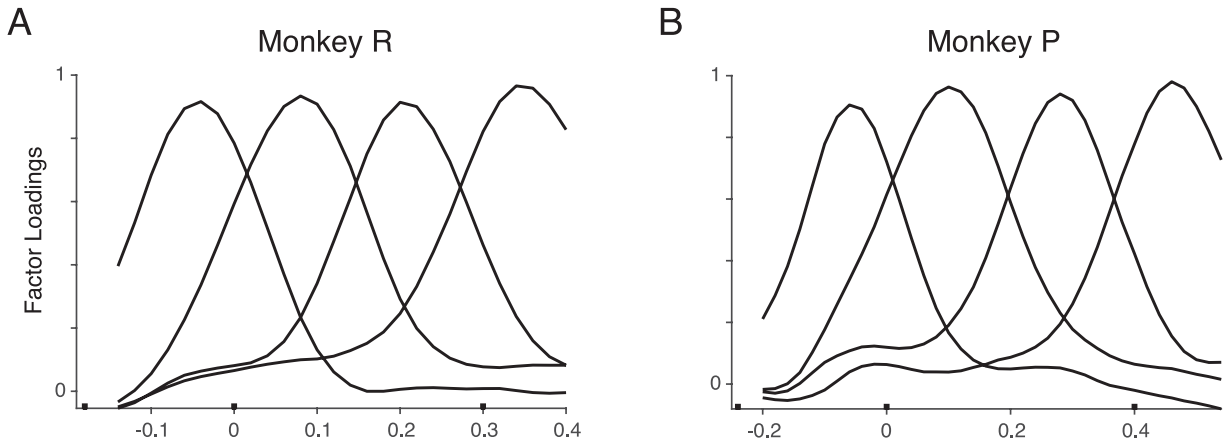

Figure 4.6 Trial epochs defined by factor analysis

Trial segmentation using FA. Black Gaussian-shaped traces show the four factor loadings over time. The factors correspond roughly to 1. reaction time and movement onset, 2. early movement, 3. late movement, and 4. movement offset. A) Factor loadings computed using data from monkey R. B) Factor loadings computed using data from monkey P.

We found that the pattern illustrated for single units in Figure 4.5 was broadly applicable to most neuronal responses: in the first two epochs, correlations between reaches were high for most units. In the later two epochs, these correlations could be much lower (Figure 4.7). Responses were most often similar between reach types during the earliest epoch, with correlation values greater than 0.7 for 86% and 89% of units (monkeys R and P, respectively). Many units maintained this high degree of correlation during the second epoch, although fewer for monkey R (65%) than for monkey P (87%). The percentage of units surpassing the 0.7 correlation threshold declined about 15-25% in each remaining epoch for both monkeys. Given the stringent inclusion criterion for this analysis, these low correlations cannot be attributed to noisy responses. In summary, activity in M1 appears to be similar during roughly the first half of both reach types but diverges during the second half depending on the hold requirement.

As a final demonstration of the prominence of this pattern, we used principal component analysis (PCA) to find linear combinations of the neuronal data that captured a large amount of population firing rate variance during a single reach type. We reasoned that, if a linear combination explained much of the variance for one reach type, it would do so for the other type as well, but only during the early part of the movement when responses were similar. We therefore computed principal components (PCs) using data from "hold" trials only, and projected data from both reach types onto the top PC (Supplemental Figure 4.12A). As expected, activation along this PC was consistently high during movement for "hold" trials (left column), since the PC was chosen to maximize variance for that reach type. Activation along the same PC for "no hold" trials was high during the early part of movement, but declined toward near-zero values starting about halfway through the reach, mirroring results of the epoch-based correlations (Figure 4.7). The same analysis was next carried out using EMG rather than neuronal data (Supplemental Figure 4.12B). Activation along the top EMG PC (computed using data only from "hold" trials) was similar for both reach types, in line with the correlation results reported earlier (Section 4.3.1).

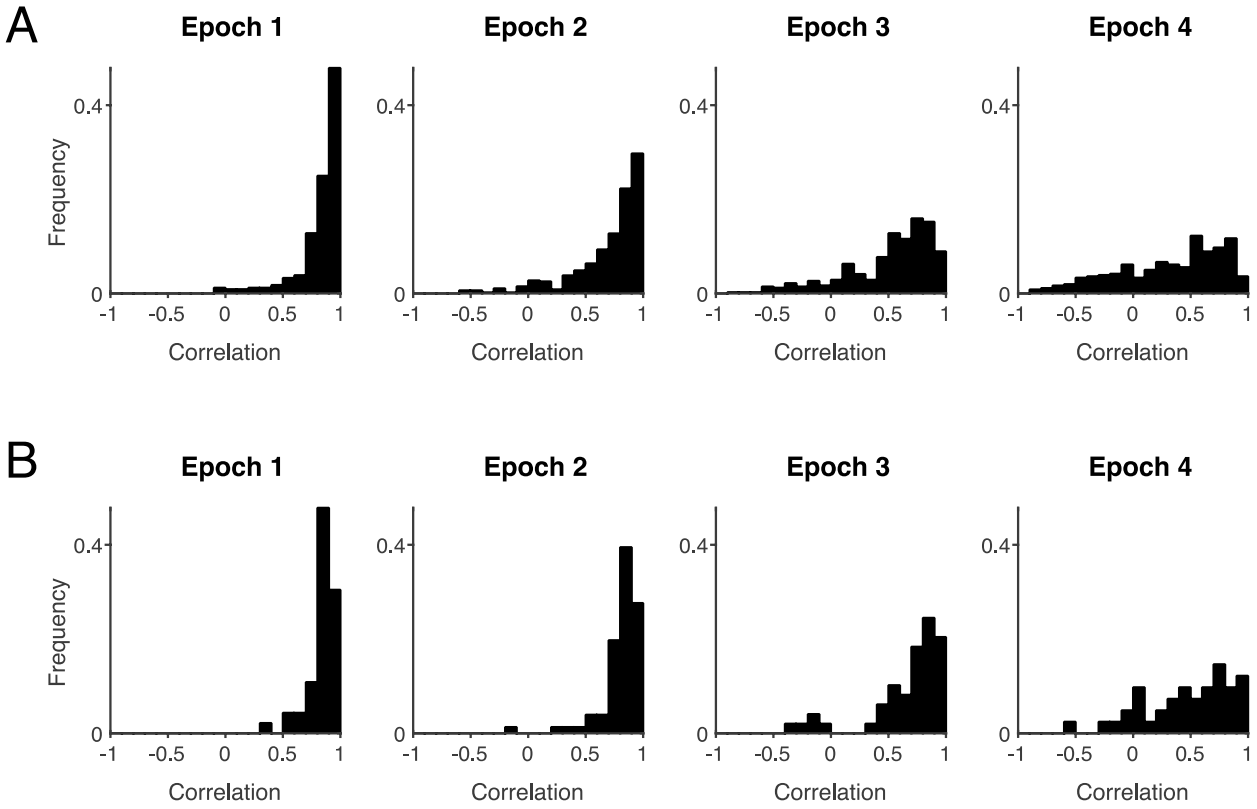

Figure 4.7 Tuning correlations between reach types during four epochs

Across-target firing rates from each unit were correlated between "hold" and "no hold" trials using rates averaged over 100 ms windows centered on each factor peak (see Figure 4.6). Histograms show correlation values during each epoch (columns) for all units recorded from A) monkey R and B) monkey P. Correlations of responses between reach types were high for most units during early epochs and lower during later epochs.

4.3.4 Reach-dependent neuronal encoding is not trivial

Might the response features unique to "no hold" trials (e.g. Figure 4.4) be trivially related to the upcoming return movement, which was delayed for "hold" trials? In that case, during "no hold" trials, the ordering of firing rates across targets early in the movement should be reversed from that seen late in the movement (Georgopoulos, Kalaska, & Caminiti 1985). Based on the single unit examples in Figures 4.4 (bottom row) and 4.5 (middle row), there seems to be a wide range of changes in firing pattern between the early and later phases of unit responses. This diversity

was quantified by measuring the preferred direction (PD) of each unit during an early and late epoch of “no hold” trials and comparing the two measurements. If late-movement PDs were related to returning to the home position, we would expect PD differences to cluster around 180°. Because the timing of modulation phases could vary from unit to unit, we repeated the comparison for average rates during each non-adjacent pair of FA epochs (first vs. third or fourth, and second vs. fourth). Angular PD changes between epochs were found to be broadly distributed around the unit circle for the population, with a slight bias toward smaller changes: for all comparisons, a majority of responses had PD changes smaller than 90° (Table 4.1). This pattern is not consistent with the hypothesis that late-movement encoding during “no hold” trials was simply related to the upcoming return movement.

	Monkey R	Monkey P
First vs. Third	71%	78%
First vs. Fourth	53%	52%
Second vs. Fourth	61%	55%

Table 4.1 Percent of units with less than 90° PD change between epochs of "no hold" trials

PDs of single unit responses were compared during non-adjacent trial epochs. PD changes between epochs were broadly distributed around the unit circle, with a slight bias toward smaller changes (less than 90°). Results pooled over datasets for each monkey.

4.3.5 Velocity encoding is prevalent early in the movement

What is the consistent M1 driver during the early part of “hold” and “no hold” trials? There is considerable evidence that M1 activity correlates strongly with reach velocity starting shortly before movement onset (Schwartz 1992, 1993; Moran & Schwartz 1999; Wang et al 2007; Inoue

et al 2018), so perhaps this kinematic signal drives unit activity in both reach types. The datasets used in the present study were well suited to testing this hypothesis because of the consistent difference in velocity between "hold" and "no hold" trials. To assess possible velocity encoding, we fit two different kinematic tuning models to single-unit rates, combining data from both reach types. The first model was the standard directional tuning model (Equation 4.2), while the second included velocity rather than direction, plus a scalar term for speed (Equation 4.3). For each unit's response, we evaluated each model and compared goodness-of-fit. We reasoned that the more complex velocity-speed model would be preferable to the direction-only model if the temporal profile and magnitude of firing rates and velocity measurements covaried. Model preference was determined by comparing the Akaike Information Criterion (AIC) of the two model fits. The AIC reflects a balance between a model's explanatory value and its complexity, such that a more complex model would be preferred if the gains in explanatory value justified the additional degrees of freedom.

Previous studies have shown that M1 activity best correlates with velocity when the movement is shifted back in time ~150 ms (Moran & Schwartz 1999; Churchland & Shenoy 2007). A similar time lag could be achieved with our data by shifting the movement start to the start of the first FA epoch (time of 20% max factor loading). To compare fits at different lags of interest, we repeated this procedure by shifting the start of the movement to the beginning of each of the four FA epochs. We found a clear pattern of lag-dependent velocity encoding that was consistent for data from both monkeys: when the movement was aligned to the start of the first epoch, the velocity-speed model was preferred for about 80% of units based on AIC. This value dropped to about 40-50% for each of the remaining epochs. It is not surprising that the velocity-speed model was still preferred over the direction-only model for many units later in the

movement, given previous reports of this effect (Churchland & Shenoy 2007). However, the effect of velocity on firing rates was clearly largest immediately before the movement began, in line with the hypothesis that a dominant kinematic signal drives early-movement firing for both reach types.

4.3.6 Neuronal covariance patterns match the movement velocity profile

In the previous section, we concluded that a velocity-like signal is present in firing rates during both reach types. The high correlations of responses between reach types early in the movement (Figures 4.5, 4.7, Supplemental Figure 4.12) were therefore likely related to this consistent kinematic encoding. Another way to illustrate this signal without using explicit models is to compute the bin-by-bin covariance of responses between reach types, in a similar manner to the sliding correlation analysis used previously (i.e. Figure 4.5). Covariance reflects the similarity of encoding patterns as well as their magnitudes, so we would expect the temporal profile of this measure to reflect the envelope of whatever signal (if any) is shared between tasks. This measure is illustrated for one set of example responses in Figure 4.8A. The top and middle panels show the response of one unit for "hold" and "no hold" trials, respectively, while the bottom panel shows the covariance between the two reach types at each time bin. The covariance profile rises and falls along with the first response phase, which was similar for both reach types. Importantly, covariance remained low during the second response phase, which differed between reach types. Qualitatively, the covariance profile looks similar to the stereotyped bell-shaped velocity profiles characteristic of reaching movements.

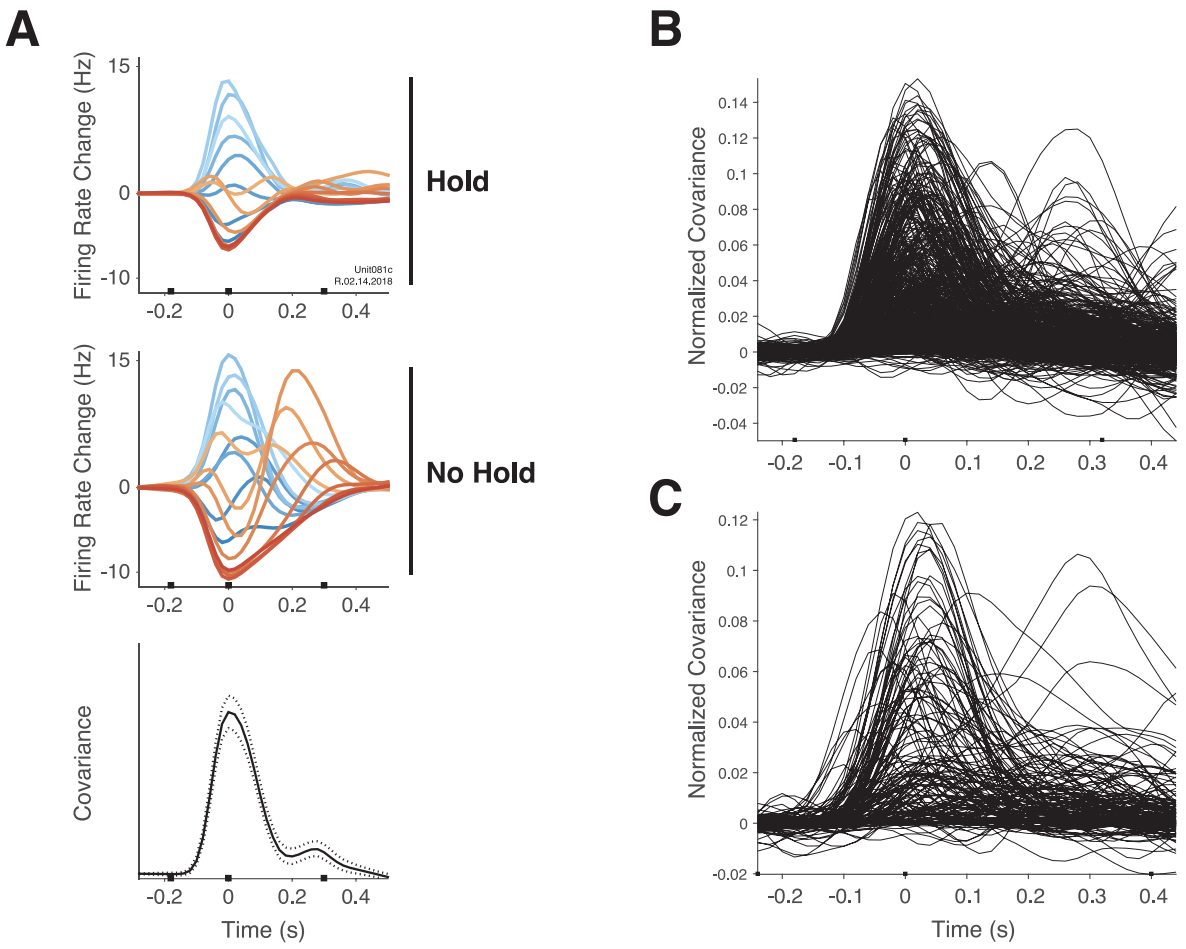

Figure 4.8 Bell-shaped covariance profiles of responses in "hold" and "no hold" trials.

A) Single unit response in “hold” (top panel) and “no hold” (middle panel) trials. Bottom row shows a sliding covariance analysis. This analysis is similar to the correlation analysis shown in Figure 4.5 but is unscaled so as to reflect the magnitude of responses as well as their similarity across targets. This unit response also appears in Figure 4.4. B) Same as in the bottom row of panel A, but shown for all unit responses recorded from monkey R. C) Same as in panel B, but for data from monkey P. For most units, the temporal profile of the covariance between reach types was bell-shaped.

Bell-shaped covariance profiles were very common for unit responses recorded from both monkeys. Figure 4.8 B-C shows these profiles overlaid for all recorded units (black traces). Because covariance depends on absolute response magnitude, rates were first normalized by their range so that data from each unit could be displayed on the same scale. Most of the covariance between reach types appeared early in the movement and had a bell-shaped temporal profile. High covariance later in the movement was uncommon, as expected from the correlation-based analyses described previously (Figure 4.7). To summarize the covariance data over the population of units, we added their values together bin-by-bin. This operation is equivalent to computing the vector field covariance (Eq. 4.5, see Section 4.2.3) of the population response between reach types, separately for each time bin. That is, for each time bin, we computed Cov(A,B) where A and B are N-by-C matrices of rates during "hold" and "no hold" trials, N is the number of units, and C is the number of reach directions. The result of this operation is illustrated in Figure 4.9A-B (blue traces) along with the speed profile, which was scaled and time-shifted for comparison (black traces). The correspondence between the covariance and speed profiles is quite good, indicating a kinematic-like signal is shared between reach types. For both monkeys, the match between the two profiles was somewhat worse toward the end of movement, likely because a small portion of responses were well-correlated between reach types at that time (Figure 4.7, fourth column).

Although units were typically well-modulated throughout a trial's duration, firing rates tended to be largest early in the movement. Therefore, a concern is that the bell-shaped covariance profile might simply reflect the modulation envelope within a single reach type, rather than the envelope of shared encoding between each type. To address this possibility, we repeated the analysis using only rates from the "no hold" trials. That is, we computed Var(B)

instead of Cov(*A*,*B*), with variables defined as before (Figure 4.9 C-D, dark gray traces). The use of data from "no hold" trials was conservative, in that late-trial activation tended to be stronger for "hold" trials (see Chapter 2). The resulting variance profiles for both monkeys were found to be largest around movement onset and initially matched speed profiles well. However, variance did not subsequently drop to the same extent as covariance, suggesting the drop in the latter was driven in part by dissimilar encoding patterns (e.g. Figure 4.7, Supplemental Figure 4.12).

As an additional comparison, we also computed the covariance profile of EMG activity between reach types (Figure 4.9 C-D, light gray traces). As expected, the initial rise in the covariance profile was comparable to that for both the neural covariance and the movement speed profiles (see Moran & Schwartz, 1999). However, the EMG covariance profile had a second large peak later in the movement, indicating a similar modulation pattern for "hold" and "no hold" trials. This result is in line with the correlation analyses described in Section 4.3.1, which suggested the muscles in our datasets were active in a comparable way during both reach types, even during the terminal reach phase.

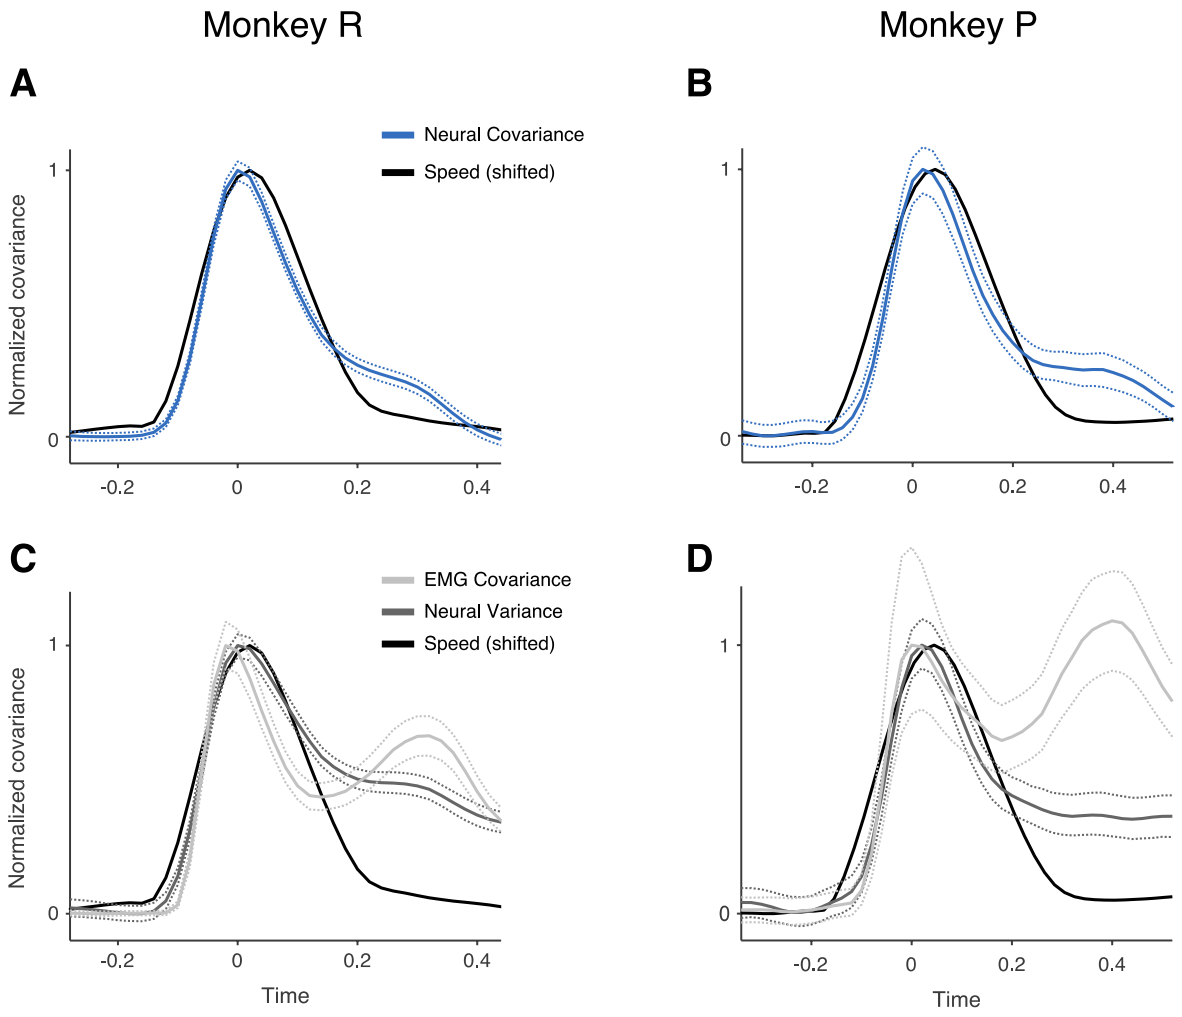

Figure 4.9 Population covariance profiles closely match speed profiles.

A) Vector field covariance (blue trace, see Equation 4.5) for firing rates computed between "hold" and "no hold" trials. The covariance temporal profile closely matched the movement speed profile (black trace, scaled and time-shifted for comparison). Data from monkey R. B) Same as in A but for data from monkey P. C) Light gray trace: vector field covariance for EMG activity computed between "hold" and "no hold" trials. Dark gray trace: vector field variance for firing rates computed during "no hold" trials. Neither measure matched the speed profile (black trace) as well as the measure in A-B. Data from monkey R. D) Same as in C but for monkey P. Dashed traces in each panel show 95% confidence interval computed over 1000 bootstrap iterations.

4.3.7 Visuomotor adaptation affected firing rates differently according to reach type

We have presented evidence that early-movement encoding is related to velocity regardless of how a reach is terminated, but the drivers of subsequent encoding differences have not been clarified. In Chapter 3, we found that directional tuning changed after monkeys adapted to a visuomotor rotation, and here we used this property to probe encoding differences during "hold" and "no hold" trials. Before comparing adaptation effects between the two reach types, we first examined results for a single type. As previously reported, we found clear encoding changes after monkeys adapted to the rotation (Supplemental Figure 4.13). Moreover, adaptation effects could differently impact each response component of single units, as previously observed (Supplemental Figure 4.14). Interestingly, when we compared these effects for the two reach types, we found a clear interaction between adaptation-related tuning changes and the hold requirement (Figure 4.10). For example, the response illustrated in Figure 4.10A had a single modulation component during "hold" trials and two components during "no hold" trials (left vs. right columns). The PD of the first component did not change after adaptation for either reach type (black bar above rates). However, the PD of the second response component, which was unique to "no hold" trials, changed by 43° after adaptation. Other examples of reach-specific adaptation effects can be seen in Figure 4.10 B-C. Those responses had two modulation components for each reach type, but the second component was adaptation-sensitive only during "no hold" reaches.

It is worth noting that the two reach types were randomly interleaved in both the "standard" and "rotated" blocks. Therefore, it appears that adaptation effects could change trial-to-trial depending on the hold requirement. These data provide further evidence that information encoded during each epoch could change discretely. A striking example of this patterning is

illustrated in Supplemental Figure 4.15, which replicates the sliding correlation technique used previously. Correlations of responses for three pairs of conditions (different colored traces) are shown for the unit response in Figure 4.10B. In each case, the correlation patterns plateaued during modulation components at values determined by task features.

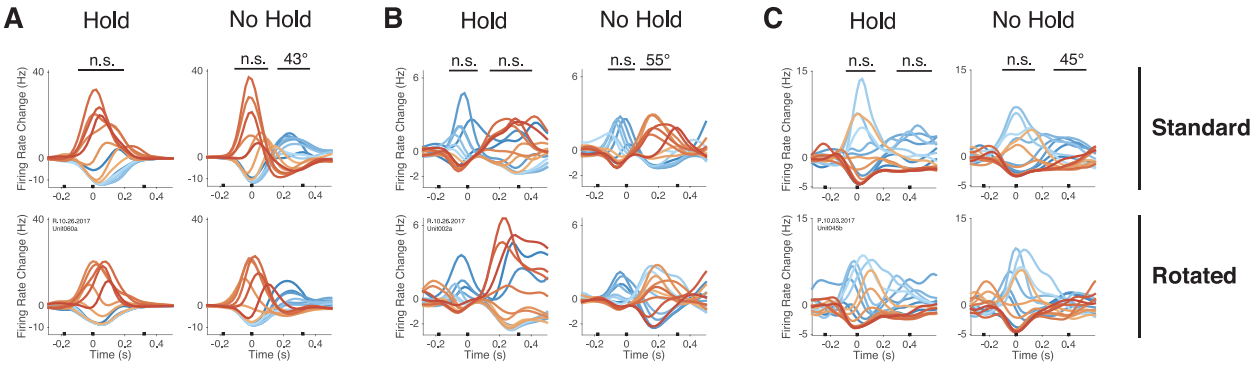

Figure 4.10 Single unit responses with an adaptation-sensitive component only during "no hold" trials.

Example single unit responses in all four task conditions. Top and bottom rows show responses during the "standard" and "rotated" blocks; left and right columns in each panel show responses in "hold" and "no hold" trials. Each panel shows a response with an adaptation-sensitive firing component only during "no hold" trials. Black bars above response components mark time periods where we compared tuning between standard and rotated blocks. Numbers above bars show the angular PD-difference between blocks, and imply that this difference was statistically significant ($p < 0.05$); n.s. indicates no significant difference ($p >= 0.05$). The response in panel B also appears in Figure 4.5.

4.3.8 Adaptation effects late in the movement were more common during "no hold" trials

In the previous section, we described several example responses with adaptation-sensitive tuning late in the movement, but only during "no hold" trials. To determine how prevalent this pattern was across the population, we computed the percentage of units with an adaptation-related tuning change separately for each reach type. Since there was a time-dependent effect of adaptation (e.g. Figure 4.10 and Supplemental Figure 4.14), we utilized the FA trial segmentation and

evaluated tuning of rates averaged over 100 ms windows centered on each factor peak (Figure 4.11). As reported in Chapter 3, the percentage of adaptation-sensitive responses was larger for 90° than 45° sessions (panel A vs. B). For both monkeys, we observed a moderately higher proportion of adaptation-related tuning changes during the third and fourth epochs for "no hold" reaches, compared to "hold" reaches (yellow and blue traces, respectively). For monkey R only, we observed a slightly larger percentage of adaptation-sensitive responses for "hold" reaches during the second epoch. This may explain the somewhat lower correlations between "hold" and "no hold" responses during that epoch for monkey R relative to monkey P (Figure 4.7 A vs. B, second column). In summary, our data suggest that encoding differences between reach types are partially explained by a difference in the degree of visuospatial encoding present for each type.

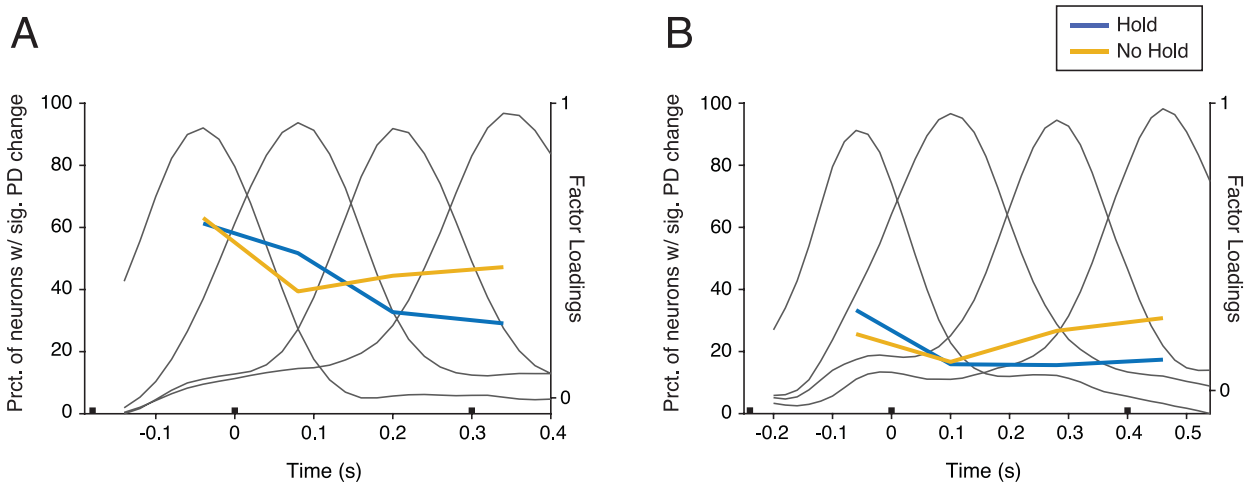

Figure 4.11 Percentage of adaptation-sensitive responses in "hold" and "no hold" trials

Tuning curves for each unit response were compared before and after adaptation for rates averaged over 100 ms windows centered on each factor peak (black Gaussian-shaped traces, scale on right ordinate). The percentage of units with a significant PD change was computed for "hold" and "no hold" trials separately (blue and yellow traces, scale on left ordinate). A) Data from monkey R (90° rotation). B) Data from monkey P (45° rotation). For both monkeys, adaptation-related PD changes were somewhat more common late in the trial for the "no hold" reach type.

4.4 DISCUSSION

In Chapter 2, we found that late-movement encoding was weaker when monkeys did not need to stop and hold at a target. We suggested this activity was therefore related to stopping the arm. Reaching is believed to consist of multiple discrete behavioral phases (Woodworth 1899; Soechting & Lacquaniti 1981; Meyer et al 1988), and each may be associated with different patterns of neural activity. Here, we found that late-movement encoding seemed to constitute a discrete component of single-unit firing, suggesting reach structure can be recognized in M1 activity. During "no hold" reaches, some units seemed to omit a response component that was seen during "hold" reaches (Figure 4.3), and this omission could be visualized as a sharp, abrupt change in tuning correlation between the responses. One might interpret this observation to mean that M1 was simply less engaged when reaches were not accurately stopped. However, we also observed units with response components unique to "no hold" reaches. Sometimes those components were appended to the response seen for "hold" reaches (Figure 4.4 A-C), and sometimes we found two distinct sequences depending on reach type (Figure 4.5 B-D). This pattern suggests M1 was differently engaged for "no hold" reaches, rather than simply less engaged.

In contrast to these late-movement differences, we found reaction-time and early movement activity was well-matched regardless of how a reach was terminated. This could be seen qualitatively for unit responses (Figures 4.3 and 4.4) and quantitatively using correlation-based analyses (Figures 4.5 and 4.7, and Supplemental Figure 4.12). This result is interesting given that M1 reaction-time activity is widely hypothesized to directly drive subsequent activity patterns, possibly via local recurrent circuits (Churchland & Shenoy 2007; Churchland et al 2010; Afshar et al 2011; Churchland et al 2012; Shenoy et al 2013; Ames, Ryu, & Shenoy 2014;

Sussillo et al 2015; Elsayed et al 2016; Lara et al 2018; Russo et al 2018). Under this "dynamical systems" hypothesis, pre-movement activity sets the "initial conditions" of a deterministic self-driven system; the early reach-independent patterns we observed therefore should not have diverged depending on the hold requirement. There are several possible conclusions to be drawn from our data. One possibility is that pre-movement M1 activity is simply not directly causal to subsequent activity. This perspective is supported by recent optogenetic work in rodents that used a laser to hold M1 in a "known" initial state, which then led to divergent activity patterns trial-to-trial when the laser was turned off (Sauerbrei et al 2018). Another possibility is that the small percentage of pre-movement responses that were weakly correlated between reach types (Figure 4.7 A-B, left-most column) was sufficient to generate divergent patterns. However, this explanation is at odds with experimental work suggesting the acceptable range of "initial conditions" is broad (Ames, Ryu, & Shenoy 2014; Michaels et al 2015). Alternatively, it may be that task-dependent inputs alter the time course of responses as behavior progresses. In that case, earlier activity may still play some causal role in determining later activity but would not be the sole driver of ongoing dynamics.

A related view is that response sequences reflect the behavioral structure of the task: the different target acquisition requirements of reaches may have driven the divergent neuronal responses. This divergence might be expected if there were large differences in motor output for the two reach types, but arm trajectories (Figure 4.1) and EMG activity (Figure 4.2) were comparable in each type (Section 4.3.1). Moreover, correlation and covariance structure of EMG patterns was similar across reach types, while the same was not true for neuronal data (Figure 4.7, 4.9, and Supplemental Figure 4.12). Although return movements back to the home position were necessarily delayed during "hold" trials, we did not find evidence that monkeys

were simply planning a return movement earlier for "no hold" trials (Table 4.1). Nor were encoding differences likely related to the tonic firing sometimes observed during hold periods (Kettner, Schwartz, & Georgopoulos 1988; Wang et al 2007), since encoding patterns diverged much earlier in the movement. Instead, our data are consistent with a two-component model of reaching (e.g. Woodworth 1899): the initial displacement may be "pre-programmed" and is associated with similar M1 activity across contexts, while subsequent target homing occurs as a distinct process that may be associated with different activity depending on target acquisition requirements.

The importance of visual feedback during the target-homing phase is well-established (Woodworth 1899; Keele & Posner 1968; Meyer et al 1988), and we therefore wondered if a difference in visuospatial encoding might help explain the different activity patterns we observed in each reach type. Indeed, when monkeys adapted to a visuomotor rotation, the prevalence of rotation-related tuning changes differed between the two types (Figures 4.10, 4.11, and Supplemental Figure 4.15). We anticipated a larger adaptation effect for "hold" trials since target-homing is reliant on vision, though we actually found the effect was slightly weaker for those trials relative to "no hold" reaches (Figure 4.11, blue vs. yellow traces). Although target acquisition occurs along with visual processing, corrective sub-movements must then be generated to carefully stop inside a target zone. Our data suggest that the "more motoric" aspects of those sub-movements may be preferentially encoded in M1 when a more accurate terminal phase is imposed by the task (cf. Rouse 2018). It is important to stress that "adaptation effects" on M1 coding do not necessarily imply a direct effect of visual feedback, though we presented evidence in Chapter 3 that visual feedback can drive M1 firing.

Numerous studies have shown that early-movement M1 activity correlates with upcoming reach velocity (Schwartz 1992; Moran & Schwartz 1999; Wang et al 2007; Inoue et al 2018), and we found this was true for our data as well (Section 4.3.5, Figures 4.8 and 4.9). To a lesser extent, late-movement encoding could also reflect velocity, a result previously reported by Churchland & Shenoy (2007). Their study replicated an analysis from Moran & Schwartz (1999) that showed the non-directional (across-target average) firing rate profile matched the movement velocity profile. In the original analysis, data collection stopped before the entire movement was completed; an extended analysis in Churchland & Shenoy (2007) found the relationship was less robust than originally reported, leading to the conclusion that M1 velocity encoding was incidental or secondary to network dynamics. Our finding that early- and late-movement responses are distinct may clarify the discrepancy between studies. That is, we found the velocity profile was most strongly represented in the early-movement response, and we leveraged the different late-movement responses in "hold" and "no hold" trials to accurately extract the kinematic signal (Figure 4.9). Computing non-directional firing rate is not sufficient to parse this signal from subsequent response components– as Churchland & Shenoy (2007) commented, the poor match with velocity was "largely caused by the ensemble neural response being broader than the velocity trace (many neurons showed response patterns that lasted longer than the movement itself)." Indeed, we observed a similar pattern when we considered only the activity for one reach type (Figure 4.9 C-D, dark gray traces, compare with Figure 8 in Churchland & Shenoy 2007).

Interestingly, a large percentage of responses during the early trial epochs showed tuning changes after monkeys adapted to the visuomotor rotation (Figure 4.11). This observation is important given our results suggesting reach velocity is strongly encoded during those early-trial

epochs. It may be that kinematic and visuospatial information simultaneously drive firing in M1. This hypothesis would be in line with previous reports of "mixed" single-unit tuning to multiple parameters (Lurito, Georgakopoulos, & Georgopoulos 1991; Shen & Alexander 1997a). A related possibility is that some neurons encode a kinematic signal in visuospatial rather than ego-centric coordinates. This would be useful, for example, for keeping a subject's "internal model" updated following adaptation (Shadmehr & Mussa-Ivaldi 1994; Miall & Wolpert 1996; Mehta & Schaal 2002).

As with previous studies, our data suggest that M1 encoding is not unitary but may relate to multiple distinct processes, concurrently and in sequence. Although single-unit firing was often temporally complex and heterogeneous, we also observed consistent structure that related to behavioral events and task constraints. Our findings support the view that M1 activity follows a "temporal parcellation scheme" as described by Johnson, Mason, & Ebner (2001). They suggest that "parameters are processed independently" in M1, and "the directional signal [observed in firing rates] is not only task dependent but also varies over time and behavioral sequence." We conclude that a behavioral perspective is useful for parsing seemingly-complex neuronal firing, and that reach structure is an important driver of M1 activity.

4.5 SUPPLEMENTAL FIGURES

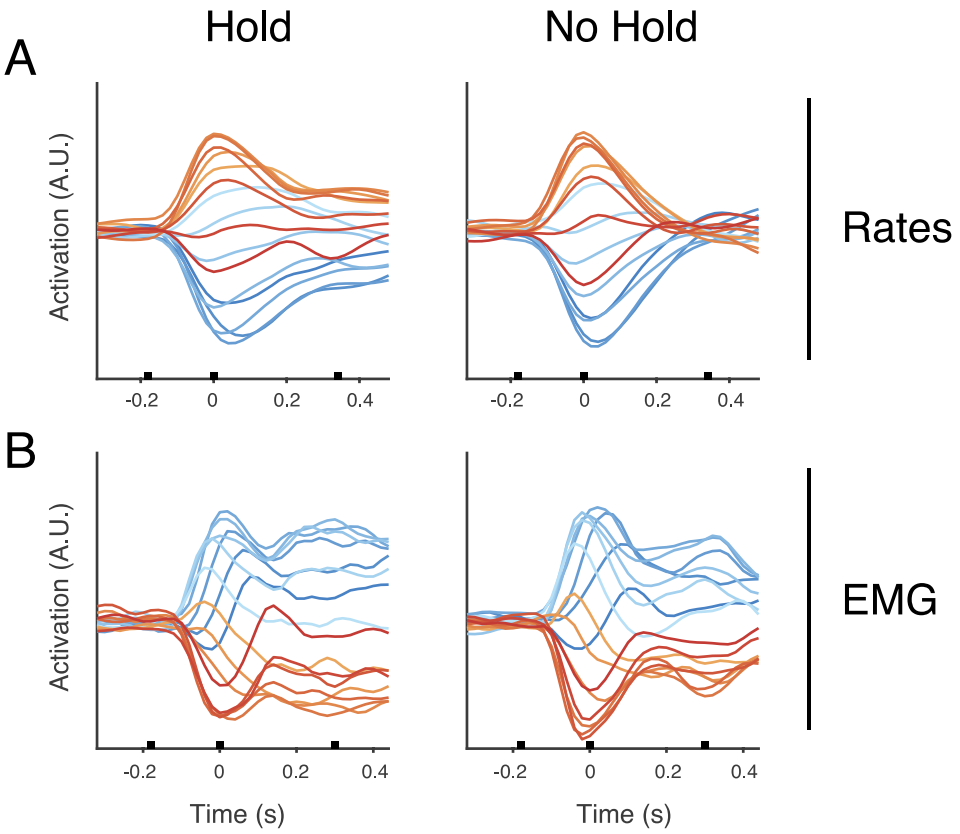

Figure 4.12 Principal component analysis of firing rates and EMG.

Principal components (PCs) were computed using data from only "hold" trials for firing rates (top row) or EMG data (bottom row). Data from "hold" or "no hold" trials were then projected onto the top PC (left and right columns, respectively). A) Firing rate data. Activation along the top PC was high throughout movement for "hold" trials, which was expected since the PC was computed to maximize variance explained for that reach type (left panel). Activation along this PC was high only during the early movement for "no hold" trials (right panel), indicating encoding patterns differed between reaches later in the movement. B) Same as in panel A but for EMG data. Activation along the top PC was comparable for both reach types, even though the PC was computed only for "hold" trials. Data in each panel from monkey R.

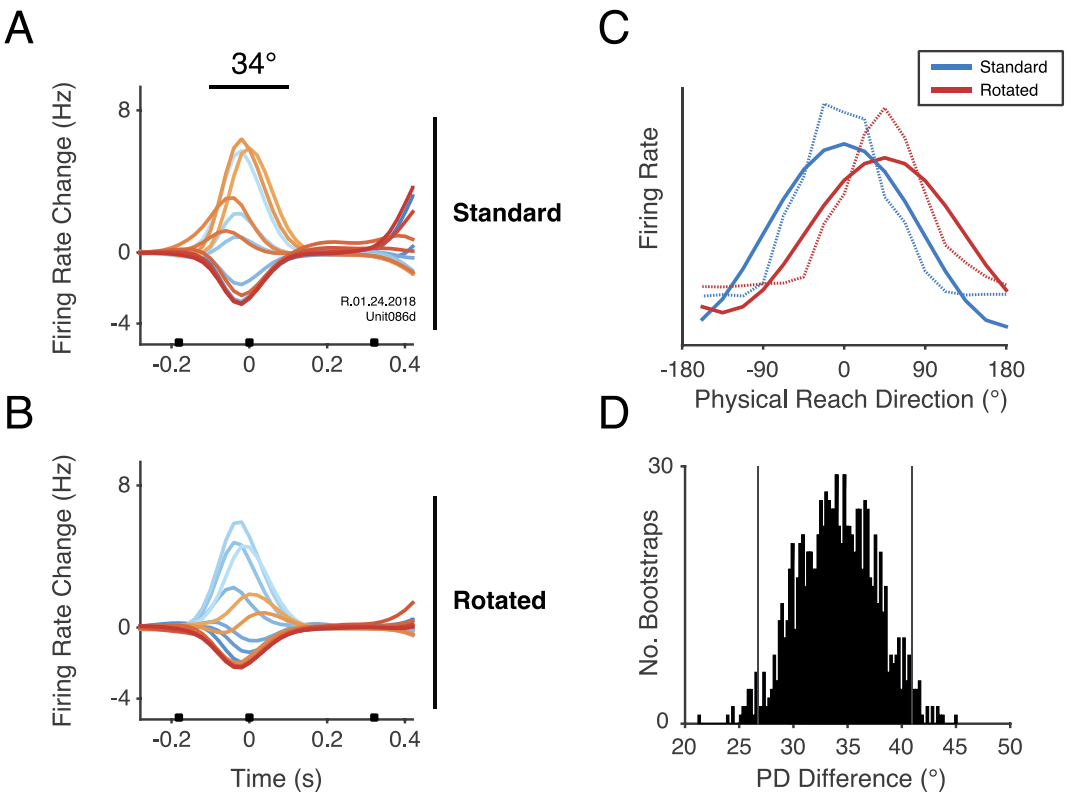

Figure 4.13 Example unit response showing tuning change following visuomotor adaptation.

A) Firing rates of one unit during the "standard" block. B) Firing rates of the same unit during the "rotated" block. C) Tuning curves computed relative to hand direction for the rates in A (blue traces) and B (red traces), averaged over the period marked by the black bar in panel A. Dashed traces show recorded rates, solid traces show cosine tuning fits (see Equation 4.1). D) Angular difference in degrees between the PDs of the cosine fits in panel C, computed 1000 times using a bootstrap procedure. PD difference was consistent over bootstrap iterations (vertical black bars bound 95% of the iterations). The median difference was 34° and is reported over the black bar in panel A.

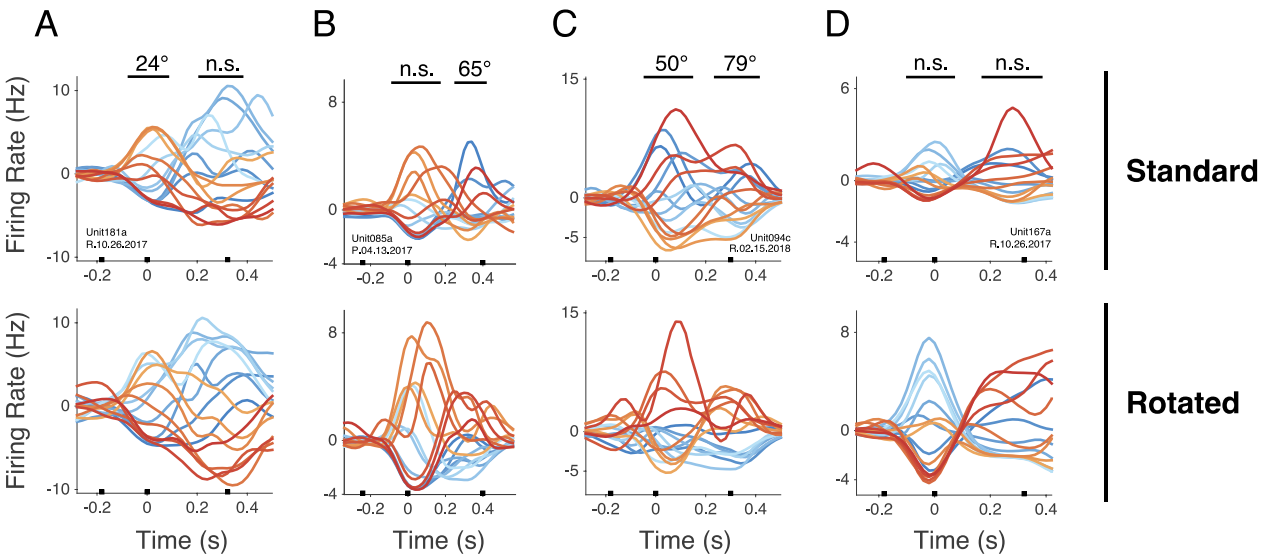

Figure 4.14 Example unit responses showing component-specific tuning related to adaptation.

Example multi-component single unit responses (columns) during the "standard" and "rotated" task blocks (top and bottom rows, respectively). PDs were compared between blocks over the periods marked by the black bars above each panel in the top row. A numerical value indicates that there was a significant difference ($p < 0.05$, bootstrap test, see methods); n.s. indicates no significant difference. A) A single unit response with a PD change after adaptation for an early firing component and not a later one. B) A response with a PD change for a late component and not an early one. C) A response with a PD change for both firing components, but each to a different degree. D) A response with two firing components, neither of which changed PD after adaptation.

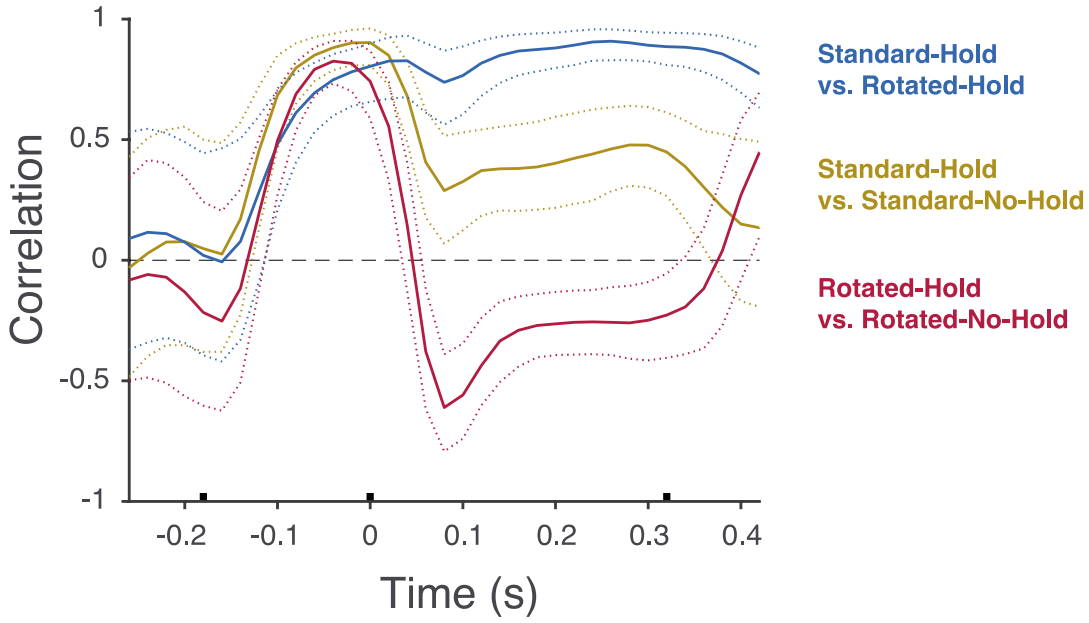

Figure 4.15 Correlation patterns change abruptly and depend on task parameters.

Sliding correlation analysis for selected pairs of conditions for the unit response shown in Figure 4.10B (main text). That response had an adaptation-sensitive firing component late in the trial only during "no hold" reaches. Since there was no PD change after adaptation for "hold" trials, correlation of responses between task blocks was high during the whole trial for that reach type (blue traces). The late-trial component during "no hold" trials was different from that during "hold" trials and changed PD after adaptation. The correlation between responses for the two reach types was therefore different before (yellow traces) and after (red traces) adaptation. Correlation patterns appeared to change in a step-like manner, suggesting the responses encoded information discretely depending on the task type.

5.0 CONCLUSIONS

5.1 SUMMARY OF FINDINGS

5.1.1 Firing patterns during reaching are structured

The central finding of the work presented here is that M1 firing patterns during reaching are structured as a sequence of discrete states, each with different encoding properties appropriate for a given behavior. In Chapter 2, we carefully characterized this structure during normal center-out reaching based on the changing directional tuning of single neurons. Using a novel statistical approach, we first confirmed that directional tuning is labile over the course of a reach. Tuning lability was not related to noise, nor was it related to poor assumptions about neurons' kinematic encoding properties. Yet, a cosine tuning function was a good approximation of firing rates throughout the entire movement, even if the PD of the tuning function was not consistent at different times in a trial. Close inspection of the time-course of single-unit tuning revealed a somewhat paradoxical property: PDs actually were, in fact, stable, but only for relatively brief (~100-200 ms) epochs. After a period of stable tuning, firing patterns could change abruptly and then stabilize again with a different PD. When rates for each movement direction were plotted together over time, we observed a coherent change in patterning that occurred around the same time during movement regardless of reach direction. Further, these abrupt, across-target pattern

changes tended to occur at similar times for most neurons in the population. We concluded that this widespread reorganization of firing patterns likely reflected a change in behavioral state as the movement unfolded.

5.1.2 Encoding structure and behavioral structure are linked

As striking as the observed pattern changes were, the strength of our conclusions in Chapter 2 was limited in a few important ways. We took single units' segmented tuning as evidence of separate, sequential neural drivers, but the independence of those segments could not be determined without additional experiments. Multi-phasic tuning changes would also be expected if oscillatory dynamics were the dominant neural driver (e.g. Churchland et al 2012), so perhaps the apparently discrete patterns were simply an artifact or epiphenomenon. In Chapter 3, we described the results of a visuomotor adaptation paradigm and found strong evidence that encoding segments were indeed related to distinct neural drivers. The dissociation between vision and movement allowed us to better characterize tuning within each response segment, and we could show clearly that each was independently impacted by the adaptation process. If segmented tuning was secondary to continuous evolution rules (for example), we would not expect to see such clear demarcation in information content between each segment. We next varied visual feedback during movement trial-to-trial and found further evidence that sequential response segments are associated with distinct drivers. For many single units, a late-movement firing component that was observed during normal reaches (i.e. with visual feedback) was weak or absent when no feedback was provided, while early-movement activity did not depend on feedback. This pattern suggested the drivers of the late- and early-movement response components were distinct.

In Chapter 4, we more directly linked segmented tuning to the structure of reaching movements. Point-to-point reaches consist of at least two phases, and accurate target acquisition occurs as a process separate from the initial arm displacement. We found that response sequences of M1 neurons depended on the target acquisition requirements of the task. Initially, sequences were very similar regardless of these requirements, in line with the idea that the initial part of movement is relatively pre-programmed. However, firing patterns abruptly diverged approximately midway through movement depending on the way the target was acquired. Accurate target acquisition relies on visual feedback, and we found that visuospatial encoding during the later part of movement also depended on the way targets were acquired. We suggest that reach structure is a major driver of M1 firing patterns.

5.2 RELATION TO CONTEMPORANEOUS WORK

Multiple studies have highlighted the temporal complexity of M1 firing during behavior (Sergio & Kalaska 1998; Hatsopoulos, Xu, and Amit, 2007; Churchland & Shenoy 2007; Churchland et al 2010; Suminski et al., 2015), and recent work has attempted to define simple rules that explain this complexity (Churchland et al 2012; Ames et al 2014; Russo et al 2018). Analysis of real M1 responses and simulated neuronal networks has suggested that activity patterns evolve smoothly in time, and that rotational dynamics dominate the population response (Churchland et al 2012; Sussillo et al 2015). Here, we found that discrete state transitions were a common feature of M1 responses, seemingly in contrast to smooth dynamics. However, these state changes are not necessarily incompatible with the rotations described by Churchland et al (2012); large-scale changes in encoding pattern (i.e. from one state to the next) would almost certainly produce a

shift in "state-space" position when population activity is visualized using dimensionality-reduction methods. It may be that these "summary" subspaces blur underlying discrete transitions, possibly due to variable timing of single-unit responses (e.g. Cisek & Scott 1999). Interestingly, some of the rotational state-space trajectories reported by Churchland et al (2012) actually appear to sharply transition at right angles, rather than smoothly evolve in time (see for example their Figure 3d, black traces). Similar abrupt state-space transitions have also been found using M1 activity during reach-to-grasp movements (Rouse & Schieber 2018, their Figure 3). Other analytical approaches may be better suited to highlighting these sharp transitions (e.g. Abeles et al 1995; Seidemann et al 1996; Latimer et al 2015).

A common hypothesis of advocates of smooth evolution rules in M1 is that intrinsic connectivity of local circuits is the primary determinant of firing rates. A related concept is that neuronal activity is confined to an intrinsic "manifold" of activity patterns (Sadtler et al 2014; Gallego et al 2017). Our work indicates that such a manifold may need to be fairly large to accommodate the encoding flexibility we observed depending on task structure (cf. Johnson & Ebner 2000). This flexibility somewhat contrasts with work described by Gallego et al (2017) suggesting that population-level activity patterns in M1 are comparable over a wide range of different tasks. It is important to note a major methodological difference between their study and ours: we typically subtracted off the condition-independent response component prior to analysis, as in Churchland et al (2012), while Gallego et al (2017) did not. Critically, this component of the response carries more variance than the condition-dependent (i.e. tuned) component (Kaufman et al 2016; Rouse & Schieber 2018). Indeed, Gallego et al (2017) reported that the vast majority of variance shared between disparate tasks did not appear tuned. Thus, our findings may not be at odds; perhaps the condition-independent response is shaped by intrinsic

connectivity, while the tuned response is more flexible and task-dependent. This knowledge could be important for developing novel control algorithms for brain-computer interfaces (BCIs).

5.3 CLINICAL IMPLICATIONS

Through much research involving BCIs, investigators have found that signals from the primary motor cortex afford the best opportunity to "read out" movement commands. In lab settings, these devices function reasonably well for computer cursor and anthropomorphic robotic limb control (Velliste et al, 2008; Wodlinger et al, 2015; Gilja et al 2012). However, speed, accuracy, and dexterity are often limited. Recent efforts to provide tactile feedback via intracortical microstimulation hold great promise for improving control, but limitations of decoding algorithms may impose a ceiling on performance. Many novel decoders have been proposed (Kemere et al 2008; Li et al 2011; Kim et al 2011; Gilja et al 2012), typically with limited performance gains over the decades-old population vector algorithm (PVA, Collinger, Gaunt, & Schwartz 2018). Yet, the PVA assumes each neuron contributes a constant directional signal– fundamentally, this assumption is incorrect. In one striking demonstration of decoding limitations, a human subject with tetraplegia controlled a prosthetic limb with ten degrees of freedom using the optimal linear estimator (OLE), which is a variant of the usual PVA (Wodlinger et al 2015). Although the subject achieved good performance in translating, orienting, and shaping the hand, her ability to interact with objects was severely limited. It is likely that non-stationarities in neural encoding, probably related to object information, interfered with the decoding process (e.g. Vargas-Irwin et al 2015; Rouse & Schieber 2016b). Investigators were able to use ad hoc re-calibration methods to improve performance, but they

made no attempt to make the decoder "aware" of neuronal temporal dynamics. It is possible that a novel decoder that explicitly models the state transitions we have described would afford subjects better, more flexible BCI control.

In support of this idea, recent work from our lab (Suway et al 2013) demonstrated that explicit state-modeling improved continuous BCI control and eliminated the need for investigator-defined task epochs. We leveraged the observation that M1 firing patterns change abruptly when subjects transition from actively moving to idly waiting (Velliste et al 2014), and this transition was detected online from firing rates using linear discriminant analysis (LDA). When the LDA reported that monkeys did not intend to move, output to the prosthetic limb was halted. Without this output-gating method, we found that, when monkeys became distracted or disinterested in the task, the usual decoder erroneously reported non-zero velocity and the arm moved in unpredictable and potentially dangerous ways. Although we calibrated the LDA to recognize only two states (active and idle), the work presented here suggests that the "active" state is not unitary but consists of several distinct components. In Chapter 2, we explicitly classified these different "active" states from firing rates, again using LDA, and found that we could robustly detect each state in sequence. This process was useful for updating the state-specific PD of each neuron for offline decoding with the PVA. Future experiments may be able to replicate this approach online as was done for idle state detection (Suway et al 2013).

5.4 LIMITATIONS AND FUTURE DIRECTIONS

We have repeatedly referred to segmented encoding epochs as neural "states," but this term is admittedly informal. Numerous analytical methods can be used to characterize functional

connectivity between neurons (e.g. Vaadia et al 1995; Okatan, Wilson, & Brown 2005; Dann et al 2016), and these methods may be valuable for formalizing our notion of functional states. For example, Abeles et al (1995) found that pairwise cross-correlations between neurons could change on short timescales along with transitions in behavioral state. These changing interactions between neurons likely reflected changes of network state, and similar measures could be assessed for the "states" we have described here. Elsayed et al (2016) reported a related phenomenon as monkeys transitioned from planning to executing a movement. Their method was based on relatively long periods of time-averaged firing rates, and so short-latency (i.e. possibly mono- or poly-synaptic) neuronal interactions could not be assessed. However, they could clearly show widespread reorganization of correlation structure between neuronal responses, which likely corresponded to a transition between two distinct network states. Future work could determine whether neuron-neuron relationships change along with the sequence of encoding epochs we have described.

A further limitation of our description of neural "states" is a reliance on trial-averaged firing rates. Although this approach effectively manages trial-to-trial variability in unit responses, slight differences in spike timing or behavioral alignment across trials can obscure underlying single-trial dynamics. For example, we showed in Chapter 2 that abrupt changes in directional tuning were associated with a gradual rise-and-fall of the modulation profile of single unit responses. It is not clear if firing rates genuinely ramp up and down over the course of tuning epochs, or if the "edges" of these modulation phases are blurred out due to timing variability over trials. Latimer et al (2015) described this problem in detail for spike-trains recorded from the lateral intraparietal area, and they challenged a common belief that firing rates ramp upward or downward while monkeys accumulate evidence during decision-making. They

argued that ramping could be an artifact of trial-averaged responses that actually changed discretely at different times on each trial. By comparing single-trial models of ramp-like or step-like spike train dynamics, they demonstrated that more than three quarters of single units produced step-like responses. Our claims that M1 responses progress through discrete states could be further supported by similar statistical approaches.

Beyond analytical limitations, it is worth noting the shortcomings in our experimental design. We have referenced neural "drivers" under the assumption that inputs to M1 generate a sequence of firing patterns, though we have not measured activity from any potential sources of input. It would be valuable for future work to include recordings from several brain regions and attempt to trace information flow between areas (e.g. Pesaran, Nelson, & Andersen 2008). We recorded neural data using microelectrode arrays, and it would be straightforward to adapt our approach to record from additional cortical structures (e.g. Dann et al 2016). For visuomotor adaptation experiments, posterior parietal cortex (Clower et al 1996; Buneo & Andersen 2006) and ventral premotor cortex (Kurata & Hoshi 1999; Schwartz, Moran, & Reina 2004) would be relevant structures to record from simultaneously along with M1. Although we did not use recording chambers in our experiments, this would be valuable for future experiments to allow recording from cerebellar structures involved in motor adaptation (Wolpert, Miall, Kawato 1998; Baizer, Kralj-Hans, & Glickstein 1999; Tseng et al 2007). These data would be useful for evaluating the functional topology of brain structures (Bullmore & Sporns 2009), which may help clarify the sources of the signals we have described in M1.

A further limitation of our experimental design was that we only observed neuronal responses during center-out reaching and not other types of movements. We therefore could not characterize state sequences during, for example, continuous drawing or random target tracking.

Schwartz & Moran (1999) used the PVA to predict hand path during figure-eight tracing and found a clear effect of path curvature on the prediction interval between the population vector and the movement. They reasoned that "curvature defines the way a movement is segmented," and concluded that "this segmentation is obvious in the neuronal firing patterns" after applying the PVA. One possibility is the state transitions we have described here align with transitions defined by movement curvature. During random target tracking, Harpaz et al (2018) found that M1 responses could be segmented into a few discrete states using hidden Markov models, and the identified states aligned with changes in acceleration. These observations agree with our conclusion that task structure defines neural state sequences. Additional work could clarify the interaction between the state transitions we identified (based on visuospatial coding, visual feedback, or endpoint strategy) and kinematic segmentation (i.e. as defined by curvature or acceleration).

As is common in related work, we studied neuronal activity during well-practiced ("over-trained") behaviors, and so there is concern that an effect of learning on activity patterns could limit the generalizability of our results. For example, our behavioral tasks strongly emphasized the importance of visuospatial information presented in the virtual reality environment, and the extreme visuomotor perturbations we imposed are outside the normal behavioral repertoire of macaques. Moreover, we did not attempt to characterize any changes in neuronal response patterns over weeks or months while monkeys gradually mastered these difficult behaviors. Prior exposure to visuomotor perturbations is associated with "savings" in terms of adaptation time during subsequent exposure, and this likely reflects learning-related changes in neural activity (Krakauer, Ghez, & Ghilardi 2005). One possibility is that segmented encoding

sequences reorganize during learning; future work could focus on determining whether the (re-)organization of these sequences constitutes a reasonable neural "substrate" for learning.

A related question is the degree to which motor adaptation can be recognized as a slow shift in encoding properties over trials, and whether this shift is correlated to behavioral performance. The data presented here focused only on the pre- and post-adaptation periods of the task, and we did not analyze responses during the adaptation period. This would have been challenging because we applied the visuomotor rotation quickly over ~20 trials, which caused monkeys to make distorted trajectories while they adapted. Therefore, it would be difficult to determine if changes in firing patterns were related to the adaptation process or to gross changes in motor output (cf. Wise et al 1998). It may be possible to study this process using an alternative paradigm in which the visuomotor rotation builds very gradually over, say, hundreds of trials. This would likely allow monkeys to produce straight movements on each trial as they adapted, and changes in encoding pattern could be evaluated over this extended adaptation period.

Many studies have examined functional throughput from M1 neurons to muscles during behavior. Here, we focused on the structure and flexibility of M1 firing patterns, but we have not determined how these patterns may influence muscle excitability. Griffin, Hoffman, & Strick (2015) demonstrated that corticomotoneuronal (CM) cells maintain a consistent functional relationship to their target muscles over a wide range of limb configurations. This finding suggests there is a consistent output channel from M1 to muscles, and it would be valuable to identify how this output is represented in the patterns we have characterized. Although we recorded from a slightly different location, it is likely that post-spike effects on EMG could be identified from units in our datasets. Spike-triggered averaging of EMG would therefore be a

useful method to apply to data during the described tasks. It is possible that, despite the changing encoding properties of units, their spiking activity affects targets muscles in a consistent way.

Alternatively, throughput from single units to muscles might change between the encoding epochs we identified. Kaufman et al (2014) reasoned that this must be true during instructed-delay tasks when monkeys transition from planning to moving; most M1 neurons are active during both task periods while muscles are active only during movement. They identified linear combinations of movement-period M1 rates that correlated well with EMG activity, and then evaluated the same combinations using delay-period rates. Delay-period combinations were near-zero valued, which led to the conclusion that certain patterns of M1 activity are structured to "cancel out" and attenuate throughput to muscles. Their methods could be adapted to our data to evaluate if a similar phenomenon applies for the different encoding segments we found. In that case, it may be that some segments reflect inputs to the system while others reflect outputs to other structures of the nervous system. Identifying how the signals we have described relate to inputs, outputs, and local computation could be a topic of future work.

BIBLIOGRAPHY

Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, Tishby N, Vaadia E. 1995. Cortical activity flips among quasi-stationary states. Proc Natl Acad Sci, 92:8616-8620

Aflalo TN, Graziano MSA. 2006. Partial tuning of motor cortex neurons to final posture in a free-moving paradigm. Proc Natl Acad Sci USA. 103:2909-2914.

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.

Alexander GE, Crutcher MD. 1990a. Preparation for movement: neural representations of intended direction in three motor areas of the monkey. J Neurophysiol. 64:133-150.

Alexander GE, Crutcher MD. 1990b. Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. J Neurophysiol. 64:164-178.

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

Asanuma H, Sakata H. 1967. Functional organization of a cortical efferent system examined with focal depth stimulation in cats. J. Neurophyiol. 30:35-54.

Ashe J, Georgopoulos AP. 1994. Movement parameters and neural activity in motor cortex and area 5. Cereb Cortex. 6:590–600.

Baizer JS, Kralj-Hans I, Glickstein M. 1999. Cerebellar lesions and prism adaptation in macaque monkeys. 81(4):1960-1965.

Bennett KM, Lemon RN. 1996. Corticomotoneuronal contribution to the fractionation of muscle activity during precision grip in the monkey. J Neurophysiol. 75:1826–1842.

Bernstein NA. 1984. The problem of the interrelation of co-ordination and localization. In: Whiting HTA, editor. Human motor actions. Bernstein reassessed (Advances in Psychology 17). Amsterdam (North Holland): Elsevier. p 77–119.

Bizzi E, Accornero N, Chapple W, Hogan N. 1984. Posture control and trajectory formation during arm movement. J Neurosci. 4(11):2738–2744.

Bortoff GA, Strick PL. 1993. Corticospinal Terminations in Two New-World Primates: Further Evidence That Corticomotoneuronal Connections Provide Part of the Neural Substrate for Manual Dexterity. J. Neurosci. 13(12): 5105-5118.

Bossom J. 1965. The effect of brain lesions on prism-adaptation in monkey. Psychon. Sci. 2:45-46.

Brinkman J, Kuypers HGJM. 1972. Splitbrain monkeys: cerebral control of ipsilateral and contralateral arm, hand, and finger movements. Science. 176:536–539.

Bullmore E, Sporns O. 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. Nature Reviews Neuroscience. 10:186-198.

Buneo CA, Andersen RA. 2006. The posterior parietal cortex: Sensorimotor interface for the planning and online control of visually guided movements. Neuropsychologia. 44:2594-2606.

Buzsáki G. 2010. Neural syntax: cell assemblies, synapsembles, and readers. Neuron. 68:362-385.

Caminiti R, Johnson PB, Urbano A. 1990. Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. J Neurosci. 10(7):2039–2058.

Cheney PD, Fetz EE. 1980. Functional classes of primate corticomotoneuronal cells and their relation to active force. J Neurophysiol. 44:773-791.

Cherian A, Krucoff MO, Miller LE. 2011. Motor cortical prediction of EMG: evidence that a kinetic brain-machine interface may be robust across altered movement dynamics. J Neurophysiol. 106:564-575.

Churchland MM, Cunningham JP, Kaufman MT, Foste 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, Scott SH. 1999. An alternative interpretation of population vector rotation in macaque motor cortex. Neuosci Letters. 272:1-4.

Clower DM, Hoffman JM, Votaw JR, Faber TL, Woods RP, Alexander GE. 1996. Role of posterior parietal cortex in the recalibration of visually guided reaching. Nature. 383:618-621.

Collinger JL, Gaunt RA, Schwartz AB. 2018. Progress towards restoring upper limb movement and sensation through intracortical brain-computer interfaces. Current Opinion in Biomedical Engineering. 8:84-92

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. 2016;5:e15719. DOI: 10.7554/eLife.15719.

di Pellegrino G, Wise SP. 1993. Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. J Neurosci. 13:1227-1243.

Diedrichsen J, Shadmehr R, Ivry RB. 2010. The coordination of movement: optimal feedback control and beyond. Trends Cogn Sci. 14:31–39.

Dum RP, Strick PL. The Origin of Corticospinal Projections from the Premotor Areas in the Frontal Lobe. J. Neurosci. 11(3):667-689.

Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP. 2016. Reorganization between preparatory and movement population responses in motor cortex. Nature Communications. 7:1-15.

Evarts EV. 1965. Relation of discharge frequency to conduction velocity in pyramidal tract neurons. J Neurophysiol. 28:216-228.

Evarts EV. 1966. Pyramidal tract activity associated with a conditioned hand movement in the monkey. J Neurophysiol. 29(6):1011-1027.

Evarts EV. 1968. Relation of pyramidal tract activity to force exerted during voluntary movement. J Neurophysiol. 31(1):14-27.

Evarts EV. 1969. Activity of pyramidal tract neurons during postural fixation. J Neurophysiol. 32(3):375-385.

Evarts EV, Fromm C. 1977. Sensory responses in motor cortex neurons during precise motor control. Neuroscience Letters. 5:267-272.

Feldman A. 1966. Functional tuning of the nervous system with control of movement or maintenance of a steady posture. II. Controllable parameters of the muscle. Biophysics. 11:565–578.

Fetz EE. 1969. Operant conditioning of cortical unit activity. Science. 163(3870):955-958.

Fetz EE. 1992. Are movement parameters recognizably coded in the activity of single neurons? Behavioral and Brain Sciences. 15:679-690.

Fetz EE, Cheney PD. 1980. Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. J Neurophysiol. 44:751-772.

Fetz EE, Finocchio DV. 1971. Operant conditioning of specific patterns of neural and muscular activity. Science. 174(4007):431-435.

Flash T. 1987. The control of hand equilibrium trajectories in multi-joint arm movements. Biol. Cyb. 57:257–274.

Flash T, Hogan N. 1985. The coordination of arm movements: an experimentally confirmed mathematical model. J Neurosci, 5(7):1688–1703.

Fraser GW, Schwartz AB. 2012. Recording from the same neurons chronically in motor cortex. J Neurophysiol. 107:1970-1978.

Fritsch G, Hitzig E. 1870. Electric excitability of the cerebrum. Arch Anat Physiol Wissen. 37:300-332.

Fu Q-G, Flament D, Coltz JD, Ebner TJ. 1995. Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. J Neurophysiol. 73:836-854.

Gallego JA, Perich MG, Miller LE, Solla SA. 2017. Neural manifolds for the control of movement. Neuron. 94:978-984.

Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE. Multiple tasks viewed from the neural manifold: Stable control of varied behavior. bioRxiv doi: https://doi.org/10.1101/176081.

Georgopoulos AP, Ashe J, Smyrnis N, Taira M. 1992. The motor cortex and the coding of force. Science. 256:1692-1695.

Georgopoulos AP, Caminiti R, Kalaska JF, Massey, JT. 1983. Spatial coding of movement: a hypothesis concerning the coding of movement direction by motor cortical populations. Exp Brain Res Suppl. 7:327–336.

Georgopoulos AP, Carpenter AF. 2015. Coding of movements in the motor cortex. Curr Opin Neurobiol. 33:34–39.

Georgopoulos AP, Crutcher MD, Schwartz AB. 1989. Cognitive spatial-motor processes 3. Motor cortical prediction of movement direction during an instructed delay period. Exp Brain Res. 75:183-194.

Georgopoulos AP, Kalaska JF, Caminiti R. 1985. Relations between two-dimensional arm movements and single-cell discharge in motor cortex and Area 5: Movement direction versus movement end point. Exp Brain Res., Suppl 10: 175-183.

Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. 1982. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci. 2(11):1527–1537.

Georgopoulos AP, Kalaska JF, Crutcher MD, Caminiti R, Massey JT. 1984. In: Edelman GM, Goll WE, Cowan WM, editors. Dynamic aspects of neocortical function. New York (NY): Neurosciences Research Foundation, Inc. p 501–524.

Georgopoulos AP, Kettner RE, Schwartz AB. 1988 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. J Neurosci. 8:2928–2937.

Georgopoulos AP, Lurito JT, Petrides M, Schwartz AB, Massey JT. 1989. Mental rotation of the neuronal population vector. Science. 243: 234-236.

Georgopoulos AP, Massey JT. 1987. Cognitive spatial-motor processes. 1. The making of movements at various angles from a stimulus direction. Exp Brain Res. 65:361-370.

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

Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Churchland MM, Kaufman MT, Kao JC, Ryu SI, Shenoy KV. A high-performance neural prosthesis enabled by control algorithm design. Nature Neuroscience. 15(12):1752-1757.

Griffin DM, Hoffman DS, Strick PL. 2015. Corticomotoneuronal cells are “functionally tuned.” Science. 350:667–670.

Hahnloser RHR, Kozhevnikov AA, Fee MS. 2002. An ultra-sparse code underlies the generation of neural sequences in a songbird. Nature. 419:65-70.

Harpaz NK, Ungarish D, Hatsopoulos NG, Flash T. 2018. Movement decomposition in the primary motor cortex. Cerebral Cortex. 2018, 1-5.

Hatsopoulos NG, Xu Q, Amit Y. 2007. Encoding of movement fragments in the motor cortex. J Neurosci. 27:5105–5114.

Heffner RS, Masterton RB. 1983. The role of the corticospinal tract in the evolution of human digital dexterity. Brain Behav. Evol. 23:165-183.

Hochberg LR. et al. 2006. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature. 442:164–171).

Hogan N. 1984. Adaptive control of mechanical impedance by coactivation of antagonist muscles. IEEE Trans. Aut. Control AC 29:681–690.

Hollerbach JMM. 1982. Computers, brains and the control of movement. TINS. 5:189–192.

Holmes G, May P. 1909. On the exact origin of the pyramidal tracts in man and other mammals. Brain. 32:1-43.

Inoue Y, Mao H, Suway SB, Orellana J, Schwartz AB. 2018. Decoding arm speed during reaching. Nature Communications. (2018)9:5243.

Inoue M, Uchimura M, Kitazawa S. 2016. Error signals in motor cortices drive adaptation in reaching. Neuron. 90:1114-1126.

Jackson A, Mavoori J, Fetz EE. 2007. Correlations between the same motor cortex cells and arm muscles during a trained task, free behavior, and natural sleep in the macaque monkey. J Neurophysiol. 97: 360–374.

Jankowska E, Padel Y, Tanaka R. 1975b. The mode of activation of pyramidal tract cells by intracortical stimuli. J. Physiol. 249:617-636.

Jarosiewicz B, et al. 2008. Functional network reorganization during learning in a brain-computer interface paradigm. Proc Natl Acad Sci USA. 105:19486–19491.

Johnson MTV, Ebner TJ. 2000. Processing of multiple kinematic signals in the cerebellum and motor cortices. Brain Res Rev. 33:155-168.

Johnson MTV, Mason CR, Ebner TJ. 2001. Central processes for the multiparametric control of arm movements in primates. Curr Opin Neurobiol. 11:684–688.

Kakei S, Hoffman DS, Strick PL. 1999. Muscle and movement representations in the primary motor cortex. Science. 285:2136–2139.

Kalaska JF, Caminiti R, Georgopoulos AP. 1983. Cortical mechanisms related to the direction of two-dimensional arm movments: relations in parietal area 5 and comparison with motor cortex. Exp Brain Res. 51:247–260.

Kalaska JF, Cohen DAD, Hyde ML, Prud'homme M. 1989. A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. J Neurosci. 9:2080-2102.

Kalaska JF, Scott SH, Cisek P, Sergio LE. 1997. Cortical control of reaching movements. Curr Op in Neurobio. 7:849-859.

Kass RE, Ventura V, Brown EN. 2005. Statistical issues in the analysis of neuronal data. J Neurophysiol. 94: 8–25.

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. The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. eNeuro. 3(4):1-25.

Keele SW, Posner MI. 1968. Processing of visual feedback in rapid movements. Journal of Experimental Psychology. 77(1):155-158.

Kemere C, Santhanam G, Yu BM, Afshar A, Ryu SI, Meng TH, Shenoy KV. 2008. Detecting Neural-State Transitions Using Hidden Markov Models for Motor Cortical Prostheses. J Neurophysiol. 100:2441-2452.

Kettner RE, Schwartz AB, Georgopoulos AP. 1988. Primate motor cortex and free arm movements to visual targets in three-dimensional space. III. positional gradients and population coding of movement direction from various movement origins. J Neurosci. 8(8):2938-2947.

Kim S-P, Simeral JD, Hochberg LR, Donoghue JP, Friehs GM, Black MJ. 2011. Point-and-Click Cursor Control With an Intracortical Neural Interface System by Humans With Tetraplegia. IEEE Transactions on Neural Systems and Rehabilitation Engineering. 19(2):193-203.

Krakauer JW, Ghez C, Ghilardi MF. 2005. Adaptation to Visuomotor Transformations: Consolidation, Interference, and Forgetting. J Neurosci. 25(2):473-478.

Krakauer JW, Pine ZM, Ghilardi M-F, Ghez C. 2000. Learning of visuomotor transformations for vectorial planning of reaching trajectories. J Neurosci. 20:8916-8924.

Kurata K, Hoshi E. 1999. Reacquisition deficits in prism adaptation after muscimol microinjection into the ventral premotor cortex of monkeys. J Neurophysiol. 81(4):1927-1938.

Kutz DF, Dannenberg W, Werner W, Hoffmann KP. 1997. Population coding of arm-movement-related neurons in and below the superior colliculus of Macaca mulatta. Biol Cybern. 76:331–337.

Lara AH, Elsayed GF, Zimnik AJ, Cunningham JP, Churchland MM. 2018. Conservation of preparatory neural events in monkey motor cortex regardless of how movement is initiated. eLife. 2018;7:e31826.

Lassek AM. 1942. The pyramidal tract: a study of retrograde degeneration in the monkey. Archives of Neurology and Psychiatry. 48(4):561-567.

Latimer KW, Yates JL, Meister MLR, Huk AC, Pillow JW. 2015. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. Science. 349:184-187.

Lawrence DG, Kuypers HGJM. 1965. Pyramidal and non-pyramidal pathways in monkeys: Anatomical and functional correlation. Science, 148:973-975.

Lemon RN, Porter R. 1976. Afferent input to movement-related precentral neurones in conscious monkeys. Proc. R. Soc. Lond. B. 194:313-339.

Levin PM. 1936. The efferent fibers of the frontal lobe of the monkey, macaca mulatta. The Journal of Comparative Neurology. 63(3):369-419.

Leyton SSF, Sherrington CS. 1917. Observations on the excitable cortex of the chimpanzee, orangutan and gorilla. Q. J. Exp. Physiol. 11:135-222.

Li C-S R, Padoa-Schioppa C, Bizzi E. 2001. Neuronal Correlates of Motor Performance and Motor Learning in the Primary Motor Cortex of Monkeys Adapting to an External Force Field. Neuron. 30:593-607.

Li Z, O'Doherty JE, Lebedev MA, Nicolelis MAL. 2011. Adaptive Decoding for Brain-Machine Interfaces Through Bayesian Parameter Updates. Neural Computation. 23:3162-3204.

Lurito JT, Georgakopoulos T, Georgopoulos AP. 1991. Cognitive spatial-motor processes. Exp Brain Res. 87:562-580.

Maier MA, Bennett KMB, Hepp-Reymond MC, Lemon RN. 1993. Contribution of the monkey corticomotoneuronal system to the control of force in precision grip. J Neurophysiol. 69: 772–785.

Mehta B, Schaal S. 2002. Forward models in visuomotor control. J Neurophysiol. 88:942-953.

Meyer DE, Abrams RA, Kornblum S, Wright CE, Smith JEK. 1988. Optimality in human motor performance: ideal control of rapid aimed movements. Psy Rev. 95:340–370.

Miall RC, Wolpert DM. 1996. Forward models for physiological motor control. Neural Networks. 9(8):1265-1279.

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

Michaels JA, 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(32):11415-11432.

Milner TE. 1992. A model for the generation of movements requiring endpoint precision. Neuroscience. 49:487–496.

Moran DW, Schwartz AB. 1999. Motor cortical activity during drawing movements: population representation during spiral tracing. J Neurophysiol. 82:2693–2704.

Moran DW, Schwartz AB. 1999. Motor cortical representation of speed and direction during reaching. J Neurophysiol. 82:2676-2692.

Morasso P. 1981. Spatial control of arm movements. Exp Brain Res. 42:223–227.

Morrow MM, Miller LE. 2003. Prediction of Muscle Activity by Populations of Sequentially Recorded Primary Motor Cortex Neurons. J Neurophysiol. 88:2279-2288.

Novak, Miller, Houk. 2002. The use of overlapping submovements in the control of rapid hand movements. Exp. Brain. Res. 144:351-364.

Okatan M, Wilson MA, Brown EN. 2005. Analyzing Functional Connectivity Using a Network Likelihood Model of Ensemble Neural Spiking Activity. Neural Computation. 17:1927-1961.

Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP. 2004. Spatiotemporal tuning of motor cortical neurons for hand position and velocity. J Neurophysiol. 19:515-232.

Paillard J. 1982. The contribution of peripheral and central vision to visually guided reaching. In: Ingle DJ, Goodale MA, Mansfield RJW, editors. Analysis of visual behavior. Cambridge: MIT Press. p 367–385.

Paillard J. 1996. Fast and slow feedback loops for the visual correction of spatial errors in a pointing task: a reappraisal. Can. J. Physiol. Pharmacol. 74:401-417.

Pellizzer G, Georgopoulos AP. 1993. Common processing constraints for visuomotor and visual mental rotations. Exp Brain Res. 93:165-172.

Pellizzer G, Sargent P, Georgopoulos AP. 1995 Motor cortical activity in a context-recall task. Science. 269:702-705.

Penfield W, Boldrey E. 1937. Somatic motor and sensory representation in the cerebral cortex of man as studies by electrical stimulation. Brain, 60:389-443.

Penfield W, Rasmussen T. 1950. The Cerebral Cortex of Man; a Clinical Study of Localization of Function. Macmillan Co., New York. 248.

Pesaran B, Nelson MJ, Andersen RA. 2008. Free choice activates a decision circuit between frontal and parietal cortex. Nature. 453:406-409.

Phillips CG, Porter R. 1997. Corticospinal neurons. London: Academic Press.

Polit A, Bizzi E. 1978. Processes controlling arm movements in monkeys. Science. 29:1235-1237.

Porter R, Lemon R. 1993. Corticospinal function and voluntary movement. Oxford: Clarendon Press.

Rathelot J-A, Strick PL. 2009. Subdivisions of primary motor cortex based on cortico-motoneuronal cells. Proc Natl Acad Sci USA. 106:918-923.

Reina GA, Moran DW, Schwartz AB. 2001. On the Relationship Between Joint Angular Velocity and Motor Cortical Discharge During Reaching. J Neurophysiol. 85(6):2576-2589,

Rouse AG. 2018. Cyclic, condition-independent activity in primary motor cortex predicts corrective movement behavior. bioRxiv doi https://doi.org/10.1101/453746.

Rouse AG, Schieber MH. 2016b. Spatiotemporal distribution of location and object effects in primary motor cortex neurons during reach-to-grasp. J Neurosci. 36(41):10640-10653.

Rouse AG, Schierber MH. 2018. Condition-Dependent Neural Dimensions Progressively Shift during Reach to Grasp. Cell Reports. 25:3158-3168.

Russo AA, Bittner SR, Perkins SM, Seely JS, London BM, Lara AH, Miri A, Marshall NJ, Kohn A, Jessell TM, Abbott LF, Cunningham JP, Churchland MM. 2018. Motor cortex embeds muscle-like commands in an untangled population response. Neuron. 97:953-966.

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

Salinas E, Abbott LF. 1994. Vector reconstruction from firing rates. J Comput Neurosci. 1:89–107.

Sauerbrei B, Guo J-Z, Mischiati M, Guo W, Kabra M, Verma N, Branson K, Hantman A. 2018. Motor cortex is an input-driven dynamical system controlling dexterous movement. bioRxiv doi: https://doi.org/10.1101/266320.

Schieber MH. 2011. Dissociating motor cortex from the motor. J Physiol. 589:5613–5624.

Schieber MH, Rivlis G. 2005. A spectrum from pure post-spike effects to synchrony effects in spike-triggered averages of electromyographic activity during skilled finger movements. J Neurophysiol. 94:3325-3341.

Schwartz AB. 1992. Motor cortical activity during drawing movements: single-unit activity during sinusoid tracing. J Neurophysiol. 68(2):528-541.

Schwartz AB. 1993. Motor cortical activity during drawing movements: population response during sinusoid tracing. J Neurophysiol. 70:28–36.

Schwartz AB. 1994. Direct cortical representation of drawing. Science. 265:540–542.

Schwartz AB. 1994. Neuronal substrate for volitional movement. In: Bennett KMB, Castiello U, editors. Insights into the reach to grasp movement. Amsterdam (North Holland): Elsevier. p 59–83.

Schwartz AB, Kettner RE, Georgopoulos AP. 1988. Primate Motor Cortex and Free Arm Movements to Visual Targets in Three-Dimensional Space. I. Relations Between Single Cell Discharge and Direction of Movement. J. Neurosci. 8(8):2913-2927.

Schwartz AB, Moran DW. 1999. Motor cortical activity during drawing movements: population representation during lemniscate tracing. J Neurophysiol. 82:2705–2718.

Schwartz AB, Moran DW, Reina A. 2004. Differential representation of perception and action in the frontal cortex. Science. 303:380-383.

Scott SH, Kalaska JF. 1995. Changes in motor cortex activity during reaching movements with similar hand paths but different postures. J Neurophysiol. 73:2563–2567.

Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E. 1996. Simultaneously Recorded Single Units in the Frontal Cortex Go through Sequences of Discrete and Stable States in Monkeys Performing a Delayed Localization Task. J Neurosci. 16(2):752-768.

Sergio LE, Hamel-Pâquet C, Kalaska JF. 2005. Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. J Neurophysiol. 94:2353–2378.

Sergio LE, Kalaska JF. 1998. Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task. J Neurophysiol. 80:1577–1583.

Sergio LE, Kalaska JF. 2003. Systematic changes in motor cortex cell activity with arm posture during directional isometric force generation. J Neurophysiol. 89:212–228.

Serruya MD, Hatsopoulos NG, Paninski L, Fellows MR, Donoghue JP. 2002. Instant neural control of a movement signal. Nature. 416:141–142.

Shadmehr R, Mussa-Ivaldi FA. 1994. Adaptive representation of dynamics during learning of a motor task. J Neurosci. 14:3208-3224.

Shen L, Alexander GE. 1997a. Neural correlates of a spatial sensory-to-motor transformation in primary motor cortex. J Neurophysiol. 77:1171-1194.

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

Shmiel T, Drori R, Shmiel O, Ben-Shaul Y, Nadasdy Z, Shemesh M, Teicher M, Abeles M. 2006. Temporally Precise Cortical Firing Patterns Are Associated With Distinct Action Segments. J Neurophysiol. 96:2645-2652.

Soechting JF, Lacquaniti F. 1981. Invariant characteristics of a pointing movement in man. J Neurosci. 1:710–720.

Stavisky SD, Kao JC, Ryu SI, Shenoy KV. 2017. Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neural state space dimensions. Neuron. 95:1-14.

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

Sussillo D, 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.

Suway SB, Orellana J, McMorland AJC, Fraser GW, Liu Z, Velliste M, Chase SM, Kass RE, Schwartz AB. 2017. Temporally segmented directionality in the motor cortex. Cerebral Cortex. 2017, 1-14.

Suway SB, Tien RN, Jeffries SM, Zohny Z, Clanton ST, McMorland AJC, Velliste M. 2013. Resting State Detection for Gating Movement of a Neural Prosthesis. 6th Annual International IEEE EMBS Conference on Neural Engineering, San Diego, California, 6 - 8 November, 2013. 665-668.

Tanji J, Evarts EV. 1976. Anticipatory Activity of Motor Cortex Neurons in Relation to Direction of an Intended Movement. J Neurophysiol. 39(5):1062-1068.

Taylor DM, Helms Tillery SI, Schwartz AB. 2002. Direct cortical control of 3D neuroprosthetic devices. Science. 296:1829–1832.

Thach WT. 1978. Correlation of Neural Discharge with Pattern and Force of Muscular Activity, Joint Position, and Direction of Intended Next Movement in Motor Cortex and Cerebellum. J. Neurophysiol. 41(3):654-676.

Tseng Y-W, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ. Sensory Prediction Errors Drive Cerebellum-Dependent Adaptation of Reaching. J Neurophysiol. 98:54-62.

Turner RS, Anderson ME. 1997. Pallidal discharge related to the kinematics of reaching movements in two dimensions. J Neurophysiol. 77:1051–1074.

Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, Slovin H, Aertsen A. Dynamics of neuronal interactions in monkey cortex in relation to behavioral events. Nature. 373:515-518.

van Hemmen JL, Schwartz AB. 2008. Population vector code: a geometric universal as actuator. Biol Cybern. 98:509–518.

Vargas-Irwin CE, Franquemont L, Black MJ, Donoghue JP. 2015. Linking Objects to Actions: Encoding of Target Object and Grasping Strategy in Primate Ventral Premotor Cortex. J Neurosci. 35(30):10888-10897.

Velliste M, Kennedy SD, Schwartz AB, Whitford AS, Sohn J-W, McMordland AJC. 2014. Motor cortical correlates of arm resting in the context of a reaching task and implications for prosthetic control. J Neurosci. 34:6011-6022.

Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. 2008. Cortical control of a prosthetic arm for self-feeding. Nature 453:1098–1101.

Wang W, Chan SS, Heldman DA, Moran DW. 2007. Motor cortical representation of position and velocity during reaching. J Neuorphysiol. 97:4258-4270.

Waters P, Strick PL. 1981. Influence of 'strategy' on muscle activity during ballistic movements. Brain Research. 207:189-194.

Wessberg J et al. 2000. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. Nature. 408:361–365.

Wiesendanger M. 1969. The pyramidal tract recent investigations on its morphology and function. Ergeb. Physiol. 61:73-136.

Wise SP, Moody SL, Blomstrom KJ, Mitz AR. 1998. Changes in motor cortical activity during visuomotor adaptation. Exp Brain Res. 121:285-299.

Wodlinger B, et al. 2015. 10 dimensional anthropomorphic arm control in a human brain-machine interface: difficulties, solutions, and limitations. J Neural Eng. 12:1-17.

Wolpaw JR. 1980. Correlations between task-related activity and responses to perturbation in primate sensorimotor cortex. J Neurophysiol. 44(6):1122-1138.

Wolpert DM, Miall RC, Kawato M. 1998. Internal models in the cerebellum. Trends in Cognitive Science. 2(9):338-347.

Woodworth RS. 1899. The accuracy of voluntary movement. The Psychological Review: Monograph Supplements, 3(3), i-114.

Woolsey CN. 1958. Organization of somatic sensory and motor areas of the cerebral cortex. In Biological and Biochemical Bases of Behavior, ed. Harlow HF & Woolsey CN. University of Wisconsin Press, Madison. pp. 63–81.

Yin P-B, Kitazawa S. 2001. Long-lasting aftereffects of prism adaptation in the monkey. Exp. Brain. Res. 141:250–253.

Zelaznik HN, Hawkins B, Kisselburgh L. 1983. Rapid visual feedback processing in single-aiming movements. Journal of Motor Behavior. 15(3):217-236.

Zhang J, Riehle A, Requin J, Kornblum S. 1997. Dynamics of single neuron activity in monkey primary motor cortex related to sensorimotor transformation. J Neurosci. 17(6): 2227-2246.

Zhang K, Sejnowski TJ. 1999. A theory of geometric constraints on neural activity for natural three-dimensional movement. J Neurosci. 19: 3122–45.