Article

Cell Reports

Activity in Primary Motor Cortex Related to Visual Feedback

Graphical Abstract

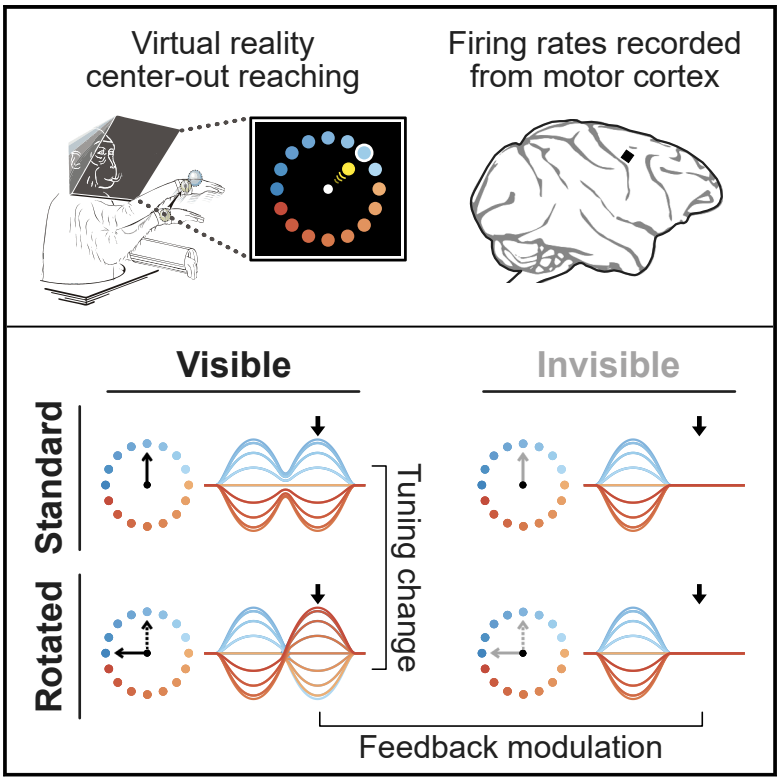

Highlights

- Single-unit encoding discretely changes between activity segments
- Effects of visuomotor adaptation on encoding differ between activity segments
- Late-movement activity segments can be weak or absent if visual feedback is absent
- If feedback is present, feedback-sensitive segments are tuned in visual coordinates

Authors

Steven B. Suway, Andrew B. Schwartz

Correspondence

abs21@pitt.edu

In Brief

Suway and Schwartz show that M1 activity is temporally segmented during reaching, visuomotor adaptation can selectively modulate early-movement or late-movement activity segments, and visual feedback selectively modulates late-movement segments.

Suway & Schwartz, 2019, Cell Reports *29*, 3872–3884
December 17, 2019
https://doi.org/10.1016/j.celrep.2019.11.069

CellPress

OPEN ACCESS
CellPress

Cell Reports
Article

Activity in Primary Motor Cortex Related to Visual Feedback

Steven B. Suway[1,2] and Andrew B. Schwartz[2,3,4,5,*]
[1]Center for Neuroscience, University of Pittsburgh, Pittsburgh, PA 15260, USA
[2]Center for the Neural Basis of Cognition, Carnegie Mellon University and University of Pittsburgh, Pittsburgh, PA 15213, USA
[3]Systems Neuroscience Center, University of Pittsburgh, Pittsburgh, PA 15260, USA
[4]Department of Neurobiology, University of Pittsburgh, Pittsburgh, PA 15260, USA
[5]Lead Contact
*Correspondence: abs21@pitt.edu
https://doi.org/10.1016/j.celrep.2019.11.069

SUMMARY

Neural modulation in primate motor cortex exhibits complex patterns. We found that modulation during reaching could be separated into discrete temporal epochs. To determine if these epochs are driven by behavioral events, monkeys performed variations of a center-out reaching task. Monkeys viewed a computer cursor matched to hand position and a radial target at 1 of 16 locations. In some trials, they performed a visuomotor rotation (the cursor moved at an angle to the hand). After adaptation, encoding changes for single units are temporally structured: adaptation could affect one temporal component of a unit's response but not another. In half the normal and perturbed trials, we removed visual feedback before movement. Adaptation-sensitive firing components toward the end of movement are often weak or absent during reaches without feedback. These results show that temporal structure in motor cortical activity is driven by behavior, with a discrete component related to visual feedback.

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 and Shenoy, 2007). Furthermore, 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 et al., 1999; Caminiti et al., 1990; Zhang et al., 1997). Investigators have attempted to improve these simple representational models by including additional features of movement, but gains in explanatory power have been modest (Paninski et al., 2004; Aflalo and Graziano, 2006; Suway et al., 2017). Although the early directional tuning work (Georgopoulos et al., 1982; Schwartz et al., 1988) used firing rates averaged over the entire movement, with the implicit assumption of stationary preferred directions (PDs), a number of studies have revealed time-dependent processing in M1 during reaching (Georgopoulos et al., 1983, 1989; Pellizzer et al., 1995; Fu et al., 1995; Zhang et al., 1997; Churchland et al., 2012; Velliste et al., 2014; Rouse and Schieber, 2016). Recent work (Suway et al., 2017; Harpaz et al., 2019) 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., 2019). 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 and 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 while the moving hand is registered in the peripheral visual field (Paillard, 1982, 1996). Accurate target acquisition takes place as a series of visually guided sub-movements as the hand's image enters the macular retina (Soechting and Lacquaniti, 1981; Meyer et al., 1988; Milner, 1992; Novak et al., 2002). Experiments dissociating vision from movement have provided evidence for visuospatial coding

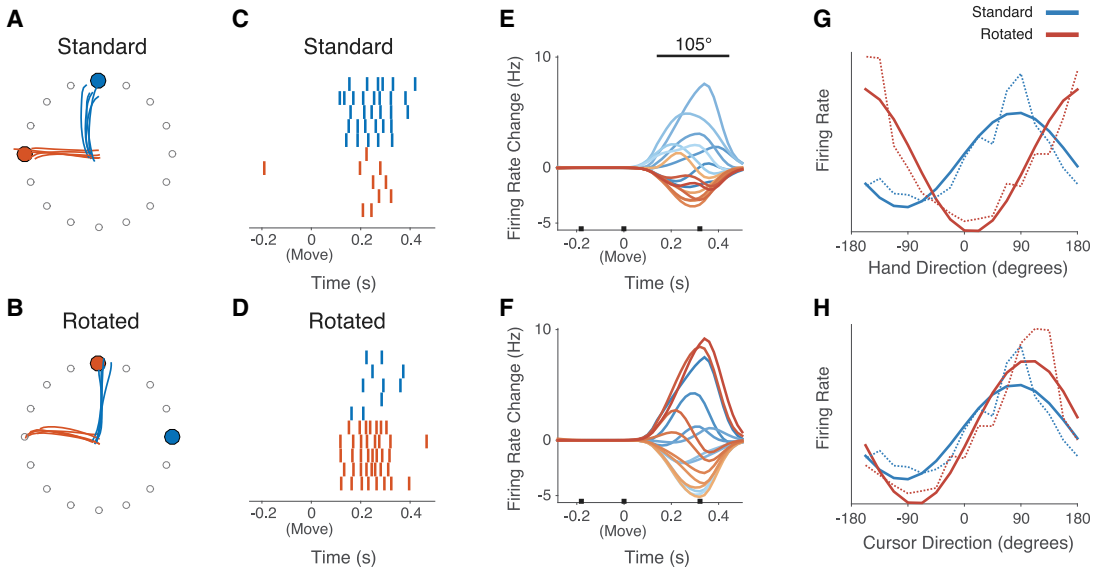

Figure 1. Firing Rate Properties Change after Visual Adaptation

(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. 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 presentation, 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 (Equation 1 in STAR Methods), and dashed lines are observed rates. Number over black bar in (E) reports the angular PD difference between cosine fits.

(H) Same as in (G) but with tuning expressed relative to cursor direction rather than hand direction.

The PDs of the responses in (E) and (F) were very similar when expressed in cursor-centric coordinates.

in M1 and premotor cortex, particularly before the movement begins (Georgopoulos et al., 1989; Pellizzer et al., 1995; Zhang et al., 1997; di Pellegrino and Wise, 1993). It is also well known that M1 is active during passive observation of movement, for example, in "movement replay" tasks (Taylor et al., 2002; Tkach et al., 2007; Palazzolo, 2015). This result is also an important aspect of the brain-computer interface paradigm used for neural prosthetics (Velliste et al., 2008). However, less is known about the relation between M1 activity and ongoing visual feedback during movement (although see Schwartz et al., 2004). One possibility is that the sequence of neuronal states is determined, in part, by visual information related to the task. To test whether discrete components of the M1 response during reaching might be related to visual feedback, we implemented several variations of the classic 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 reach visibility 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. This component of the response 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 and Posner, 1968). Our findings suggest there is a discrete temporal component of M1 activity related to visual feedback.

RESULTS

Single-Unit Directional Encoding Changes after Visuomotor Adaptation

We used the property of directional tuning of neurons in M1 to examine possible effects of visuomotor adaptation on neural firing. Monkeys made reaches to the same physical targets before and after adaptation (Figures 1A and 1B). An example unit recorded during those trials was found to fire strongly during upward reaches (Figure 1C, blue rasters). If this tuning property were 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.

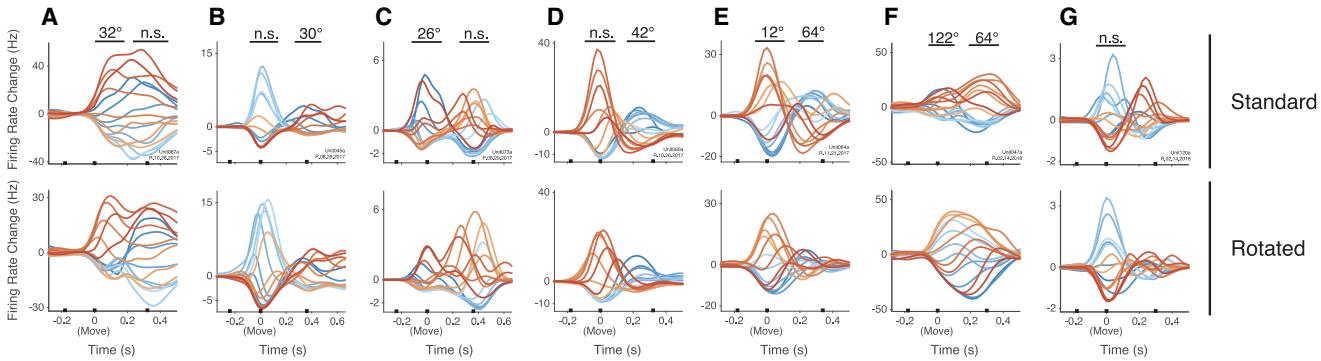

Figure 2. Components of Single Unit Responses Are Independently Affected by Visuomotor Adaptation

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 (A)–(D) were from 45° sessions, the response in (E) was from a 67.5° session, and the responses in (F) and (G) were from 90° sessions. Black bars above response components mark time periods during which we applied the same tuning comparison method as in Figures 1E–1G (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., no significant difference ($p \geq 0.05$). See also Figure S7.

However, during the rotated block, this neuron fired at a much lower rate for upward reaches (Figure 1D, blue rasters). One possibility is that this neuron was directionally tuned in virtual reality (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 did not physically reach in that direction. This is what was observed for the example response in Figure 1. The perturbation was 90° clockwise, such that the upward target cued a reach to the left (Figure 1B, red traces and red circle). Leftward reach trials had rates that were similar to those for upward reaches in the standard block (Figures 1C and 1D, red versus blue rasters), suggesting that cursor-related information was encoded in this neuron's activity.

We next examined the firing rate profile of this neuron across targets and found the profile had a similar shape and timing in both blocks (Figures 1E and 1F, different colored traces). Note that the color scheme used for these plots was defined relative to physical movement direction in both blocks. Encoding properties previously seen for the raster plots can therefore be recognized qualitatively in the ordering of rates across targets in Figures 1E and 1F (e.g., light blue versus dark red traces). Across-target rates averaged over the period marked in Figure 1E (black bar) were fit well by cosine functions in each task block (Figure 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 1.

Using a bootstrap procedure, we evaluated the across-trial consistency of the PD change (Figure 1G) by resampling individual trials with replacement and re-computing PD change for each bootstrap iteration. The 95% confidence interval (CI) was [96°, 113°], meaning that the PD change was consistent over trials within about 17°, suggesting that the effect demonstrated in Figures 1C–1G 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 (Figure 1H). The similar phase of the cosine functions in Figure 1H 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, again supporting the idea that this neuron's tuning was driven by cursor-related information. In subsequent figures, firing rates and tuning curves are plotted with respect to physical movement.

The activity displayed in Figure 1 had only a single modulation peak. However, M1 neurons often fire with multiple modulation peaks during reaching (Churchland and 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 often showed a variety of patterns following adaptation (Figure 2). Some patterns displayed an early modulation component related to the adaptation, while a subsequent component remained unchanged (Figures 2A and 2C). Others showed the opposite temporal pattern (Figures 2B and 2D). Both components in Figures 2E and 2F were affected by adaptation, to different degrees. Figure 2G shows a response with an early component unaffected by adaptation and a second with a highly reduced amplitude after adaptation. These observations support the hypothesis that distinct peaks of modulation correspond to distinct encoding epochs. We also found responses with no relation to the adaptation (not shown), indicating a constant encoding of the physical movement direction. Although the example response shown in Figure 1 had a tuning change of similar magnitude to the visuomotor perturbation, the examples in Figure 2 indicate that the magnitude of change was variable across responses.

Adaptation Effects Are Widespread and Depend on Rotation Magnitude

Because the effect of adaptation had a strong time-dependent effect on single-neuron firing rates (Figure 2), we segmented trials into four epochs using factor analysis (FA; see STAR Methods). Previously, our group identified three similar epochs

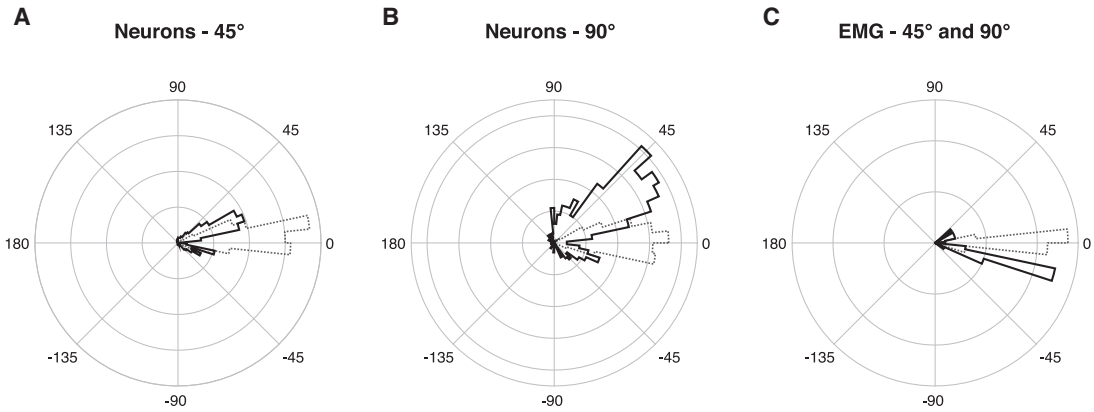

Figure 3. Neuronal PD Changes Depend on Rotation Magnitude, but Muscle PDs Are Similar after Adaptation

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. See also Figures S2 and S3.

on the basis of the timing of modulation peaks of single neurons (Suway et al., 2017). An additional epoch in the current FA procedure targeted activity during the reaction time and early movement and was included for greater temporal resolution. The four epochs corresponded to similar behavioral events in both monkeys; the first two epochs for monkey P were slightly shifted forward in time compared with monkey R (Figure S2A versus Figure S2B). The shift could 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 retained responses well fit by cosine functions ($R^2 > 0.5$) in both the standard and rotated blocks. For both monkeys, the proportion of neurons meeting this criterion was ~60% in the first epoch, ~50% in the second, and 30%–40% in the last two. These values are similar to those previously reported (Suway et al., 2017). We computed PD differences between the two blocks for each neuron and epoch (Figure 3; black traces indicate statistically 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 3A versus Figure 3B). This property can also be seen for single-unit responses (Figures 2A–2D versus Figures 2E and 2F). The proportion of neurons with a significant PD change (regardless of magnitude) after adaptation was also larger for 90° sessions (Figures 3A and 3B, black versus gray histograms). The majority of PD differences were in the positive direction, like the responses shown in Figures 1 and 2.

For responses following this pattern, the most common observation was a PD change of roughly half the perturbation angle. To test the possibility that intermediate-valued changes may reflect an incomplete adaptation process, 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, those changes should better match the perturbation angle following additional exposure (i.e., during the second half of the block). However, the distributions of changes were similar to those in Figure 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 et al., 1991; Shen and Alexander, 1997a), rather than an incomplete adaptation process.

To compare M1 responses to the concurrent motor output, we repeated this analysis for electromyography (EMG) activity (Figure 3C). The EMG factors corresponded to three clear epochs around movement onset, peak velocity, and movement offset. A fourth factor had small loadings and contained several noisy peaks, suggesting that three factors were sufficient for describing these data. PDs of EMG data computed within the first three EMG epochs were not substantially different from those computed within the four epochs defined by the neuronal data. For a more direct EMG comparison with the neuronal data, we used the neuronal epochs. The PDs computed from EMG activity tended to be similar between the two task blocks, though some small differences were observed 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, and (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 magnitude similar to those observed for EMG. This subpopulation may be considered "muscle-like" (Kakei et al., 1999), while responses with a

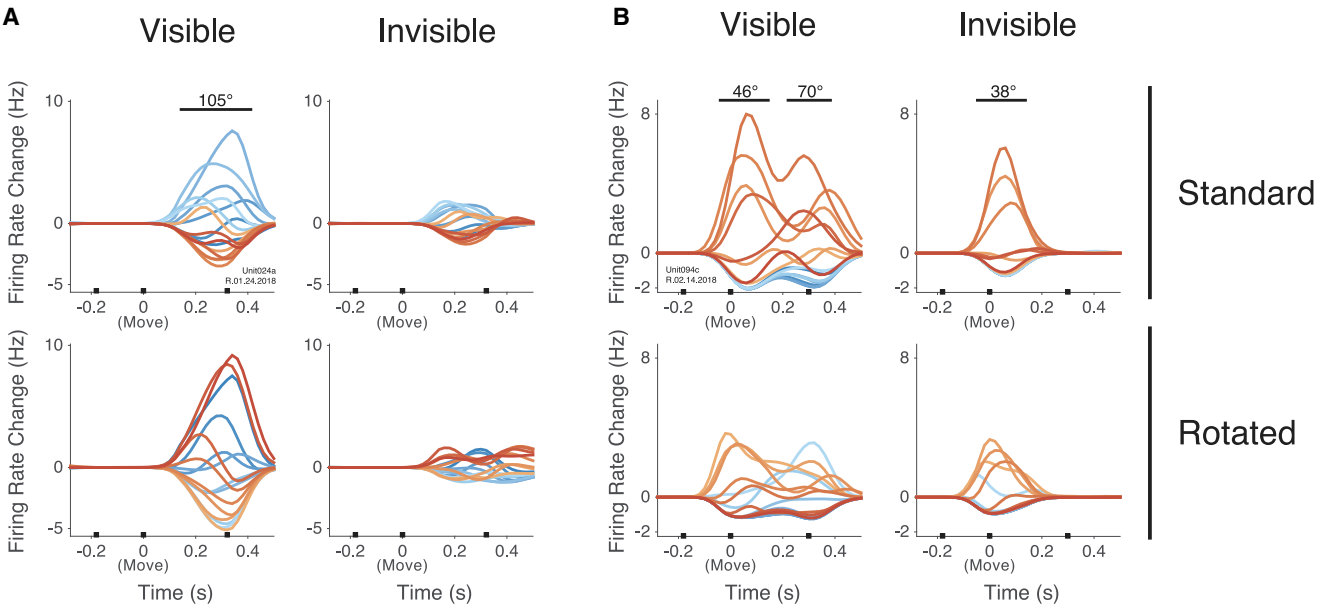

Figure 4. Reach Visibility Affects Neuronal Firing

(A) Firing rates of the same unit shown in Figure 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.

See also Figure S7.

positive PD change between blocks may be related to 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 were partitioned into two groups, and the PD difference between them was computed for each trial epoch. We then repeated the procedure 1,000 times with new randomly chosen partitions and displayed the PD differences as angular histograms (Figure S3). We found that these differences were always near zero, suggesting the larger changes seen between blocks (Figure 3) were related to the task.

Features of M1 Responses Are Related to Reach Visibility

The effects of visuomotor adaptation on the firing patterns of M1 neurons may be 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 4 shows firing rates of two 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 for visible trials, and the right column shows those for the invisible trials. For the response in Figure 4A, during visible trials, the PD changed after the adaptation (left column). This neuron's firing decreased greatly when reaches were not visible (right column). A different unit's response (Figure 4B) had two modulation peaks during visible trials (left column). Both peaks were affected by the visual perturbation, although the second peak was affected to a greater degree (46° versus 70°). Interestingly, the second modulation peak was absent when the cursor was invisible (right column). The first peak was present regardless of cursor visibility, with a similar PD change following adaptation (38° versus 46°; black bars above peaks).

The responses in Figure 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. To determine whether these findings applied to the sampled population, we again used the FA trial segmentation procedure. Within each of the four 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 5A (blue lines, solid and dashed, respectively). The factor loadings marking the four epochs are also displayed for reference (black traces). If the visibility-related firing patterns shown in Figure 4 were common across the population, we would expect to find more neurons with significant PD changes between blocks for visible trials compared with 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

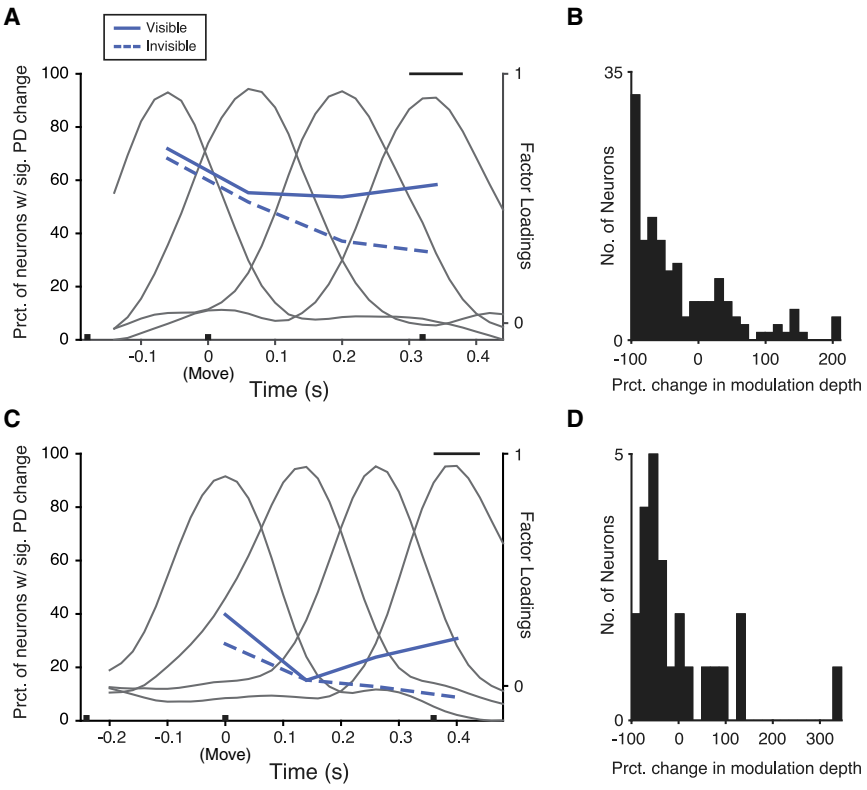

Figure 5. Visibility-Related Tuning Changes Occur toward the End of Movement

(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 A). Histogram includes responses with a positive significant PD change between blocks in the visible condition. Data in (B) 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.

See also Figures S4 and S6.

around peak movement velocity and the end 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 it could not. Activity related to the perturbation but not to visual feedback (i.e., during epochs 1 and 2) may be a hallmark of the adaptation process. Activity related to both task parameters, occurring toward the end of behavior (i.e., epochs 3 and 4), may be consistent with a feedback signal.

The example responses in Figure 4 showed decreased modulation amplitude around the end of the movement when reaches were not visible, and this decrease might explain why fewer PD changes were observed toward the end of these trials. To address this possibility, we focused on the fourth epoch of visible trials and selected neurons with positive PD changes after adaptation. We then computed the percentage change in modulation depth between visible and invisible trials (Figure 5B). Tuning amplitude commonly decreased in the invisible trials, with the firing of many units showing a nearly complete loss of modulation. This pattern is similar to the responses shown in Figure 4.

The results in Figures 5A and 5B 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 5C), though the overall percentage of PD changes after adaptation was smaller (see also Figure 3A versus Figure 3B). Changes in modulation depth were also similar between monkeys: neurons with positive PD changes in epoch four of visible trials commonly showed lower tuning modulation during invisible trials (Figure 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 vision is absent. However, we also found other patterns. From Figure 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 Figure S4.

Decoding Cursor Direction Captures Activity Related to Reach Visibility

Next, we investigated whether the displayed (VR) movement direction could be extracted or decoded from the population. Reverse regression was used to build a decoder using linear combinations of the neural data that captured the displayed movement direction (see STAR 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 weighted responses with this property more heavily than those without. This procedure yielded two linear combinations of rates: one for decoding the x-component and another for the y-component of the displayed direction.

We identified these decoder weights using only the visible trials. Figures 6A and 6B show 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 6A, blue traces), and negative for downward reaches (red and orange traces). This is the expected pattern for a representation of the y-component. In the rotated condition, the activation pattern changed (Figure 6B). Note that the color scheme of the 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. Figure 6 shows data from a 90°

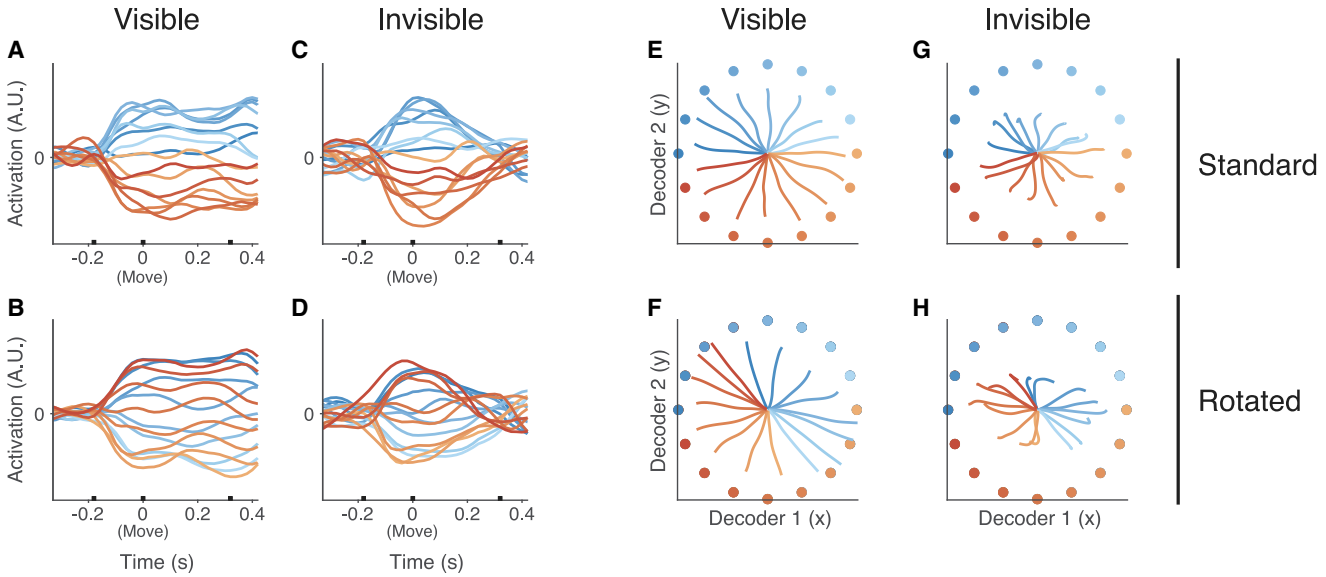

Figure 6. Accurate Decoder Readout of Displayed Direction Depends on Cursor Visibility

(A) Decoder output for the y-component of displayed direction 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 that the ordering across targets has changed, though the decoder weights were the same.

(C and D) Same as in (A) and (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). Trajectories during the standard block are shown 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 that the trajectories have rotated 90° clockwise in this block, though the physical targets are the same.

(G and H) Same as in (E) and (F) but for invisible trials. Trajectories during invisible trials were shortened and distorted toward the end of movement, relative to visible trials.

session; in the rotated block, leftward physical reaches were associated with upward displayed movement. Because 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 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 Figures 6A and 6B, the decoded components were nearly orthogonal, matching the y-component of displayed movement direction closely in both blocks: the 5-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.

Next, the rates recorded during invisible trials were used in the decoder constructed from the visible trials (Figures 6C and 6D). For those invisible trials, the decoder result was comparable with that for visible trials shortly after the target appeared, during the early part of the movement. Over that time span, the across-target pattern and amplitude were comparable regardless of visibility. However, later in the invisible trials, decoder output decreased in amplitude, and the ordering across targets became less clear (Figures 6C and 6D). This pattern matches closely with that described for the tuning analysis of single-unit responses in Figure 5.

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 (Figures 6E–6H). This integration procedure is similar to that used in the classical population vector algorithm (Schwartz, 1994; Schwartz and Moran, 1999). In the standard block, the trajectories were relatively straight and accurate (Figure 6E). A similar result was found during the rotated block (Figure 6F), in which the trajectories were rotated nearly 90°, closely matching the displayed direction. The same procedure for invisible trials yielded trajectories that initially matched those of the visible trials (Figures 6G and 6H). However, in both the standard and rotated blocks, the trajectories did not extend as far and became distorted at the end of the task.

We next quantified this temporal structure in the decoder output. Averaged over epochs one and two, y-component decoder output across targets closely fit a sine function, peaking near 90° (upward) in the standard block (Figure 7A, blue traces). The sinusoid shifted phase during the rotated block, becoming centered near 180° (leftward; Figure 7A, red traces). Using a bootstrap procedure, we found this phase difference to be consistent over trials (95% CI = [78°, 92°]; Figure 7B). In contrast, during epoch four, across-target activation was poorly fit by sinusoidal functions and was inconsistent over trials (Figures 7C and 7D). This result reflects the single-neuron analyses presented in Figure 5: late in the movement, a component of the response related to the perturbation was less prominent if the movement was hidden from view.

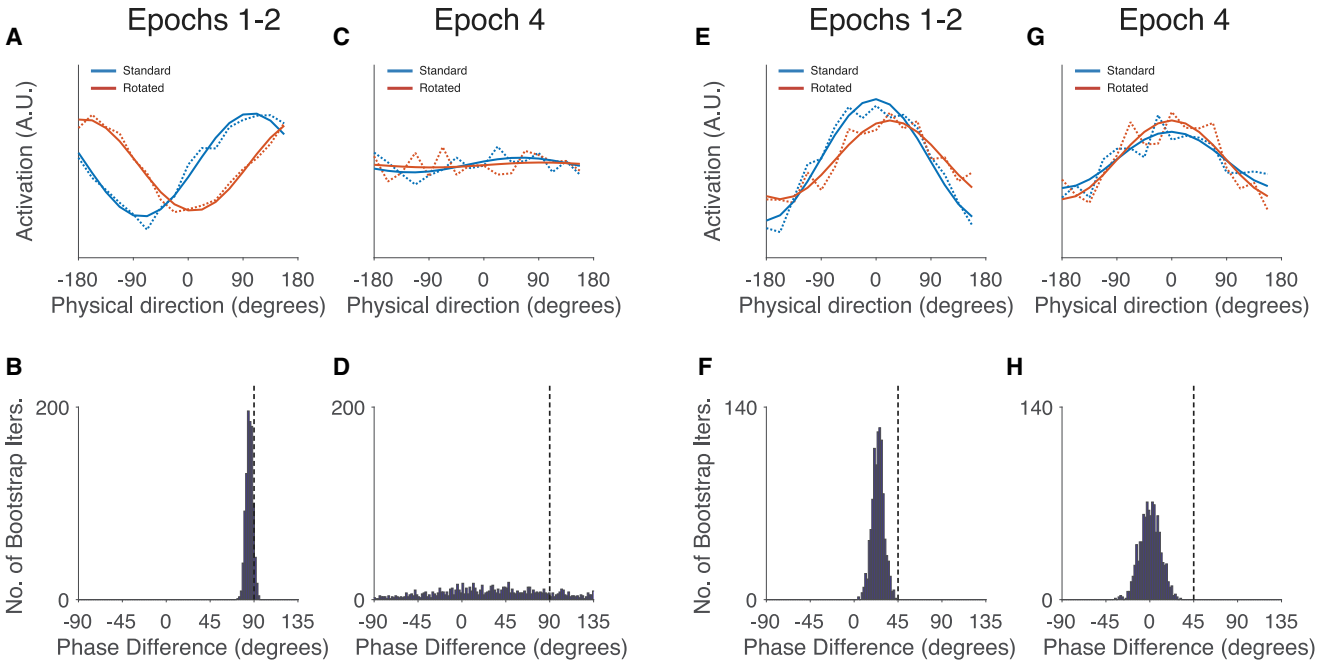

Figure 7. Decoder Output Is Inconsistent Late in the Movement for Invisible Trials

(A) Same decoder output shown in Figures 6C and 6D, averaged over epochs 1 and 2, and plotted against physical reach direction. Blue traces are from the standard block (i.e., Figure 6C), red traces from the rotated block (i.e., Figure 6D). Dashed traces show observed data, solid traces show model fits (Equation 1 in STAR Methods).

(B) Phase difference between the two model fits in (A), shown over bootstrap iterations. The difference was consistently near 90°.

(C and D) Same as in (A) and (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). Readout for the x-component of displayed direction is shown.

Data in Figures 7A–7D are from a 90° session from monkey R. For monkey P, we recorded data only for 45° sessions; the physical and displayed movements were not perfectly orthogonal. The correlation between them ($\rho \approx 0.7$) likely causes decoder output to represent a mixture of activity related to displayed and executed movement. Nonetheless, we found similar trends for monkey P: averaged over epochs 1 and 2, the x-component decoder output was well fit by a sinusoid centered near 0°(rightward) in the standard block (Figure 7E, blue traces), which shifted 25° during the rotated block (95% CI = [12°, 37°]; Figures 7E and 7F). During the fourth epoch, no phase change was observed (95% CI = [−19°, 21°]; Figures 7G and 7H). Decoder output in epoch four was slightly noisier and weaker than for epochs 1 and 2. Compared with the results for the 90° session, these effects were more moderate (Figures 7C and 7D versus Figures 7G and 7H). A similar result was observed for a 45° session from monkey R, including a slight undershoot of the rotation angle between blocks for the early epochs.

Physical Reach Direction Is Well Represented Regardless of Task Condition

Although many neurons responded in different ways before and after adaptation, and during 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 described above but used the reverse regression procedure to decode the physical reach direction, rather than the displayed reach direction. Units were included regardless of whether their responses had adaptation-related tuning changes. Regression was again performed using only visible trials. Accuracy was relatively high: the cross-validated R^2 values for the y-component were 0.89 and 0.93, with 0.87 and 0.93 for the x-component in the standard and rotated blocks, respectively. Moreover, decoder output was similar both before and after adaptation. This is the expected pattern for a decoder that predicts physical reach direction, which was identical in each block. Interestingly, output in both blocks was relatively consistent even during invisible trials (albeit slightly weaker for invisible trials in the standard block). These results suggest that the physical movement was consistently represented in M1 activity regardless of visual context, in line with the finding that a considerable portion of responses did not change PD in the rotated block (Figures 3A and 3B).

Kinematics and Muscle Activity Were Generally Comparable across Task Conditions

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 possible pair of conditions (Shadmehr and Mussa-Ivaldi, 1994). Pairwise correlations for all sessions were very high, ranging from 0.96 to 0.99. In general, patterns

of muscle activation were also comparable across conditions (Figure S5). From Figure 3C, we would expect the EMG data to have similar tuning properties, relative to kinematic direction, before and after adaptation. Indeed, the across-target ordering of activity was comparable in both blocks (Figure S5, top versus bottom rows). An exception can be seen in the antagonist burst of the medial deltoid.

EMG differences between visible and invisible trials were also generally subtle (Figure S5, left versus right columns in each panel). One important exception was the second agonist burst for some muscles. For example, triceps activity around movement offset was weaker during invisible trials than during visible trials (Figure S5C). This can be seen to a lesser degree in the medial deltoid (Figure S5A) and wrist flexor (Figure S5B). One possibility is that the later EMG activity during visible trials is related to online error correction or trajectory stabilization. Reaching movements are composed of a primary arm displacement 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 neuronal modulation 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 3C). Furthermore, the magnitude of modulation change between trial types was not as large for EMG activity. To quantify this, we computed tuning functions of each muscle during epochs 3 and 4 and compared their amplitudes in visible and invisible trials (Figure S6). Although many muscles were less strongly activated when reaches were not visible, this decrease was only ~14% smaller, on average, across all observations. When we restricted the analysis to just the fourth epoch, the decrease was slightly larger, ~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. Nonetheless, it is possible that the decreased neuronal modulation we observed around the end of movement for invisible trials was in part related to the decreased EMG modulation around the same time.

Previous results show that 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 2 shows several responses with multiple modulation epochs, and these could be differently affected following visuomotor adaptation. Figures 4 and S4 show multi-peaked responses that were affected differently by the visibility of the reach. Although EMG activity patterns were also multi-phasic during our reaching tasks, their encoding patterns did not change abruptly. To illustrate the nature of neuronal encoding changes, and to contrast them with EMG patterns, we compared the activity of each between task types using a sliding correlation analysis (Figure S7). We found that abrupt changes in the correlation of neuronal responses between tasks were common, whereas this was not true for the recorded EMG activity.

DISCUSSION

Numerous lines of experimentation show that M1 modulation is correlated to surprisingly diverse aspects of motor control. This activity is related to muscle excitability (Fetz and Cheney, 1980; Schieber and Rivlis, 2005; Griffin et al., 2015), force production (Cheney and Fetz, 1980; Kalaska et al., 1989; Georgopoulos et al., 1992), limb kinematics and geometry (Georgopoulos et al., 1982; Moran and Schwartz, 1999; van Hemmen and Schwartz, 2008), visuospatial processing and visuomotor transformation (Zhang et al., 1997; Alexander and Crutcher, 1990a; Shen and Alexander, 1997a), and cognitive processing (Georgopoulos et al., 1989; Pellizzer et al., 1995). M1 modulation occurs even in the absence of any active movement, for example, during passive "movement replay" tasks (Taylor et al., 2002; Tkach et al., 2007; Palazzolo, 2015). In one study, M1 responses during replay of a monkey's movement were found to be similar to those observed during the actual movement. Replay activity matched movement activity when both the visual target and the feedback cursor were visible, though displaying the visual target alone was often sufficient for evoking a movement-like response. The authors suggested that "neural activity during [passive] observation is attributable to covert generation of a motor command" (Tkach et al., 2007). In contrast, during active movement, we found that a major response component was absent or reduced when we displayed the visual target but not the feedback cursor (Figures 4, 5, and 6). This result suggests that the feedback-dependent response component was related to active monitoring of the ongoing movement, without being critical for movement generation. Therefore, the role of visual feedback in driving M1 seems to depend on context (i.e., whether the subject is actively moving).

Several investigators have also 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 used a "context recall" version of the task in which the monkey relied on a sequence of stimuli to determine the correct movement direction (Pellizzer et al., 1995). In this task, M1 neurons first coded for the direction of the stimulus in the sequence that cued the correct response and then abruptly changed 100–150 ms later to code the direction of the upcoming movement. Here, we also commonly observed abrupt changes in the encoding coordinate system, though we found these changes occurred at multiple points during a movement (Figures 2, 4, S4, and S7).

Other related dissociation paradigms have involved comparing neuronal responses during two types of movements: those made toward 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 and Crutcher, 1990a), many M1 neurons initially coded the stimulus direction and later coded the

movement direction, a finding comparable with 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 (Figures 5A and 5C). 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. 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 (Figures 5A and 5C, blue solid traces). A similar result was found in a four-target joystick-controlled movement task with a similar visuomotor rotation (Shen and Alexander, 1997a), suggesting that 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 (Figures 3A and 3B). Most of the observed tuning changes were in the positive direction, corresponding to adaptation to the perturbation. For example, the neuronal response in Figure 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. Although 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 et al., 1991; Shen and 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), a process 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. The tuning changes might represent a mixture of both visuospatial-related coding and adaptation-related effects on firing properties, though our experiments did not dissociate these possibilities.

Tuning changes late in the movement are related to the visibility of the cursor trajectory, and although many studies of ''sensorimotor transformation'' have focused on 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 et al., 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 4, 5, 6, and 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 (Figures 4, S4B, S4C, andS7C).

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, 2012; Shenoy et al., 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 curved 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 6) or the physical movement. 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 (Figures 6C and 6D). Although this visualization conveyed a relatively gradual evolution toward zero, the underlying changes in single-unit firing were often abrupt (see especially Figures 4B, S4B, S4C, and S7C). 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 and Schwartz, 1999; Cisek and Scott, 1999).

The concept of smoothly evolving lawful neuronal dynamics, although appealing, may simply reflect the smooth changes characteristic of movement kinematics (Flash and Hogan, 1985). A ''temporal parcellation scheme'' may constitute a useful compromise between these two descriptions of M1 activity (Johnson et al., 2001). Much like the ''dynamical systems'' perspective, this scheme predicts that firing rates covary with task-related parameters in a time-varying sequence (Shenoy et al., 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 et al., 1995; Zhang et al., 1997; Alexander and Crutcher, 1990a), reach-to-grasp requirements (Rouse and Schieber, 2016), or feedback constraints (Inoue et al., 2016; Stavisky et al., 2017) of a given behavior, rather than by fixed evolution rules. Similarly, the abrupt changes in state signified

by changes in tuning that we report here are likely input driven. Importantly, recent dynamical systems methods explicitly model network inputs (e.g., Pandarinath et al., 2018), in contrast to other approaches that have placed emphasis on locally driven network state (Churchland et al., 2012; Russo et al., 2018).

A major finding reported here is that neuronal encoding transitioned abruptly from one parameter to the next. Other 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 et al., 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., 2019). Discrete, step-like transitions in firing have also been reported in other macaque brain regions (Latimer et al., 2015). In the songbird brain region RA (robust nucleus of the arcopallium), neurons produce sequences of spike bursts associated with discrete song syllables (Yu and Margoliash, 1996; Hahnloser et al., 2002). Broadly, these types of discrete activity patterns may be a hallmark of ''cell assemblies,'' groups of neurons that are transiently and simultaneously active during behavior. These assemblies are hypothesized to underlie a general ''neural syntax'' for brain operations (Abeles et al., 1993; 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 and 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 (Figures 3A and 3B versus Figure 3C). Specifically, neuronal tuning changes were most often in the positive direction and could be very large, whereas 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 et al., 1999). We also observed changes in modulation strength from visible to invisible trials for both firing rates (Figures 5B and 5D) and EMG activity (Figure S6), although those decreases were typically much larger for neurons. Classically, reaching movements have been viewed as composed of two phases: an initial ''transport'' phase, considered to be ballistic, 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, whereas 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. A related possibility is that without feedback, some of the changes in neuronal activity were related to the corresponding absence of visually guided sub-movements, though this explanation alone is unlikely to account for all of the changes we observed.

Our 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 and that this input is received toward the end of movement. 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.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- LEAD CONTACT AND MATERIALS AVAILABILITY
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Behavioral Task
  - Neuronal Recordings
  - EMG Recordings
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Data Preprocessing
  - Tuning Curve Analysis
  - Trial Segmentation
  - Reverse Regression Analysis
- DATA AND CODE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j.celrep.2019.11.069.

ACKNOWLEDGMENTS

We wish to thank Jordan Williams for assistance collecting EMG data and Scott Kennedy, Rex Tien, and Andrew Whitford for helpful discussions. S.B.S. was supported by the NIH (NS062019).

AUTHOR CONTRIBUTIONS

S.B.S. and A.B.S. conceived of the study. S.B.S. recorded and analyzed the data and wrote the initial draft of the manuscript. S.B.S. and A.B.S. wrote the final draft of the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: February 12, 2019
Revised: September 16, 2019
Accepted: November 15, 2019
Published: December 17, 2019

REFERENCES

Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. J. Neurophysiol. *70*, 1629–1638.

Abeles, M., Bergman, H., Gat, I., Meilijson, I., Seidemann, E., Tishby, N., and Vaadia, E. (1995). Cortical activity flips among quasi-stationary states. Proc. Natl. Acad. Sci. U S A *92*, 8616–8620.

Aflalo, T.N., and Graziano, M.S.A. (2006). Partial tuning of motor cortex neurons to final posture in a free-moving paradigm. Proc. Natl. Acad. Sci. U S A *103*, 2909–2914.

Alexander, G.E., and Crutcher, M.D. (1990a). Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. J. Neurophysiol. *64*, 164–178.

Ashe, J., and Georgopoulos, A.P. (1994). Movement parameters and neural activity in motor cortex and area 5. Cereb. Cortex *4*, 590–600.

Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. Neuron *68*, 362–385.

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

Cheney, P.D., and Fetz, E.E. (1980). Functional classes of primate corticomotoneuronal cells and their relation to active force. J. Neurophysiol. *44*, 773–791.

Churchland, M.M., and Shenoy, K.V. (2007). Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. J. Neurophysiol. *97*, 4235–4257.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I., and Shenoy, K.V. (2010). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? Neuron *68*, 387–400.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. Nature *487*, 51–56.

Cisek, P., and Scott, S.H. (1999). An alternative interpretation of population vector rotation in macaque motor cortex. Neurosci. Lett. *272*, 1–4.

di Pellegrino, G., and Wise, S.P. (1993). Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. J. Neurosci. *13*, 1227–1243.

Elsayed, G.F., Lara, A.H., Kaufman, M.T., Churchland, M.M., and Cunningham, J.P. (2016). Reorganization between preparatory and movement population responses in motor cortex. Nat. Commun. *7*, 13239.

Fetz, E.E., and Cheney, P.D. (1980). Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. J. Neurophysiol. *44*, 751–772.

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

Fu, Q.-G., Flament, D., Coltz, J.D., and Ebner, T.J. (1995). Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. J. Neurophysiol. *73*, 836–854.

Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., and Massey, J.T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J. Neurosci. *2*, 1527–1537.

Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., and Massey, J.T. (1983). Interruption of motor cortical discharge subserving aimed arm movements. Exp. Brain Res. *49*, 327–340.

Georgopoulos, A.P., Lurito, J.T., Petrides, M., Schwartz, A.B., and Massey, J.T. (1989). Mental rotation of the neuronal population vector. Science *243*, 234–236.

Georgopoulos, A.P., Ashe, J., Smyrnis, N., and Taira, M. (1992). The motor cortex and the coding of force. Science *256*, 1692–1695.

Griffin, D.M., Hoffman, D.S., and Strick, P.L. (2015). Corticomotoneuronal cells are "functionally tuned". Science *350*, 667–670.

Hahnloser, R.H.R., Kozhevnikov, A.A., and Fee, M.S. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. Nature *419*, 65–70.

Harpaz, N.K., Ungarish, D., Hatsopoulos, N.G., and Flash, T. (2019). Movement decomposition in the primary motor cortex. Cereb. Cortex *29*, 1619–1633.

Inoue, M., Uchimura, M., and Kitazawa, S. (2016). Error signals in motor cortices drive adaptation in reaching. Neuron *90*, 1114–1126.

Inoue, Y., Mao, H., Suway, S.B., Orellana, J., and Schwartz, A.B. (2018). Decoding arm speed during reaching. Nat. Commun. *9*, 5243.

Johnson, M.T.V., and Ebner, T.J. (2000). Processing of multiple kinematic signals in the cerebellum and motor cortices. Brain Res. Brain Res. Rev. *33*, 155–168.

Johnson, M.T.V., Mason, C.R., and Ebner, T.J. (2001). Central processes for the multiparametric control of arm movements in primates. Curr. Opin. Neurobiol. *11*, 684–688.

Kakei, S., Hoffman, D.S., and Strick, P.L. (1999). Muscle and movement representations in the primary motor cortex. Science *285*, 2136–2139.

Kalaska, J.F., Cohen, D.A.D., Hyde, M.L., and 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, J.F., Scott, S.H., Cisek, P., and Sergio, L.E. (1997). Cortical control of reaching movements. Curr. Opin. Neurobiol. *7*, 849–859.

Kass, R.E., Ventura, V., and Brown, E.N. (2005). Statistical issues in the analysis of neuronal data. J. Neurophysiol. *94*, 8–25.

Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. Nat. Neurosci. *17*, 440–448.

Keele, S.W., and Posner, M.I. (1968). Processing of visual feedback in rapid movements. J. Exp. Psychol. *77*, 155–158.

Krakauer, J.W., Pine, Z.M., Ghilardi, M.-F., and Ghez, C. (2000). Learning of visuomotor transformations for vectorial planning of reaching trajectories. J. Neurosci. *20*, 8916–8924.

Lara, A.H., Elsayed, G.F., Zimnik, A.J., Cunningham, J.P., and Churchland, M.M. (2018). Conservation of preparatory neural events in monkey motor cortex regardless of how movement is initiated. eLife *7*, e31826.

Latimer, K.W., Yates, J.L., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2015). Neuronal modeling. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. Science *349*, 184–187.

Lurito, J.T., Georgakopoulos, T., and Georgopoulos, A.P. (1991). Cognitive spatial-motor processes. 7. The making of movements at an angle from a stimulus direction: studies of motor cortical activity at the single cell and population levels. Exp. Brain Res. *87*, 562–580.

Meyer, D.E., Abrams, R.A., Kornblum, S., Wright, C.E., and Smith, J.E.K. (1988). Optimality in human motor performance: ideal control of rapid aimed movements. Psychol. Rev. *95*, 340–370.

Milner, T.E. (1992). A model for the generation of movements requiring endpoint precision. Neuroscience *49*, 487–496.

Moran, D.W., and Schwartz, A.B. (1999). Motor cortical representation of speed and direction during reaching. J. Neurophysiol. *82*, 2676–2692.

Novak, K.E., Miller, L.E., and Houk, J.C. (2002). The use of overlapping submovements in the control of rapid hand movements. Exp. Brain Res. *144*, 351–364.

Paillard, J. (1982). The contribution of peripheral and central vision to visually guided reaching. In Analysis of Visual Behavior, D.J. Ingle, M.A. Goodale, and R.J.W. Mansfield, eds. (MIT Press), pp. 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.

Palazzolo, M. (2015). Classifying neural signals related to action perception. PhD thesis (University of Pittsburgh).

Pandarinath, C., O'Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. Nat. Methods *15*, 805–815.

Paninski, L., Fellows, M.R., Hatsopoulos, N.G., and Donoghue, J.P. (2004). Spatiotemporal tuning of motor cortical neurons for hand position and velocity. J. Neurophysiol. *91*, 515–532.

Pellizzer, G., Sargent, P., and Georgopoulos, A.P. (1995). Motor cortical activity in a context-recall task. Science *269*, 702–705.

Porter, R., and Lemon, R. (1993). Corticospinal Function and Voluntary Movement (Oxford: Clarendon Press).

Rouse, A.G., and Schieber, M.H. (2016). Spatiotemporal distribution of location and object effects in primary motor cortex neurons during reach-to-grasp. J. Neurosci. *36*, 10640–10653.

Russo, A.A., Bittner, S.R., Perkins, S.M., Seely, J.S., London, B.M., Lara, A.H., Miri, A., Marshall, N.J., Kohn, A., Jessell, T.M., et al. (2018). Motor cortex embeds muscle-like commands in an untangled population response. Neuron *97*, 953–966.e8.

Schieber, M.H., and 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, A.B. (1994). Direct cortical representation of drawing. Science *265*, 540–542.

Schwartz, A.B., and Moran, D.W. (1999). Motor cortical activity during drawing movements: Population representation during lemniscate tracing. J. Neurophysiol. *82*, 2705–2718.

Schwartz, A.B., Kettner, R.E., and Georgopoulos, A.P. (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*, 2913–2927.

Schwartz, A.B., Moran, D.W., and Reina, G.A. (2004). Differential representation of perception and action in the frontal cortex. Science *303*, 380–383.

Shadmehr, R., and Mussa-Ivaldi, F.A. (1994). Adaptive representation of dynamics during learning of a motor task. J. Neurosci. *14*, 3208–3224.

Shen, L., and Alexander, G.E. (1997a). Neural correlates of a spatial sensory-to-motor transformation in primary motor cortex. J. Neurophysiol. *77*, 1171–1194.

Shenoy, K.V., Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. Annu. Rev. Neurosci. *36*, 337–359.

Soechting, J.F., and Lacquaniti, F. (1981). Invariant characteristics of a pointing movement in man. J. Neurosci. *1*, 710–720.

Stavisky, S.D., Kao, J.C., Ryu, S.I., and Shenoy, K.V. (2017). Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neural state space dimensions. Neuron *95*, 195–208.e9.

Suway, S.B., Orellana, J., McMorland, A.J.C., Fraser, G.W., Liu, Z., Velliste, M., Chase, S.M., Kass, R.E., and Schwartz, A.B. (2017). Temporally segmented directionality in the motor cortex. Cereb. Cortex *2017*, 1–14.

Taylor, D.M., Tillery, S.I., and Schwartz, A.B. (2002). Direct cortical control of 3D neuroprosthetic devices. Science *296*, 1829–1832.

Tkach, D., Reimer, J., and Hatsopoulos, N.G. (2007). Congruent activity during action and action observation in motor cortex. J. Neurosci. *27*, 13241–13250.

Van Hemmen, J.L., and Schwartz, A.B. (2008). Population vector code: A geometric universal as actuator. Biol. Cybern. *98*, 509–518.

Velliste, M., Perel, S., Spalding, M.C., Whitford, A.S., and Schwartz, A.B. (2008). Cortical control of a prosthetic arm for self-feeding. Nature *453*, 1098–1101.

Velliste, M., Kennedy, S.D., Schwartz, A.B., Whitford, A.S., Sohn, J.-W., and McMorland, A.J. (2014). Motor cortical correlates of arm resting in the context of a reaching task and implications for prosthetic control. J. Neurosci. *34*, 6011–6022.

Wise, S.P., Moody, S.L., Blomstrom, K.J., and Mitz, A.R. (1998). Changes in motor cortical activity during visuomotor adaptation. Exp. Brain Res. *121*, 285–299.

Woodworth, R.S. (1899). The accuracy of voluntary movement. Psy. Rev. Monograph Supp *13*, i–i114.

Yu, A.C., and Margoliash, D. (1996). Temporal hierarchical control of singing in birds. Science *273*, 1871–1875.

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

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and Algorithms		
MATLAB	Mathworks	https://www.mathworks.com/products/matlab.html
Offline Sorter	Plexon Inc.	https://plexon.com/products/offline-sorter/
Other		
MAP Data Acquisition System	Plexon Inc.	https://plexon.com/products/map-data-acquisition-system-plexon/
Cerebus System	Blackrock Microsystems LLC	https://www.blackrockmicro.com/neuroscience-research-products/neural-data-acquisition-systems/cerebus-daq-system/
Utah Array	Blackrock Microsystems LLC	https://www.blackrockmicro.com/electrode-types/utah-array/
RZ2 Bioamp Processor	Tucker-Davis Technologies	https://www.tdt.com/component/rz2-bioamp-processor/

LEAD CONTACT AND MATERIALS AVAILABILITY

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Andrew B Schwartz (abs21@pitt.edu). The data supporting the current study have not been deposited in a public repository but are available from the corresponding author upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Two adult male monkeys (*Macaca mulatta*, R and P) were used in the present study. All procedures for the care and use of these animals were in accordance with the guidelines of the NIH and were approved by the Institutional Animal Care and Use Committee of the University of Pittsburgh.

METHOD DETAILS

Behavioral Task

Monkeys 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 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 rapid 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 and Alexander, 1997a). The processing of the scene displayed in VR during a reach may be another factor driving changes in M1 firing patterns. That is, visual feedback from the cursor movement may drive firing. 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. Center-out reaching tasks commonly require stopping and holding within a target zone, which is difficult without visual feedback. We therefore omitted the terminal hold period for tasks used in this study; monkeys were rewarded upon touching the target with the cursor. The terminal hold period was omitted in all trials, regardless of whether feedback was provided.

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.

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 into the muscles 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.

QUANTIFICATION AND STATISTICAL ANALYSIS

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 and Georgopoulos, 1994; Moran and 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). The magnitude of the across-target mean could vary across trial types (standard versus rotated, visible versus invisible), though a comprehensive analysis of this variation was outside the scope of our study. For visualization only, we applied "PC smoothing" to firing rates (Churchland et al., 2010, 2012). This was applied to neuronal responses separately within each task type to prevent artificially mixing response properties.

Figure S1 shows a comparison of firing rates processed using these steps along with a less-processed version of the same rates. For panels A and B, the left-hand column shows rates computed using a 10 ms bin width and smoothed using a Gaussian with a 20 ms SD. The across-target mean (black line) was not subtracted at each time step, and PC smoothing was not applied. The right-hand column shows the same rates but processed as described above. As previously reported (Suway et al., 2017), using a narrower smoothing kernel did not generally reveal important higher-frequency response components.

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 were discarded. These trials primarily occurred during the adaptation period between blocks. We chose to analyze only straight reaches so that the time-course of behavior was consistent on each trial. Monkeys were well-trained on each task and only a small proportion of trials was discarded: we retained 85% and 80% of successful trials for monkeys R and P, respectively.

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.

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:

$$y = k + B \cdot \cos(\theta - \theta_{PD}) + \varepsilon \quad (1)$$

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 1 to EMG activations in place of firing rates. In that case, direction was shifted backward 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.

Trial Segmentation

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. Factor loadings spanned the entire trial, though we found that loadings for a given factor were large only during a short time period and near-zero at other times. Thus, the factors corresponded to sequential temporal features which were then 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.

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 et al., 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:

$$V_x = k_x + B_x R + \varepsilon \quad (2)$$

$$V_y = k_y + B_y R + \varepsilon \quad (3)$$

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 2 and 3 were fit using ridge regression rather than ordinary least-squares regression. Ridge regression is useful when 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 2 and 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.

DATA AND CODE AVAILABILITY

Data and code used in this manuscript will be available from the corresponding author upon reasonable request.