Area 2 of primary somatosensory cortex encodes kinematics of the whole arm

Raeed H Chowdhury[1], Joshua I Glaser[2,3,4], Lee E Miller[1,5,6,7]

[1] Department of Biomedical Engineering, Northwestern University, Evanston, IL
[2] Interdepartmental Neuroscience Program, Northwestern University, Chicago, IL
[3] Department of Statistics, Columbia University, New York City, NY
[4] Zuckerman Mind Brain Behavior Institute, Columbia University, New York City, NY
[5] Department of Physiology, Northwestern University, Chicago, IL
[6] Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL
[7] Shirley Ryan AbilityLab, Chicago, IL

Abstract

Proprioception, or the sense of body position, movement, and associated forces, remains poorly understood, despite its clear importance to making coordinated movements. In the classic model of neurons in area 2, a proprioceptive area of primary somatosensory cortex, neural activity is related simply to the movement of the hand and interaction forces it encounters during movement, similar to our conscious experience of proprioception. However, in two separate experiments, examining 1) active and passive movements and 2) reaches to targets in two different workspaces, we found that a model of area 2 activity built on kinematics of the whole arm successfully predicted how features of neural activity changed across movement conditions, while the classic hand-based model was unable to. This suggests that, unlike our conscious experience and the classic model of proprioception in cortex, neural activity in area 2 represents movement of the whole arm.

1 Introduction

Moving around in an uncontrolled environment is a remarkably complex feat. In addition to the necessary computations on the efferent side to generate movement, an important aspect of motor control is the afferent information we receive from our limbs, essential both for movement planning and for the feedback it provides during movement. Of the sensory modalities we receive, proprioception, or the sense of body position, movement and forces, is arguably the most critical for making coordinated movements (Ghez and Sainburg 1995; Gordon et al. 1995; Sainburg et al. 1995; Sainburg et al. 1993; Sanes et al. 1984). However, despite its importance, few studies have explicitly addressed how proprioception is represented in the brain during natural movement. In comparison, touch, vision, and the motor areas of the brain have received far more attention.

Our conscious experience of proprioception typically focuses on where our hands are going, rather than the rotations of our joints. This also matches with psychophysical data showing that humans are better at estimating the location of the hand than estimating joint angles (Fuentes and Bastian 2010). Perhaps consequently, of the few studies that examine central nervous system (CNS) activity underlying proximal limb proprioception, most assume that individual neurons represent the movement of the limb's endpoint, rather than joint angles or muscle lengths (Bosco et al. 2000; Bosco et al. 1996; London and Miller 2013; London et al. 2011; Prud'homme and Kalaska 1994; Weber et al. 2011). However, to achieve such a hand-centric view, neurons would need to integrate the proprioceptive signals from sensors in the muscles, joints, and skin of the whole limb (Goodwin et al. 1972; Proske and Gandevia 2012).

The proximal arm representation within area 2 of primary somatosensory cortex (S1) receives a combination of muscle and cutaneous information (Hyvärinen and Poranen 1978; Padberg et al. 2018; Pons et al. 1985) and is thought to be important for proprioception (Jennings et al. 1983; Kaas et al. 1979; London and Miller 2013). Interestingly, recent computational studies have shown that while neural activity may appear to be tuned to the state of the limb's endpoint, features of this tuning might be a direct consequence of the biomechanics of the limb (Chowdhury et al. 2017; Lillicrap and Scott 2013). Consistent with those results, we have recently observed, using artificial neural networks, that that muscle lengths were better predictors of area 2 activity than were hand kinematics (Lucas et al. 2019). Thus, in this study, we set out to understand what information is represented in area 2, with two experiments that altered the relationship between hand and whole limb kinematics. Using both classic single neuron analysis techniques like tuning curves and preferred movement directions (Georgopoulos et al. 1982; Prud'homme and Kalaska 1994; Sergio and Kalaska 2003), as well as more recently developed neural population analyses (Churchland et al. 2012; Cunningham and Byron 2014), we discovered several features of neural activity that could not be explained by the classic hand-based model. However, using biomechanical modeling, we show that models relating neural activity to kinematics of the whole arm can explain these features. Our results indicate that if there is a transformation towards representing reaching in terms of only the hand, it likely occurs beyond area 2.

2 Results

For the experiments detailed in this paper, we recorded neural signals from three monkeys (H, C, and L) using Utah multi-electrode arrays (Blackrock Microsystems) implanted in the arm representation within Brodmann's area 2 of S1 (Figure 1A). We trained each of these monkeys to grasp a two-link planar manipulandum and make reaching movements to targets presented on a screen in front of them (Figure 1B). During these behavioral sessions, we also tracked the locations of markers on the monkey's arm using a custom motion tracking system and registered these markers to a musculoskeletal model in OpenSim (SimTK) to estimate joint angles and muscle lengths. Our experiments included two components: one comparing active and passive movements and one comparing reaching movements in two different workspaces.

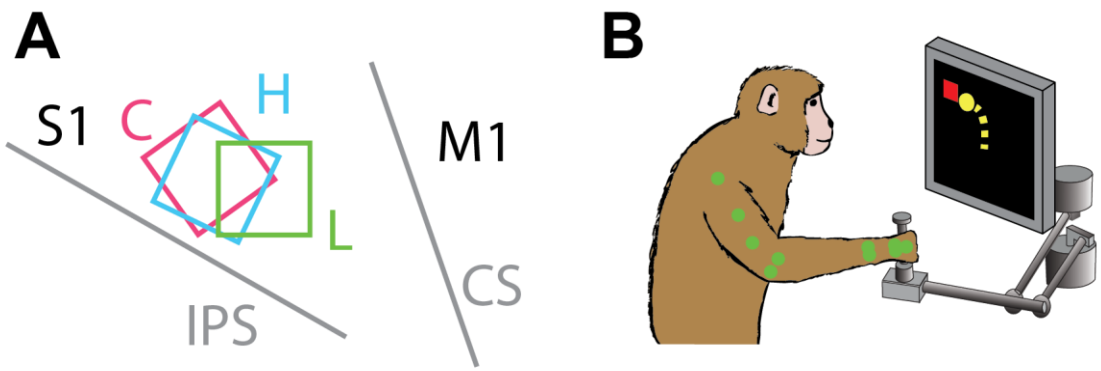

Figure 1: General experimental setup. A –locations of Utah arrays implanted in area 2 of primary somatosensory cortex (S1) of Monkeys C, H, and L. M1, primary motor cortex; IPS, intraparietal cortex; CS central sulcus. B – Behavioral task. Monkey controls a cursor on screen (yellow) with a two link manipulandum to reach to visually presented targets (red). We track the locations of markers on the monkey's arm (green) during the task.

2.1 Whole arm, rather than hand-based models, explain somatosensory area 2 representation of active and passive movements

We reported previously that the direction of maximal response (the "preferred direction"; PD) of single neurons in area 2 is typically similar during active reaching movements and passive perturbations of the hand (London and Miller 2013). However, this section shows that despite the apparent similarity in directional tuning of individual neurons, information from a population of neurons can be used to clearly distinguish active and passive movements, which is unexpected given our previous results. We set out to explore whether this separation could be explained by modeling the neural activity in area 2 purely in terms of behavioral variables.

In this experiment, the monkey performed a center-out reaching task to four targets. On half of these reaching trials, the monkey's hand was bumped by the manipulandum during the center-hold period in one of the four target directions (Figure 2). See Methods section or (London and Miller 2013) for task details. This experiment involved two sessions with each of Monkeys C and H.

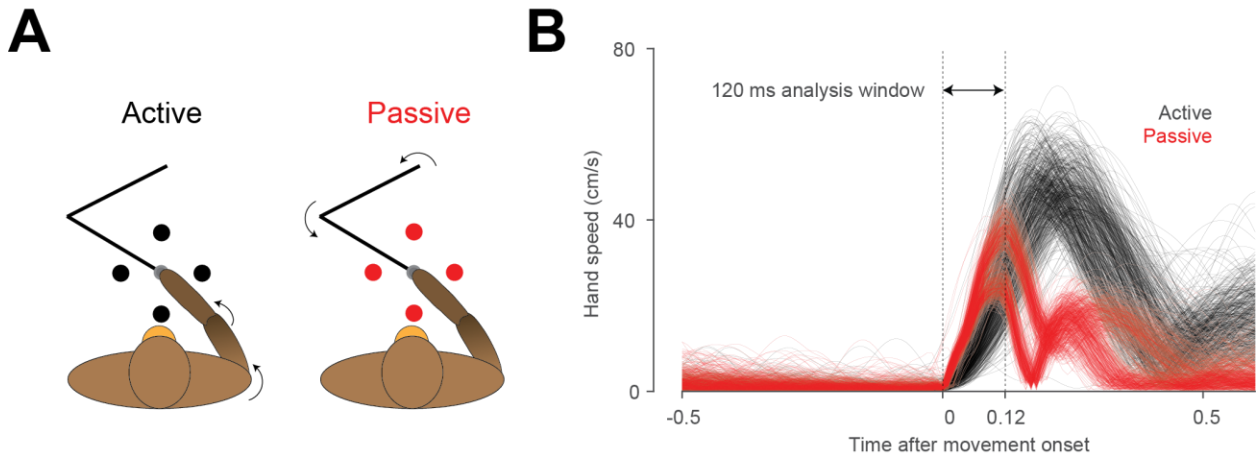

Figure 2: Active vs. passive behavior. A – Schematic of task. On active trials (black), monkey reaches from center target to a visually presented target in one of four directions. On passive trials, manipulandum bumps monkey's hand in one of the four target directions. B – Speed of hand during active (black) and passive (red) trials, plotted against time, for one session. Starting around 120 ms after movement onset, a bimodal distribution in passive movement speed emerges. This bimodality reflects differences in the stiffness of the arm for different directions of movement. Perturbations towards and away from the body tended to result in a shorter overall movement than those to the left or right. However, average movement speed was similar between active and passive trials in the 120 ms after movement onset, the window used for neural analysis.

Of necessity, the similar active and passive tuning described by London and Miller included only neurons that had significant sinusoidal tuning to both active and passive movement. Consequently, it was based on fewer than 30% of the recorded cells. Similarly, in this study, only 29 of 57 recorded neurons were significantly tuned to both types of movements for Monkey C and only 12 of 83 for Monkey H (Figure 3). The low percentage of tuned neurons is at least partially due to measurement error (Stevenson et al. 2011), exacerbated by the small amount of data available in the first 120 ms after movement onset when the kinematics of the two types of movements were similar. Thus, there are clear limitations inherent to single neuron tuning curve-based analyses.

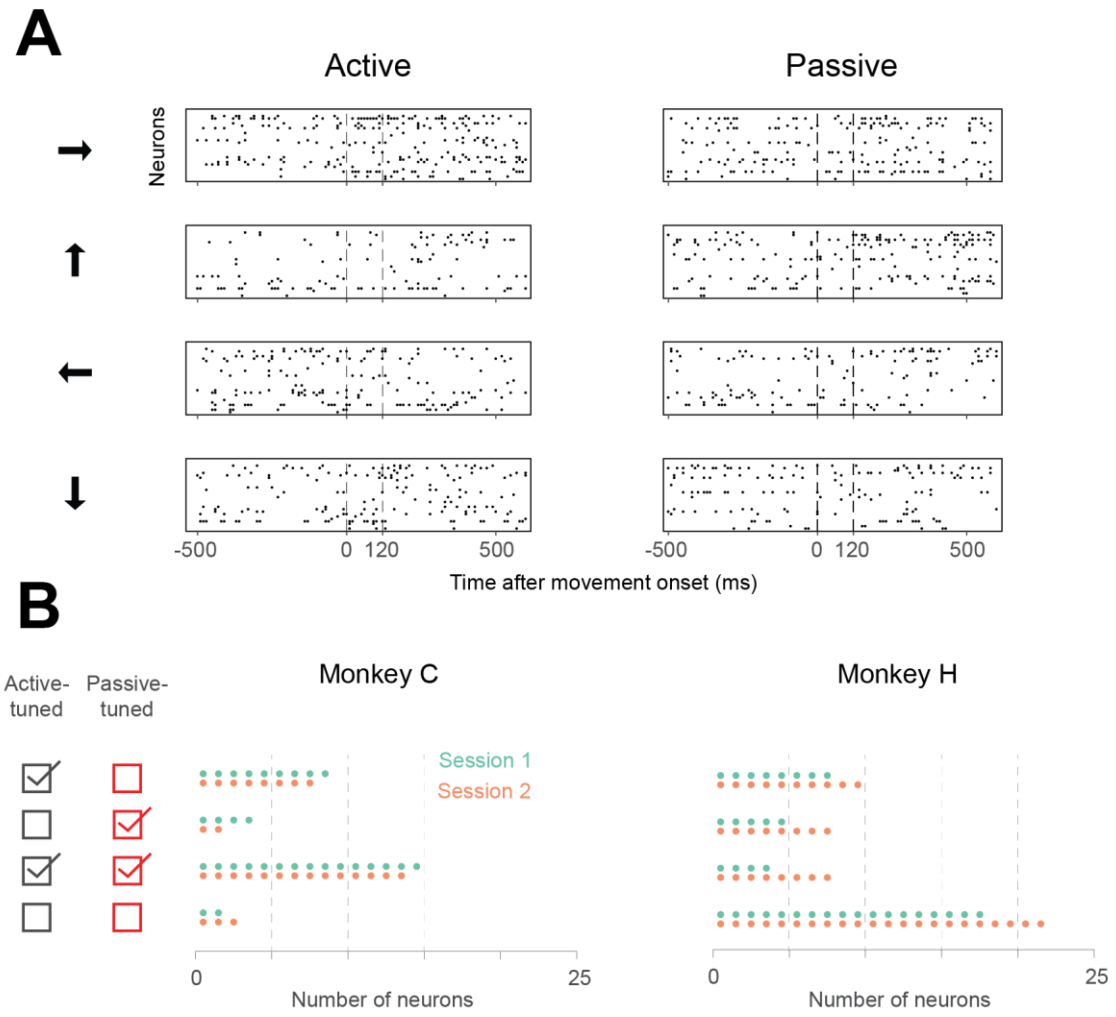

Figure 3: Neural Activity during active vs. passive task. A – Neural raster plots for example active and passive trials in each movement direction (indicated by arrows to the left of each raster plot). In each plot, rows indicate spikes recorded from different neurons, plotted against time. Vertical dashed lines indicate boundaries of analysis window. B – Tuning properties for neurons. Dot histograms indicate counts of neurons tuned only to active movements, those tuned only to passive movements, those tuned to both, and those tuned to neither. Note that tuning estimates from Monkey H have more error due to the fewer trials that were typically available in a given session, leading to a higher proportion of apparently untuned neurons.

Because we had simultaneous recordings from many neurons, we could circumvent this limitation by analyzing the *population* responses to active and passive movements. Moreover, by aggregating information across neurons instead of averaging across trials, the neural population analysis allows the analysis of single trials. Recent studies have uncovered a host of neural coding properties that are difficult to infer from single neuron analyses, yet are revealed with single-trial, population-level analyses (Cunningham and Byron 2014; Dekleva et al. 2018; Perich et al. 2018). In this case, our population analysis treats each active or passive trial as a point in a high-dimensional neural state space, where the value of each dimension corresponds to the

activity of a single neuron. Using Principal Components Analysis (PCA), which extracts the sequential, orthogonal dimensions of highest covariance from the neural population activity, we uncovered an unexpected linear boundary between active and passive movements within the first three PCs (Figure 4). To quantify this separation, we trained a linear discriminant analysis (LDA) model to classify the movement type in each of the four sessions. On average, for data not used to train the LDA model, 84% of movements were correctly classified as active or passive, indicating that despite the similarity of hand-movement coding by single neurons in area 2, the neural population distinguishes the different movement types.

To explore potential causes of this separation in neural state space, we attempted to replicate it using synthetic state spaces constructed from simplified neuron activity models that assume different representations in area 2 (see Methods for modeling details). As a baseline, we fit a purely hand kinematics model, in which neural activity was driven by the position and velocity of hand movement. Classification rate of this simple model was slightly above chance but well under the actual data, with only 65% correct movement classification (Figure 4). Thinking that the difference between active and passive movements may have been related to the differing forces on the hand in the two conditions, we also fit a hand kinematics+force model, like the one described in (Prud'homme and Kalaska 1994), where neural activity was driven by both the velocity of the hand and the forces applied to it. Reaching only 64% correct classification, the addition of force made no difference, indicating that area 2 contains some additional information that allows it to distinguish between active and passive movements.

Finally, we fit a model that represented the Cartesian positions and velocities of the elbow as well as the hand. Unlike the others, this model resulted in an average movement classification accuracy of 82% (Figure 4). This result provides a possible explanation of how area 2 separates the two types of movement: while hand movement is similar, the movement of the whole arm differs between active and passive conditions. While there might be additional explanations for this separability, like an efference copy signal from motor areas specifically during active movements (Bell 1981; London and Miller 2013; Nelson 1987), this finding suggests that the parsimonious hand/elbow model is sufficient to explain it, unlike the hand-based models.

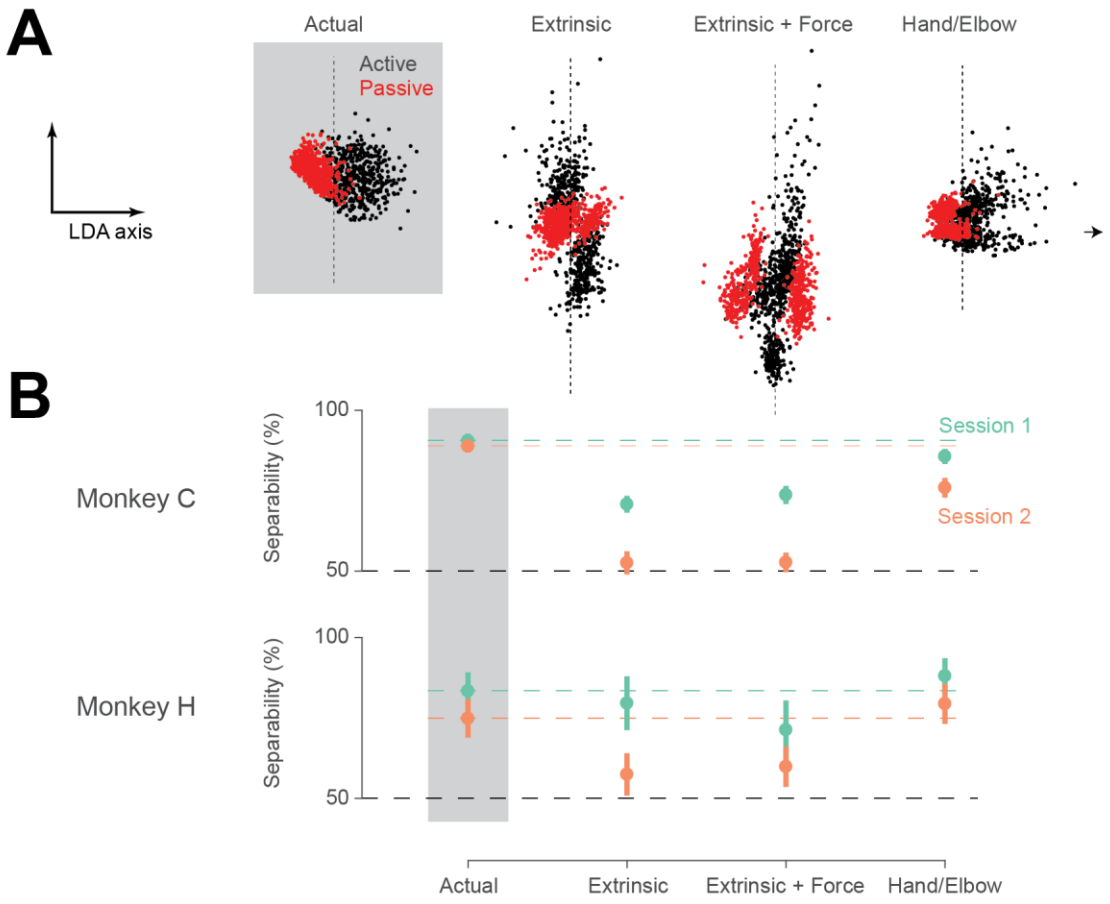

Figure 4: Area 2 population separability. A – Population activity plots for an example session with Monkey C. Each dot in the scatter plots represents population neural activity during one trial, projected onto the LDA axis of highest separation between movement types (within the first three PCs of neural population activity). Vertical axis is an arbitrary axis orthogonal to the LDA axis in the three-dimensional PC space. Black dots indicate active trials, and red dots indicate passive trials. Leftmost panel shows actual population activity, while other three panels show population activity predicted by extrinsic, extrinsic + force, and hand/elbow models, from left to right. Black arrow in hand/elbow plot indicates vertical axis location of one active trial projected far to the right of the main cluster; this point was excluded from the figure for visual clarity of the main cluster. Note that actual population activity and activity predicted by the hand/elbow model separate active and passive movements, while activity predicted by the extrinsic and extrinsic + force models do not. B – Percent separability of active and passive movements for actual neural population activity and activity predicted by the three models. Error bars indicate 95% confidence intervals (derived from cross-validation – see Methods), and colored dashed lines indicates actual population separability for a given session. Black dashed line indicates chance separability. Extrinsic and extrinsic + force models generally yielded lower separability than the hand/elbow model, which predicts similar separability to that found in actual population activity.

2.2 Whole arm, rather than hand-based models, explain somatosensory area 2 representation of movements in two separate workspaces

Given the importance of whole-arm kinematics for explaining neural population activity revealed by the active vs. passive experiment, we conducted a second experiment to further examine how

single neurons represent reaching movements that are constrained to two disjoint areas in front of the monkey, henceforth called workspaces (Figure 5A). We found that whole-arm models and hand-based models made different predictions of neural activity across these two conditions, allowing us to further characterize the extent to which area 2 represented whole-arm movements rather than just the movements of the hand.

In this experiment, we tested four different models of area 2 encoding, which we titled “extrinsic”, “egocentric”, “muscle”, and “hand-elbow”, schematized in Figure 5D. The extrinsic and egocentric models stem from classic, endpoint models of limb movement-related neural activity (Bosco and Poppele 2001; Bosco et al. 1996; Caminiti et al. 1990; Georgopoulos et al. 1982; Prud'homme and Kalaska 1994). Both represent the kinematics of the hand, but in different coordinate frames: the extrinsic model assumes neurons relate to Cartesian coordinates of hand position and velocity, while the egocentric model assumes neurons relate to spherical coordinates with respect to the shoulder. On the other hand, the muscle and hand-elbow models account more fully for how the whole arm moves. In the muscle model, neural activity was assumed to be driven by the length and stretch velocity of the muscles (see Methods for details). In the hand-elbow model, the activity was related to the Cartesian kinematics (position and velocity) of both the hand and the elbow markers. As in the active vs. passive experiment, we aimed to test how well these models predicted features of neural activity during reaching.

However, a challenge in comparing these models of neural representation is that for the typical, center-out reaching task in a small workspace, kinematic signals in the various coordinate frames are highly correlated. Because a high correlation means that a linear transform can accurately convert one coordinate frame into another, all four models make very similar predictions of neural activity. To deal with this problem, we trained the monkeys to reach to randomly generated targets presented in two different workspaces: one close to the body and contralateral to the reaching hand, and one distant and ipsilateral. This had two important effects. First, the random locations of the targets lessened the stereotopy of the movements, allowing for the collection of more varied movement data than from the center-out paradigm. Second, the average posture in each workspace was significantly different, such that while signals in the different model coordinate frames were still correlated within each workspace, this correlation (and the mapping between coordinate frames) changed significantly between workspaces. This forced the models to make different predictions of neural activity across the two workspaces. Figure 5A schematizes the two-workspace random target behavior, and Figure 5B shows an example raster plot of neurons while a monkey performed the task. We recorded three sessions with each of Monkeys C and H and two sessions with Monkey L.

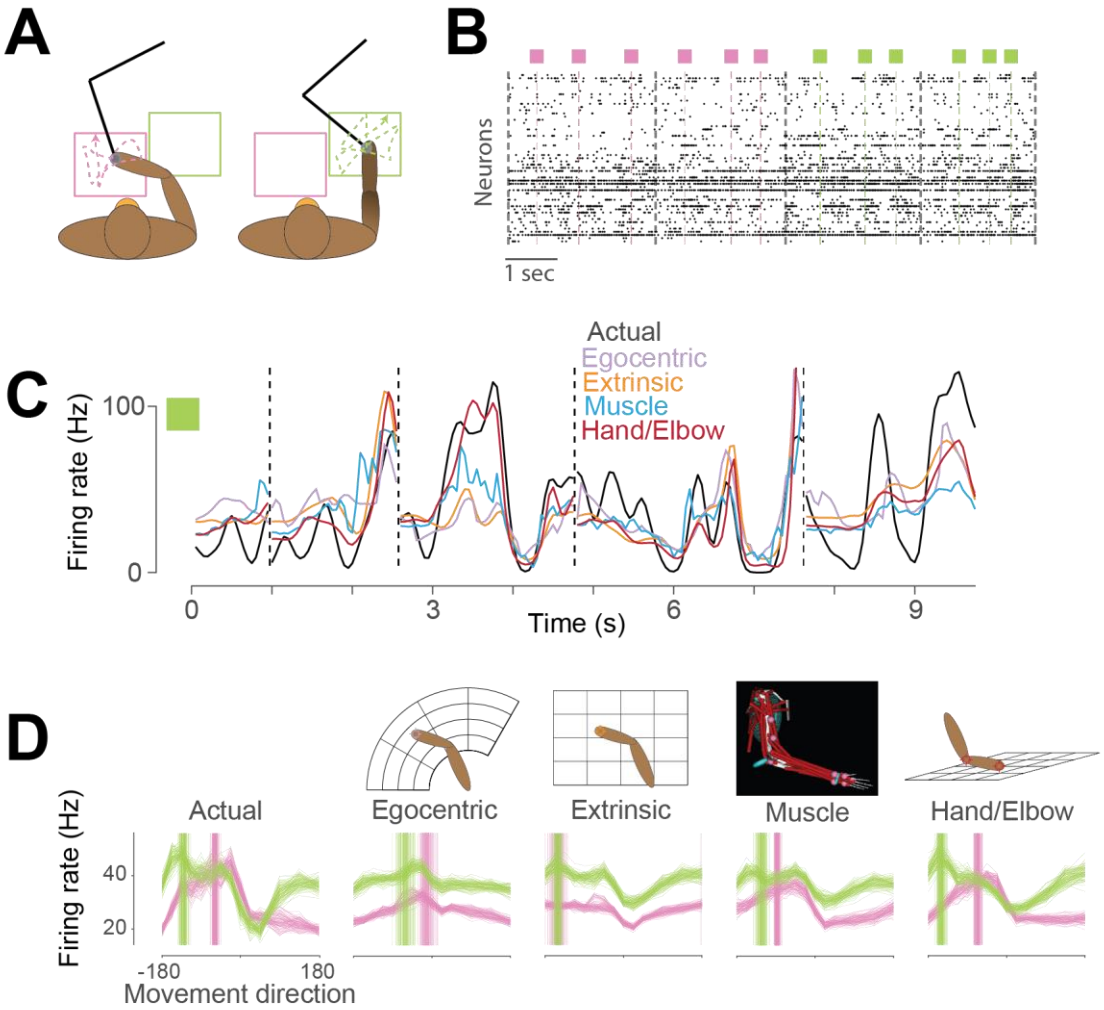

Figure 5: Behavior and example neural activity for two-workspace task. A – Two-workspace behavior. On each trial, monkey makes reaches to random targets in one of two workspaces: one close to the body and contralateral to the reaching hand (pink) and the other distant and ipsilateral (green). B – Example neural raster plot from one session for two randomly drawn trials in each workspace. Dots in each row represent activity for one of the simultaneously recorded neurons. Black dashed lines indicate starts and ends of trials, and colored dashed lines and boxes indicate times of target presentation, with color indicating the workspace for the trial. C – Firing rate plot for an example neuron during five randomly drawn trials from the distal (green) workspace. Black trace represents neural firing rate, smoothed with a 50 ms gaussian kernel. Colored traces represent firing rates predicted by egocentric (purple), extrinsic (orange), muscle (blue), and hand/elbow (red) models (not smoothed). D – Actual and predicted tuning curves and preferred directions (PDs) for an example neuron in the two workspaces. Leftmost plot shows tuning curves and PDs for actual neural firing rate, while other plots show curves and PDs for activity predicted by each of the models. Each panel plots firing rate against direction of hand movement for each workspace. Vertical bars indicate preferred directions calculated in each workspace. Individual traces and vertical bars are tuning curves and preferred directions calculated during one cross-validation run. Note that both the actual tuning curve and PD change between workspaces. Neither the egocentric nor extrinsic models predict these changes, but the muscle and hand/elbow models can.

The simplest evaluation of these models is to compare how well they predicted actual neural firing rates (Figure 5C). To assess this, we used repeated k-fold cross-validation of a goodness-of-fit metric (see Methods for more details). Here, normal goodness-of-fit metrics like R^2 or variance-accounted-for (VAF) are ill-suited to the Poisson-like statistics of neural activity; instead, we used the likelihood-based pseudo-R^2 (Cameron and Windmeijer 1997; 1996; McFadden 1977). Like VAF, psuedo-R^2 has a maximum value of 1, but can also be negative for models that do worse than predicting the mean firing rate during cross-validation. In general, the values corresponding to what constitutes a good fit are lower for pR^2 than for either R^2 or VAF. Of our four models, the whole-arm models out-performed the hand models (Figure 6). Generally, the hand-elbow model was the most predictive of actual neural firing rate across the two workspaces, winning the great majority of the pairwise comparisons with the other models (Figure 6B, S1). Of the 288 neurons recorded across the 8 sessions, in only 47 did either of the hand-based models win a pairwise comparison against either of the whole-arm models (using $\alpha = 0.05$ and a Bonferroni correction to account for six total pairwise comparisons for each neuron). For the other 241 cells, neither hand-based model could out-predict either whole-arm model.

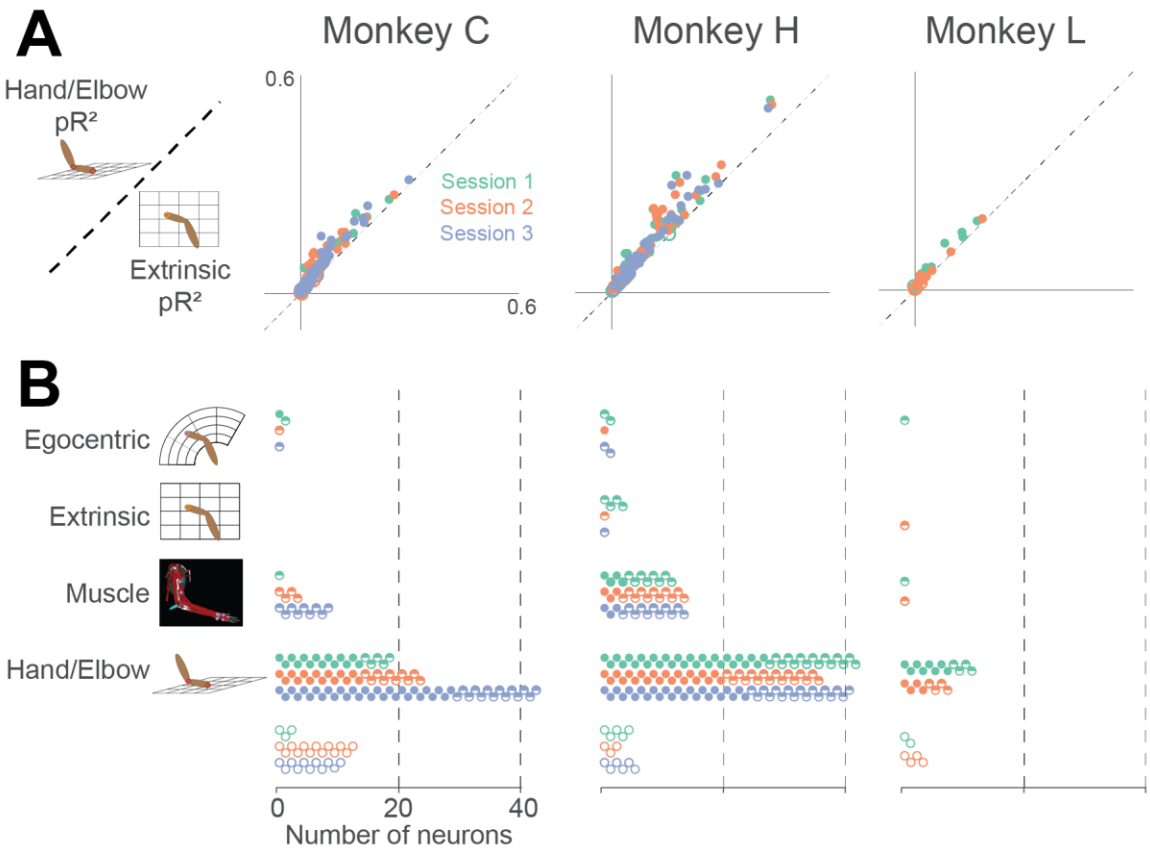

Figure 6: Goodness-of-fit comparison analysis, with one column for each monkey. A – Scatter plots comparing the pseudo-R^2 of the hand/elbow model to that of the extrinsic model for each monkey. Each point in the scatter plot represents the pseudo-R^2 values of one neuron, with hand/elbow pseudo-R^2 on the vertical axis and extrinsic pseudo-R^2 on the horizontal. Different colors represent neurons recorded during different sessions. Filled circles represent neurons for which one model's pseudo-R^2 was

significantly higher than that of the other model. Conversely, open circles represent neurons for which the model pseudo-R^2 values were not significantly different. In this comparison, most filled circles lie above the dashed unity line, indicating that the hand/elbow model performed better than the extrinsic model. All six pairwise comparisons for pseudo-R^2 are shown in Figure S1. B – Dot plot of pairwise comparison winners. Each filled circle in this plot represents a neuron for which a model won all pairwise comparisons with the other three models. Each half-filled circle represents a neuron for which a model won all but one comparison. Empty circles, in the bottom group, correspond to neurons for which there was no model that won at least two of the comparisons. Color indicates the session in which a neuron was recorded. This figure shows that the hand/elbow model won at least two out of three pairwise comparisons for most neurons.

We further tested our models on how well features (e.g., tuning curves and PDs) computed from the model-predicted firing rates matched those of the actual data. Figure 5D shows the directional tuning curves for an example neuron, along with the tuning curves predicted by each model. We calculated the correlation between the predicted and the actual tuning curves in different workspaces as a measure of their similarity. With this measure, the hand-elbow model resulted in the best reconstruction of tuning curves, once again winning most of the pairwise comparisons with the other models (Figure 7, S2). In this case, for only 37 neurons did either of the hand-based models win any pairwise comparison against a whole-arm model (using $\alpha = 0.05$ and a Bonferroni correction to account for six total pairwise comparisons for each neuron) For the other 251 cells, neither hand-based model could out-predict either whole-arm model. Therefore, this experiment further suggests that area 2 of S1 encodes whole-arm rather than just hand kinematics.

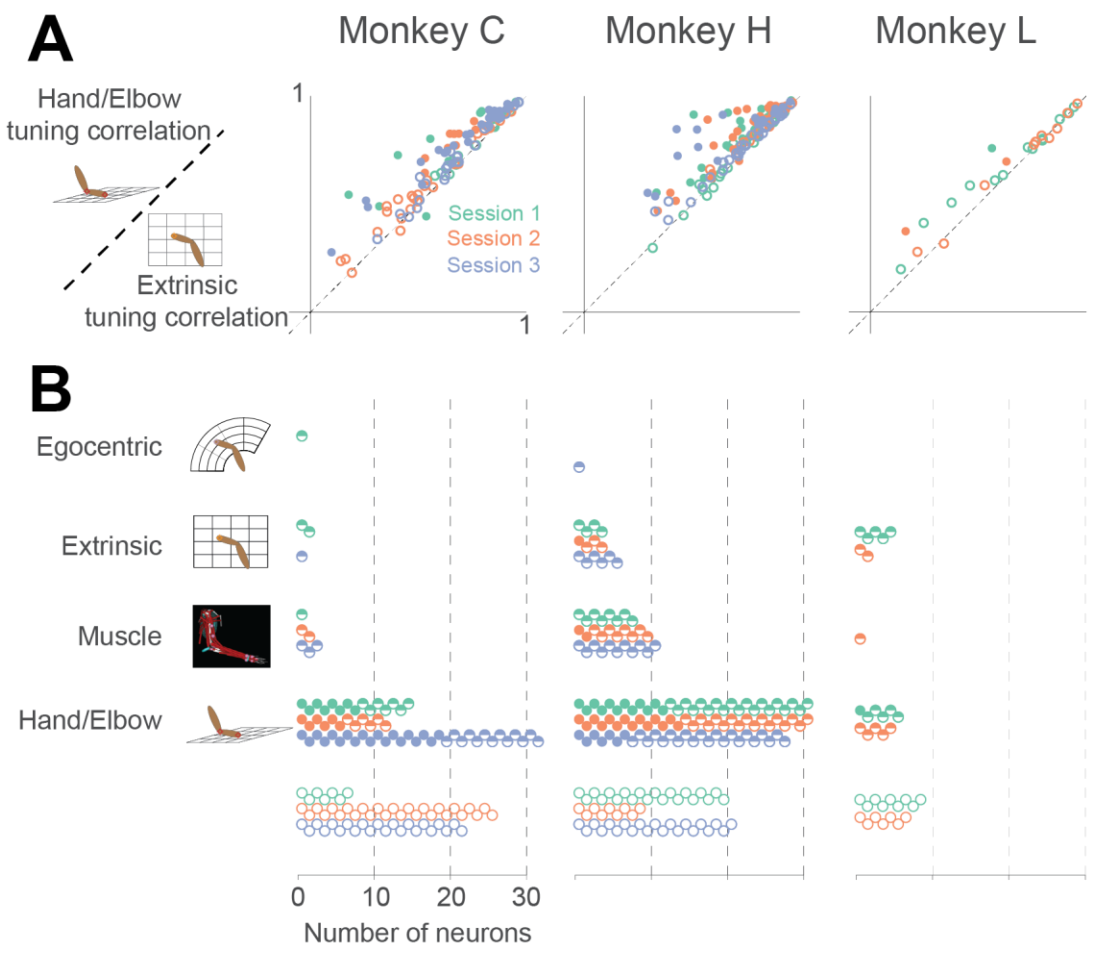

Figure 7: Tuning curve shape correlation analysis, with one column for each monkey, in the same layout as Figure 6. A – Scatter plot for one of the six pairwise comparisons of tuning curve shape correlation. Once again, most filled circles lie above the dashed line of unity, indicating that the hand/elbow model was better at predicting tuning curve shape than the extrinsic model. All six pairwise comparisons for tuning curve correlation are shown in Supplementary Figure 2. B – Dot plot of pairwise comparison winners. Compared to the pseudo-R^2 winner plot (Figure 6B), many more neurons fell into the 'no winner' category for this analysis, where there was no model that won two or more pairwise comparisons. Still, for most of the remaining neurons, the hand/elbow model appeared to predict tuning curve shape better than the other models.

Of the 288 recorded neurons, 260 were significantly tuned to movement direction in both workspaces. This fraction of tuned neurons is much higher than the fraction we found in the active vs. passive task, likely for two reasons. First, unlike with the active vs. passive task, we used the entirety of each trial in the two-workspace experiment to calculate PDs. Second, monkeys tended to move faster during this random target task than during the center-out task, which likely led to higher modulation depths and less measurement error, as predicted by (Stevenson et al. 2011). Thus, in addition to the goodness-of-fit and tuning curve correlation analyses, we were also able to examine the properties of the neural PDs in the two workspaces.

An interesting feature of this task is that for many neurons, the PD of movement changed significantly between workspaces, exemplified by the difference between the vertical bars in the

leftmost panel of Figure 5D. Figure 8A shows the actual PD shifts for these neurons plotted against the PD shifts predicted by each model. The large changes in PD, shown on the vertical axes of the scatter plots are a clue that the extrinsic model does not explain neural activity correctly; if it did, the preferred direction changes should have been insignificant (in principle, zero), as shown by the generally small extrinsic model-predicted changes (second column of Figure 8A). Additionally, and perhaps counterintuitively, the actual changes included both clockwise and counter-clockwise rotations, so it is also unlikely that they arose from a rotation of the extrinsic coordinate frame about the shoulder or the egocentric model. However, we found that the whole-arm models did predict both clockwise and counter-clockwise PD changes. Based on the circular VAF (cVAF) of the PD change prediction, Figure 8B shows that the hand-elbow model once again out-predicted the other models, with hand/elbow having the highest average cVAF (0.73), followed in order by muscle (0.63), extrinsic (0.55), and egocentric (0.38). We made pairwise comparisons between model cVAF for each session. In every session but one, the hand-based models lost pairwise comparisons to at least one of the whole-arm models, and in the outlier session, no pairwise comparisons showed significance (again using $\alpha = 0.05$ and a Bonferroni correction to account for six comparisons for each session).

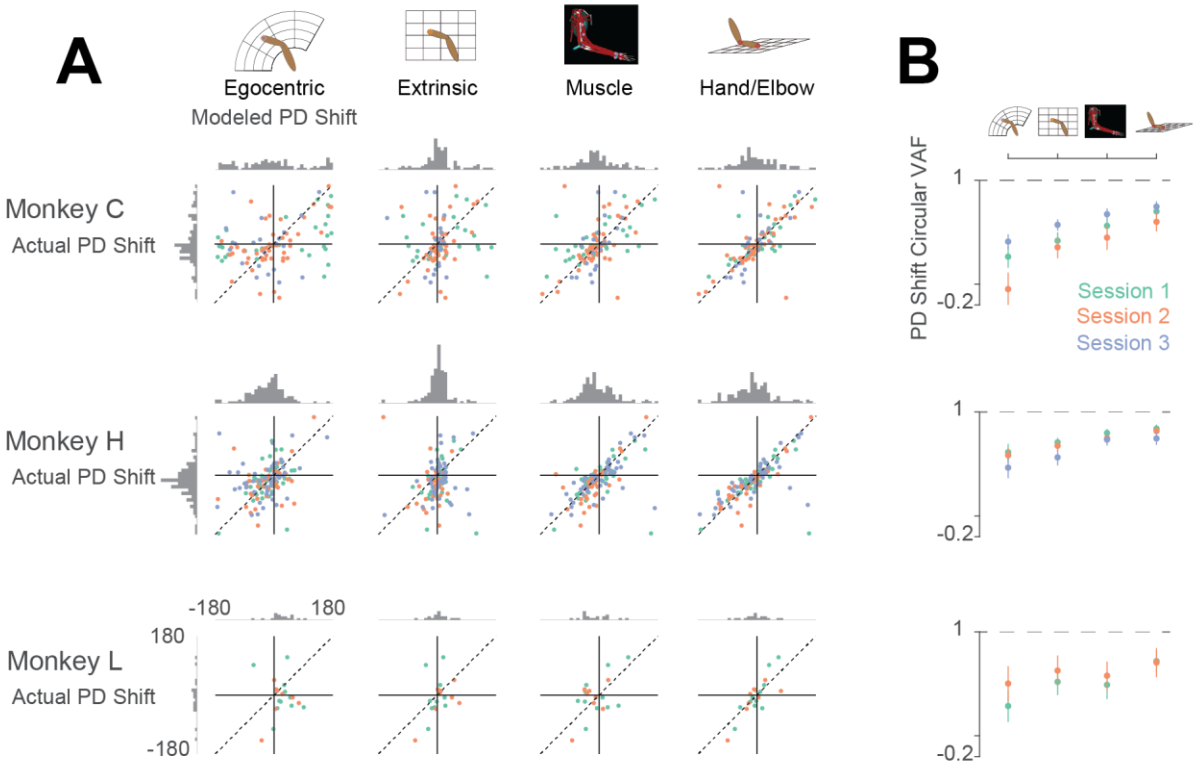

Figure 8: Model predictions of PD shift. A – Scatter plots of actual PD shifts plotted against model-predicted shifts. Each dot represents the actual and modeled PD shifts of a single neuron, where different colors correspond to neurons recorded during different sessions. Dashed diagonal line shows ideal prediction. Vertical histograms indicate distributions of actual PD shifts for each monkey. Horizontal histograms indicate distributions of modeled shifts. Note that both horizontal and vertical axes are circular, meaning that opposing edges of the plot (top/bottom, left/right) are the same. Vertical histograms show that the distribution of actual PD shifts spanned both clockwise and counter-clockwise

shifts. Clustering of scatter plot points on the diagonal line for the hand/elbow model indicates that it was most predictive of PD shift. B – plot showing circular VAF (cVAF) of scatter plots in A, an indicator of how clustered points are around the diagonal line (see Methods for details). Each point corresponds to the average cVAF for a model in a given session (indicated by color), and the horizontal dashed lines indicate the cVAF for perfect prediction. Error bars show 95% confidence intervals (derived from cross-validation – see Methods). Pairwise comparisons between model cVAFs showed that in no sessions did either of the hand-based models significantly out-predict either of the whole-arm models.

As a control for errors introduced into the muscle model by processing marker data with OpenSim, we also performed this cVAF analysis on a hand/elbow model where hand and elbow kinematics were derived from joint angles of the musculoskeletal model, rather than directly from the marker locations captured by the motion tracking system. We re-ran the model prediction analysis for only the muscle model, marker-derived hand/elbow model, and OpenSim-based hand/elbow model. Unsurprisingly, we found average cVAFs similar to those from the main analysis for the marker-derived hand/elbow model (0.75) and the muscle model (0.67). However, the average cVAF for the OpenSim-based hand/elbow model (0.67) dropped to that for the muscle model. This suggests that the difference in predictive capability between the muscle and hand/elbow models stems at least in part from errors introduced in OpenSim modeling, rather than from the hand/elbow model being the better model for area 2 neural activity. Even so, as shown through the main goodness-of-fit (Figure 6, S1), tuning curve correlation (Figure 7, S2), and cVAF (Figure 8) analyses, both of these whole-arm models are better than the hand-based models, supporting the fact that area 2 encodes whole arm kinematics.

3 Discussion

3.1 Summary

In this study, we explored how somatosensory area 2 represents reaching movements using two separate experiments. In the first experiment, we found that although single neuron directional tuning is largely preserved between active and passive movements, the two types of movements can be differentiated in the area 2 neural state space. While the classic hand-centric models cannot explain this separation, we found that a model of neural activity built from the kinematics of both the hand and elbow did. We further explored area 2 encoding of whole-arm kinematics in a second experiment, in which the monkey reached to targets in a pair of workspaces that changed the relationship between hand-based and whole-arm models. Across these two conditions, features of neural activity, including the dynamics of neural discharge, tuning curve shape and preferred direction, were better explained by a model incorporating the kinematics of the whole arm than by the classic hand-based models. Altogether, these results suggest that area 2 of S1 represents movements in terms of not just the hand, but the whole arm.

3.2 Model complexity

A significant difference between the hand and whole-arm models is their number of parameters, which make the whole-arm models more complex and expressible. There are two concerns with

testing models of differing complexity, the first dealing with model training and evaluation, and the second with interpretation of the results.

In training and evaluating our models, we had to make sure that the complex models did not overfit the data, resulting in artificially high performance on the training dataset but low generalizability to new data. However, because we found through cross-validation that the more complex models generalized to test data better than the simpler models, they were not overfitting. Consequently, the hand-based models are clearly impoverished compared to the whole-arm models.

The second concern is in interpreting what it means when the more complex models performed better. One interpretation is that this is an obvious result; if the added degrees of freedom have anything at all to do with area 2 neural activity, then the more complex models should perform better. However, the choice of the two less complex hand-based models was not arbitrary. Both are classic models within the canon of the cortical representation of limb state (Bosco et al. 2000; Prud'homme and Kalaska 1994). Additionally, as shown in the active/passive study, if the classic extrinsic model is to be believed, then the separation between neural representations of active and passive movement must have an explanation not rooted in the behavioral kinematics. Thus, the point of this study was to determine whether either of these now classic models adequately encapsulates what area 2 represents. Because the whole-arm models outperformed the classic models, we conclude that the classic models are incomplete, and in the worst cases, potentially incorrect or misleading.

3.3 Influence of modeling errors

One surprising result from the two-workspace experiment was that the hand/elbow model, based on the addition of one arbitrary point on the proximal limb to the classic hand model, performed as well or better than the muscle-based model. As proprioceptive signals originate in the muscles, arising from muscle spindles and Golgi tendon organs, we expected to find that the muscle model would outperform the other models. However, there are several potential reasons why this was not so. The most important ones can be divided into two categories loosely tied to 1) errors in estimating the musclulotendon lengths, through motion tracking and musculoskeletal modeling, and 2) the fidelity of the muscle model to the actual signals sent by the proprioceptors.

In the first category, the main issue is that of error propagation. The extra stages of analysis required to compute musculotendon lengths (registering markers to a musculoskeletal model, performing inverse kinematics to find joint angles, and using modeled moment arms to estimate musculotendon lengths) introduce errors not present when simply using the positions of markers on the arm. As a control, we ran the hand/elbow model through two of these extra steps by computing the hand and elbow positions from the joint angles of the scaled model, estimated from inverse kinematics. The results of this analysis showed that the performance of the hand/elbow model with added noise dropped to that of the muscle model, indicating that there are, in fact, errors introduced in even this portion of the processing chain.

The other potential source of error in this processing chain stems from the modeled moment arms, which might not accurately reflect those of the actual muscles. In developing their musculoskeletal model, Chan and Moran collected muscle origin and insertion point measurements from both cadaveric studies and existing literature (Chan and Moran 2006). However, due to the complexity of some joints, along with ambiguity of how the muscle wraps around bones and other surfaces, determining moment arms purely by bone and muscle geometry is a difficult problem (An et al. 1984). Because moment arms are irrelevant for determining hand and elbow kinematics, we could not subject the hand/elbow model to the error introduced by this step.

In addition to the questions of error propagation and musculoskeletal model accuracy is the question of whether our muscle model was truly representative of the signals sensed by the proprioceptors. The central complication is that spindles sense the state of the intrafusal fibers in which they reside, and have a complex, nonlinear relation to the musculotendon length that we used in our muscle model. Factors like load-dependent fiber pennation angle (Azizi et al. 2008), or tendon elasticity (Rack and Westbury 1984) can decouple muscle fiber length from musculotendon length. Additionally, intrafusal fibers receive motor drive from gamma motor neurons, which continuously alters muscle spindle sensitivity (Loeb et al. 1985; Prochazka and Wand 1981; Prochazka et al. 1976) and spindle activity also depends on the history of strain on the fibers (Haftel et al. 2004; Proske and Stuart 1985). Altogether, this means that while the musculotendon lengths we computed provide a reasonably good approximation of what the arm is doing, they may not be a good representation of the spindle responses themselves. Spindle activity might be more accurately modeled when given enough information about the musculotendon physiology. However, to model the effects of gamma drive, we would either have to record directly from gamma motor neurons or make assumptions of how gamma drive changes over the course of reaching. In developing models of neural activity, one must carefully consider the tradeoff between increased model complexity and the extra error introduced by propagating through the additional requisite measurement and analysis steps. Given our data obtained by measuring the kinematics of the arm with motion tracking, it seems that the coordinate frame with which to best explain area 2 neural activity is simply the one with the most information about the arm kinematics and the fewest steps in processing. However, this does not rule out the idea that area 2 more nearly represents a different whole-arm model that may be less abstracted from physiology, like musculotendon length or muscle spindle activity.

3.4 Coordinate frame vs. informational content

Because of their differing dimensionality, the signals from hand-based models and those from whole-arm models do not have a one-to-one relationship: there are many different arm configurations that result in a given hand position. Thus, a comparison between the hand-based and whole-arm models is mainly a question of information content (do area 2 neurons have information about more than just the hand?). In contrast, signals of the two whole-arm models do have a one-to-one (albeit nonlinear) relationship to each other. Knowledge of the hand and elbow position should completely determine the estimated musculotendon lengths, indicating that the two models have the same informational content. As such, a comparison between the muscle

model and the hand/elbow is purely one of coordinate frame. While the interpretation for a comparison of information content is straightforward, interpreting the results of a comparison between coordinate frames is not. One major issue is that these comparisons only make sense when using linear models to relate neural activity to behavior. Once nonlinear models are considered, as in our study with artificial neural networks (Lucas et al. 2019), coordinate frames with one-to-one correspondence become nearly equivalent, and much more difficult to compare meaningfully.

Clear parallels exist between this and earlier studies seeking to find a unique representation of movement in motor areas. Over the last few decades, a controversy involving the exact nature of the neural representation of movement has played itself out in the literature surrounding motor cortex, with some advocating a hand-based representation of motor control (Georgopoulos et al. 1982; Georgopoulos et al. 1986; Moran and Schwartz 1999) and others a muscle-based representation (Evarts 1968; Fetz et al. 1989; Morrow et al. 2007; Oby et al. 2012). Recently, the motor control field started turning away from questions of coordinate frame and towards questions of neural population dynamics and information processing (Churchland et al. 2010; Elsayed et al. 2016; Gallego et al. 2017; Kaufman et al. 2014; Perich et al. 2018; Russo et al. 2018; Sussillo et al. 2015). Part of the motivation for this pivot in viewpoint is that it became increasingly clear that a “pure” coordinate frame of movement representation is unlikely to exist (Fetz 1992; Kakei et al. 1999). Further, studies tended to use correlation between neural activity and behavioral variables as evidence that the neurons represent movements in a particular coordinate frame. However, as noted above, these correlations could often be explained by multiple coordinate frames, casting doubt on the conclusiveness of the exact coordinate frame of representation (Mussa-Ivaldi 1988). Consequently, in our study, we put aside the question of the coordinate frame of area 2, focusing instead on the fact that area 2 contains information not included in the classic model for proprioceptive neural activity during reaching; it relates to movements of the whole arm, rather than just the hand.

A major question this study leaves open is that of how information about reaching is processed by different areas of the proprioceptive neuraxis. While we might expect a muscle spindle-like representation at the level of the dorsal root ganglia (DRG) or the cuneate nucleus, removed from the receptors by one and two synapses, respectively, this representation likely changes as the signals propagate through thalamus and into S1. Even different areas of S1 may have different representations. Area 3a, which receives input mostly from muscle afferents (Heath et al. 1976; Kaas et al. 1979; Phillips et al. 1971; Yamada et al. 2016), seems more likely to retain a muscle-like representation than is area 2, which integrates muscle afferent input with that from cutaneous receptors (Hyvärinen and Poranen 1978; Padberg et al. 2018; Pons et al. 1985). Likewise, area 5 may have an even higher-level representation, as it receives input from both somatosensory (Mountcastle et al. 1975) and motor cortices (Padberg et al. 2018), and appears to depend on attention (Chapman et al. 1984; Omrani et al. 2016). As it becomes increasingly feasible to record from several of these areas simultaneously (Richardson et al. 2016; Suresh et al. 2017; Weber et al. 2006), future experiments could examine how these areas project information to each other, as has been explored in motor and premotor cortices (Churchland et al. 2010; Elsayed et al. 2016; Kaufman et al. 2014; Perich et al. 2018), without modeling the

more complex cortical areas explicitly in terms of particular behavioral variables “encoded” by single neurons.

This study shows that even after proprioceptive signals reach area 2, neural activity can still be predicted well by a convergence of muscle-like signals, even though the signals have been processed by several sensory areas along the way. One potential explanation for this is that at each stage of processing, neurons simply spatially integrate information from many neurons of the previous stage, progressively creating more complex response properties. This idea of hierarchical processing was first used to explain how features like edge detection and orientation tuning might develop within the visual system from spatial integration of the simpler photoreceptor responses (Felleman and Van Essen 1991; Hubel and Wiesel 1959; 1962). This inspired the design of deep convolutional artificial neural networks, now the state of the art in machine learning for image classification (Krizhevsky et al. 2012). Unlike previous image recognition methods, these feedforward neural networks are not designed to extract specific, human-defined features of images. Instead, intermediate layers learn to integrate spatially patterned information from earlier layers to build a library of feature detectors. In the proprioceptive system, such integration, without explicit transformation to some intermediate movement representation, might allow neurons in area 2 to serve as a general-purpose library of limb-state features, whose activity is read out in different ways for either perception or use in motor control.

3.5 Relevance for BCI

One motivation for this work is its potential to augment brain-computer interfaces (BCI) for restoring movement to persons with spinal cord injury or limb amputation. As BCI for motor control gets more advanced (Collinger et al. 2013; Ethier et al. 2012; Kao et al. 2015; Young et al. 2018), it will become more necessary to develop a method to provide feedback about movements to the brain, potentially using intracortical microstimulation (ICMS) to activate somatosensory areas. While ICMS in S1 has seen some success in providing feedback about touch (Flesher et al. 2016; Romo et al. 1998; Salas et al. 2018; Tabot et al. 2013), the path towards providing proprioceptive feedback remains relatively unexplored. At least one study did use electrical stimulation in S1 for feedback during movement, using the stimulation to specify target direction with respect to the evolving hand position (Dadarlat et al. 2015). In that study, monkeys used the ICMS to reach to targets, even in the absence of visual feedback. However, target-location information is very different from the information normally encoded by S1, and the monkeys required several months to learn to use it. To our knowledge, no study has yet shown a way to use ICMS to provide more biomimetic proprioceptive feedback during reaching. Previously, our lab attempted to address this gap by stimulating a small number of electrodes in area 2 based on neural activity recorded from them during normal reaching movements. In that experiment, the monkey reported the direction of a mechanical bump to his arm that occurred simultaneously with the ICMS. The ICMS biased one monkey’s reports of the mechanical bump direction toward the PDs of the stimulated electrodes. Key to this finding was the fact that any bias in reporting actually decreased the reward rate, suggesting that the ICMS was indistinguishable from the perception of the bump itself (Tomlinson and Miller 2016).

Unfortunately, the result could not be replicated in other monkeys; while the ICMS often biased their reports, the direction of the bias could not be explained by the PDs of the stimulated electrodes. One potential reason may be that the stimulation paradigm in those experiments was derived from the classic, hand-based model and the assumption that area 2 represents active and passive movements similarly. As this paper has shown, both of these assumptions have important caveats. It is possible that a stimulation paradigm based on a whole-arm model may be more successful, due to its greater accuracy at predicting neural PDs (Figure 8). It is also possible that the stimulus model would need to include information about forces in addition to kinematics. Regardless of the exact model, prospects for stimulating S1 to create natural proprioceptive sensations would likely improve given a more accurate generative model of S1 activity.

In addition to developing better models for S1 activity, it will be important to consider the implications of the difference between sensation for perception versus action. These two broad purposes for sensation are thought to involve distinct pathways in both vision and touch (Dijkerman and De Haan 2007; Mishkin and Ungerleider 1982; Sedda and Scarpina 2012). It is quite plausible that this distinction exists for proprioception as well (Dijkerman and De Haan 2007). However, studies of the effects of ICMS in S1 tend to use perceptual reporting to test the effect of stimulation (Salas et al. 2018; Tomlinson and Miller 2016; Zaaimi et al. 2013), thereby not directly addressing how effectively ICMS can be used as feedback for action. Even in the study conducted by Dadarlat et al., movements guided by ICMS were slower and contained more submovements that those guided by even a noisy visual signal, suggesting that monkeys used the ICMS as a learned sensory substitute, rather than as a biomimetic replacement for proprioception. As such, that study was also likely a cognitive one, engaging the perceptual stream rather than the action stream of proprioception (see (Deroy and Auvray 2012; Elli et al. 2014) for discussion of the limits of sensory substitution). As we better characterize how S1 represents movements, we hope to develop a stimulation paradigm in which we can engage both streams, to enable users of a BCI both to perceive their limb, and to respond rapidly to movement perturbations.

4 Conclusion

This study began with an observation: the classic, hand-based cortical model of proprioception could not explain the separability of active and passive movements we observed in the neural state space within area 2. We found, however, that this feature could be explained by extending the classic model to include the kinematics of the whole arm. In a second experiment, we found that predictions of area 2 neural activity from such whole-arm models generalized to different behavioral conditions better than those of the classic model. This suggests that even though our perception of our arm is typically centered on the hand, area 2 of S1 still appears to represent movement of the whole arm.

5 Acknowledgements

We would like to thank Brian London for initial discussions of the active vs. passive result and Tucker Tomlinson, Christopher VerSteeg, and Joseph Sombeck for their help with training and caring for the research animals. Additionally, we would like to thank them, along with Matt Perich, Juan Gallego, Sara Solla, and the entire Miller Limb Lab for discussions and feedback that greatly improved this work.

6 Grants

This research was funded by National Institute of Neurological Disorders and Stroke Grant No. NS095251 and National Science Foundation Grant No. DGE-1324585.

7 Methods and Materials

All surgical and experimental procedures were fully consistent with the guide for the care and use of laboratory animals and approved by the institutional animal care and use committee of Northwestern University under protocol #IS00000367.

7.1 Behavior

We recorded data from a monkey while it used a manipulandum to reach for targets presented on a screen within a 20 cm x 20 cm workspace. After each successful reaching trial, the monkey received a pulse of juice or water as a reward. We recorded the position of the handle using encoders on the manipulandum joints. We also recorded the interaction forces between the monkey's hand and the handle using a six-axis load cell mounted underneath the handle.

For the active vs. passive experiment, we had the monkey perform a classic center-out (CO) reaching task, as described in (London and Miller 2013). Briefly, the monkey held in a target at the center of the full workspace for a random amount of time, after which one of four outer targets was presented. The trial ended in success once the monkey reached to the outer target. On 50% of the trials (deemed "passive" trials), during the center hold period, we used motors on the manipulandum to deliver a 2 N perturbation to the monkey's hand in one of the four target directions. After the bump, the monkey returned to the center target, after which the trial proceeded like an active trial. From only the successful passive and active trials, we analyzed the first 120 ms after movement onset. Movement onset was determined by looking for the peak in handle acceleration either after the motor pulse (in the passive condition) or after 200 ms post-go cue (in the active condition) and sweeping backwards in time until the acceleration was less than 10% of the peak.

For the two-workspace experiment, we partitioned the full workspace into four 10cm x 10cm quadrants. Of these four quadrants, we chose the far ipsilateral one and the near contralateral one in which to compare neural representations of movement. Before each trial, we chose one of the

two workspaces randomly, within which the monkey reached to a short sequence of targets randomly positioned in the workspace. For this experiment, we only analyzed the portion of data from the end of the center-hold period to the end of the trial.

7.2 Motion tracking

Before each reaching experiment, we painted 10 markers on the outside of the monkey's arm, marking bony landmarks and a few points in between, a la Chan and Moran 2006. Using a custom motion tracking system built from a Microsoft Kinect, we recorded the 3D locations of these markers with respect to the camera, synced in time to the other behavioral recordings. We then aligned the Kinect-measured marker locations to the lab frame by aligning location of the Kinect hand marker to the location of the handle in the manipulandum coordinate frame. Code for motion tracking can be found at *https://github.com/limblab/KinectTracking.git*.

7.3 Musculoskeletal modeling

We registered the Kinect marker locations to a monkey arm musculoskeletal model in OpenSim (SimTK), based on a model published by (Chan and Moran 2006). After scaling the limb segments of the model to match those of each monkey, we used the inverse kinematics analysis tool provided by OpenSim to estimate the joint angles (and corresponding muscle lengths) required to match the model's virtual marker positions to the positions of the actual recorded markers. The OpenSim model we used can be found at *https://github.com/limblab/monkeyArmModel.git*.

7.4 Neural recordings

We implanted 100-electrode arrays (Blackrock Microsystems) into the arm representation of area 2 of S1 in these monkeys. For more details on surgical techniques, see (Weber et al. 2011). In surgery, we roughly mapped S1 by recording from the cortical surface while manipulating the arm and hand to localize their representations. To record neural data for our experiments, we used a Cerebus recording system (Blackrock). This recording system sampled signals from each of the 96 electrodes at 30 kHz. To conserve data storage space, the system detected spikes online using a threshold set at -5 x signal RMS, and only wrote to disk a time stamp and the 1.6 ms snippet of signal surrounding the threshold crossing. After data collection, we used Plexon Offline Sorter to manually sort these snippets into putative single units, using features like waveform shape and inter-spike interval. In addition to these recording sessions, we also occasionally performed sensory mapping sessions to identify the neural receptive fields by manipulating the monkey's arm while listening to neural activity. In all monkeys, we found a roughly equal mix of cutaneous and deep (muscle) receptive fields, suggesting that we were recording primarily from area 2 (Hyvärinen and Poranen 1978; Padberg et al. 2018; Pons et al. 1985; Seelke et al. 2011). It is possible that the more posterior regions of the array were at least in transition regions of area 5.

7.5 Neural analysis

Code for the following neural analyses can be found at *https://github.com/raeedcho/s1-kinematics.git*.

7.5.1.1 Preferred directions

We used a simple bootstrapping procedure to calculate PDs for each neuron. On each bootstrap iteration, we randomly drew timepoints from the reaching data, making sure that the distribution of movement directions was uniform to mitigate the effects of any potential bias. Then, as in (Georgopoulos et al. 1982), we fit a cosine tuning function to the neural activity with respect to the movement direction, using equations 1a-b.

$$f_i(\tau) = b_0 + b_1 * \sin\big(\theta_m(\tau)\big) + b_2 * \cos(\theta_m(\tau)) \qquad (1a)$$

$$= b_0 + r_i * \cos(\theta_m(\tau) - PD_i) \qquad (1b)$$

where

$$PD_i = atan2(b_1, b_2) \text{ and } r_i = sqrt(b_1^2 + b_2^2)$$

Here, $f_i(\tau)$ is the average firing rate of neuron i for a given time point τ, and $\theta_m(\tau)$ is the corresponding movement direction, which for the active/passive task was the target or bump direction, and for the two-workspace experiment was the average movement direction over a time bin. We took the circular mean of PD_i and mean of r_i over all bootstrap iterations to determine the preferred direction and the modulation depth respectively, for each neuron.

As the PD analysis is meaningless for neurons that don't have a preferred direction of movement, we only analyzed the PDs of neurons that were significantly tuned. We assessed tuning through a separate bootstrapping procedure, described in (Dekleva et al. 2018). Briefly, we randomly sampled the timepoints from reaching data, again ensuring a uniform distribution of movement directions, but this time also randomly shuffled the corresponding neural activity. We calculated the r_i for this shuffled data on each bootstrap iteration, thereby creating a null distribution of modulation depths. We considered a neuron to be tuned if the true r_i was greater than the 95th percentile of the null distribution.

7.5.2 Models of neural activity

For the active/passive analyses, we averaged behavioral variables and neural firing rates over the 120 ms period following movement onset in each trial. For the two-workspace analyses, both behavioral variables and neural firing rate were averaged over 50 ms bins. We modeled neural activity with respect to the behavior using Poisson generalized linear models (outline in (Truccolo et al. 2005)) shown in equation 2a, below.

$$f \sim Poisson(\lambda), \lambda = \exp(X\beta) \qquad (2a)$$

In this equation, f is a T (number of time points) x N (number of neurons) matrix of average firing rates, X is a T x P (number of behavioral covariates, explained below) matrix of behavioral correlates, and β is a P x N matrix of model parameters. We fit these GLMs by finding

maximum likelihood estimation of the parameters, $\hat{\beta}$. With these fitted models, we predicted firing rates ($\hat{f}$) on data not used for training, shown in equation 2b, below.

$$\hat{f} = \exp(X\hat{\beta}) \qquad (2b)$$

We tested five firing rate encoding models, detailed below. Note that each model also includes an offset term, increasing the number of parameters, P, by one.

- Extrinsic kinematics: behavioral covariates were position and velocity of the hand, estimated by using the location of one of the hand markers, in three-dimensional Cartesian space, with origin at the shoulder ($P = 7$).
- Extrinsic kinematics+force: behavioral covariates were position and velocity of the hand, as well as forces on the hand, in three-dimensional Cartesian space (P =10).
- Egocentric kinematics: behavior covariates were position and velocity of the hand marker in spherical coordinates (θ, ϕ, and ρ), with origin at the shoulder ($P = 7$).
- Hand/elbow kinematics: behavior covariates were position and velocity of both the hand and elbow markers in three-dimensional Cartesian space, with origin at the shoulder. This is the simplest extension of the extrinsic model that incorporates information about the configuration of the whole arm ($P = 13$)
- Muscle kinematics: behavioral covariates were derived from the length of the 39 modeled muscles (Chan and Moran 2006) and their time derivatives. However, because this would result in almost 78 (highly correlated) covariates, we used PCA to extract 5-dimensional orthogonal basis sets for both the lengths and their derivatives. On average, five components explained 99 and 96 percent of the total variance of lengths and length derivatives, respectively. Behavioral covariates of this model were the projections of the muscle variables into these spaces during behavior ($P = 11$).

We used repeated 5-fold cross-validation to evaluate our models of neural activity, given that the models had different numbers of parameters, P. On each repeat, we randomly split trials into five groups (folds) and trained the models on four of them. We used these trained models to predict neural firing rates ($\widehat{f_i}$) in the fifth fold. We then compared the predicted firing rates from each model to the actual firing rates in that test fold, using analyses described in the following sections. This process (including random splitting) was repeated 20 times, resulting in n=100 sample size for each analysis result. Thus, if a more expressive model with more parameters performs better than a simpler model, it would suggest that the extra parameters do provide relevant information about the neural activity not accounted for by the simpler models.

7.5.3 Statistical tests and confidence intervals

To perform statistical tests on the output of repeated 5-fold cross-validation, we used a corrected resampled t-test, outlined in (Ernst 2017) and (Nadeau and Bengio 2003). Here, sample mean and variance are calculated as in a normal t-test, but a correction factor needs to be applied to the standard error, depending on the nature of the cross-validation. Equation 3a-c shows a general case of this correction for R repeats of K-fold cross-validation of some analysis result d_{kr}.

$$\hat{\mu}_d = \frac{1}{K \times R} \sum_{k=1}^{K} \sum_{r=1}^{R} d_{kr} \qquad (3a)$$

$$\hat{\sigma}_d^2 = \frac{1}{(K \times R) - 1} \sum_{k=1}^{K} \sum_{r=1}^{R} (\, d_{kr} - \widehat{\mu_d})^2 \qquad (3b)$$

$$t_{stat} = \frac{\hat{\mu}_d}{\sqrt{(\frac{1}{K \times R} + \frac{1/K}{1 - 1/K})\hat{\sigma}_d^2}} \qquad (3c)$$

We then compare the t-statistic here (t_{stat}) to a t-distribution with $K \times R - 1$ degrees of freedom. Note that the correction applied is an extra term (i.e., $\frac{1/K}{1-1/K}$) under the square root, compared to the typical standard error calculation.

7.5.4 Active/passive analyses

7.5.4.1 Estimating neural state-space separability

The neural response to movement, whether active or passive, can be represented as a single datapoint in a neural population space defined by the activity of all neurons in the relevant 120 ms period following movement onset. As described above, we estimated the separability of the active and passive movements in neural space using 20x repeated, 5-fold cross-validation. We did this in three steps. First, for each training set, we characterized the population response to each trial by finding the first three modes of neural activity using principal component analysis (PCA). We then projected the neural activity onto these three principal components and trained a linear discriminant analysis (LDA) model to find the axis of maximal separation between active and passive trials. Finally, we sequentially projected each test fold's neural data into the PC space and then onto the LDA axis. This resulted in a scalar value for each trial, with the sign indicating whether LDA classified the trial as active or passive. We took the average classification accuracy of the test fold's data as the percent separability for the fold, giving us 100 total samples from the 20x5-fold cross-validation. By averaging these samples, we estimated the overall neural separability of active and passive movements in a given session.

7.5.4.2 Estimating model-predicted separability

We also trained encoding models to predict neural firing from behavior (see equation 2a for procedure) using three different models of neural activity: the extrinsic kinematics, extrinsic kinematics+force, and hand/elbow kinematics models. See "Models of neural activity" section for more details on the specific models. After training each model over the four training folds, we estimated firing rates in both the training and the test folds. Subsequent analysis mirrored that of the actual neural data: we found the three leading PCs and LDA axis of highest separability in the training folds and then sequentially projected the test-fold data through the PC space and onto the LDA axis. This resulted in 100 samples with which to estimate the model-predicted separability of active and passive movements.

7.5.4.3 Neural space dimensionality reduction

To visualize the population neural activity for figures, we used a combination of LDA and PCA. For the horizontal axis, we used LDA to find an axis in the three PCs of neural population space along which active and passive trials were most separated. For the vertical axis, we projected all activity onto the hyperplane orthogonal to the LDA axis and used PCA again to find the remaining axis of highest variance.

7.5.5 Two-workspace analyses

These analyses examined how well models of neural activity could predict neural activity as the monkey reached to targets in different workspaces. As such, we analyzed firing rate goodness-of-fit, along with how well the models could replicate the tuning curves and preferred directions (PDs) of neurons.

7.5.5.1 Goodness-of-fit

We evaluated goodness-of-fit of these models for each neuron by using a pseudo-R^2 (pR^2) metric. We used a formulation of pseudo-R^2 based on a comparison between the deviance of the full model and the deviance of a "null" model, i.e., a model that only predicts the overall mean firing rate (Cameron and Windmeijer 1997; 1996; Heinzl and Mittlböck 2003; Perich et al. 2018).

$$pR^2 = 1 - \frac{D\left(f_i; \widehat{f_i}\right)}{D\left(f_i; \overline{f_i}\right)} \qquad (4a)$$

$$= 1 - \frac{logL(f_i) - logL\left(\widehat{f_i}\right)}{logL(f_i) - logL\left(\overline{f_i}\right)} \qquad (4b)$$

When computing the likelihood of a Poisson statistic, this is:

$$= 1 - \frac{\sum_{\tau=1}^{T} f_i\,(\tau) \log\left(\frac{f_i(\tau)}{\widehat{f_i}(\tau)}\right) - \left(f_i(\tau) - \widehat{f_i}(\tau)\right)}{\sum_{\tau=1}^{T} f_i\,(\tau) \log\left(\frac{f_i(\tau)}{\overline{f_i}}\right) - \left(f_i(\tau) - \overline{f_i}\right)} \qquad (4c)$$

This pR^2 metric ranges from $-\infty$ to 1, with a value of 1 corresponding to a perfectly fit model and a value of 0 corresponding to a model that only fits as well as the "null" model. In contrast with the general intuition for regular R^2, a pR^2 of ~0.2 is considered a "good" fit (McFadden 1977).

7.5.5.2 Tuning curves

We binned the trajectory into 16 bins, each 22.5 degrees wide, based on the mean direction across 50 ms of hand motion. For each directional bin, we calculated the sample mean and 95% confidence interval of the mean. In figures, we plotted this mean firing rate against the center-point of the bin.

7.5.5.3 Preferred direction shift

We calculated PDs for each neuron in each workspace and found the predicted change in PD from the contralateral workspace to the ipsilateral workspace, given each model. We compared these changes to those observed for each neuron. The values of these PD shifts are shown in Figure 8 for all neurons tuned to movements in both workspaces, averaged over all 100 test folds.

We computed a variance-accounted-for (VAF) metric, here called the "circular VAF" (cVAF) for each neuron (i) in each fold as:

$$cVAF_i = cos(\Delta\theta_{PD,i} - \Delta\hat{\theta}_{PD,i}) \qquad (5)$$

As the cVAF metric is essentially the inner product of unit vectors with direction $\Delta\theta_{PD,i}$ and $\Delta\hat{\theta}_{PD,i}$, it accounts for the circular domain of the PD shifts. Like regular VAF, the cVAF has a maximum value of 1 when $\Delta\theta_{PD,i}$ and $\Delta\hat{\theta}_{PD,i}$ are the same, and decreases in proportion to the squared difference between $\Delta\theta_{PD,i}$ and $\Delta\hat{\theta}_{PD,i}$. We took the average cVAF over all neurons as the cVAF for the fold. In total, given the 20 repeats of 5-fold cross-validation, this gave us 100-samples of the cVAF for each model in a given session.

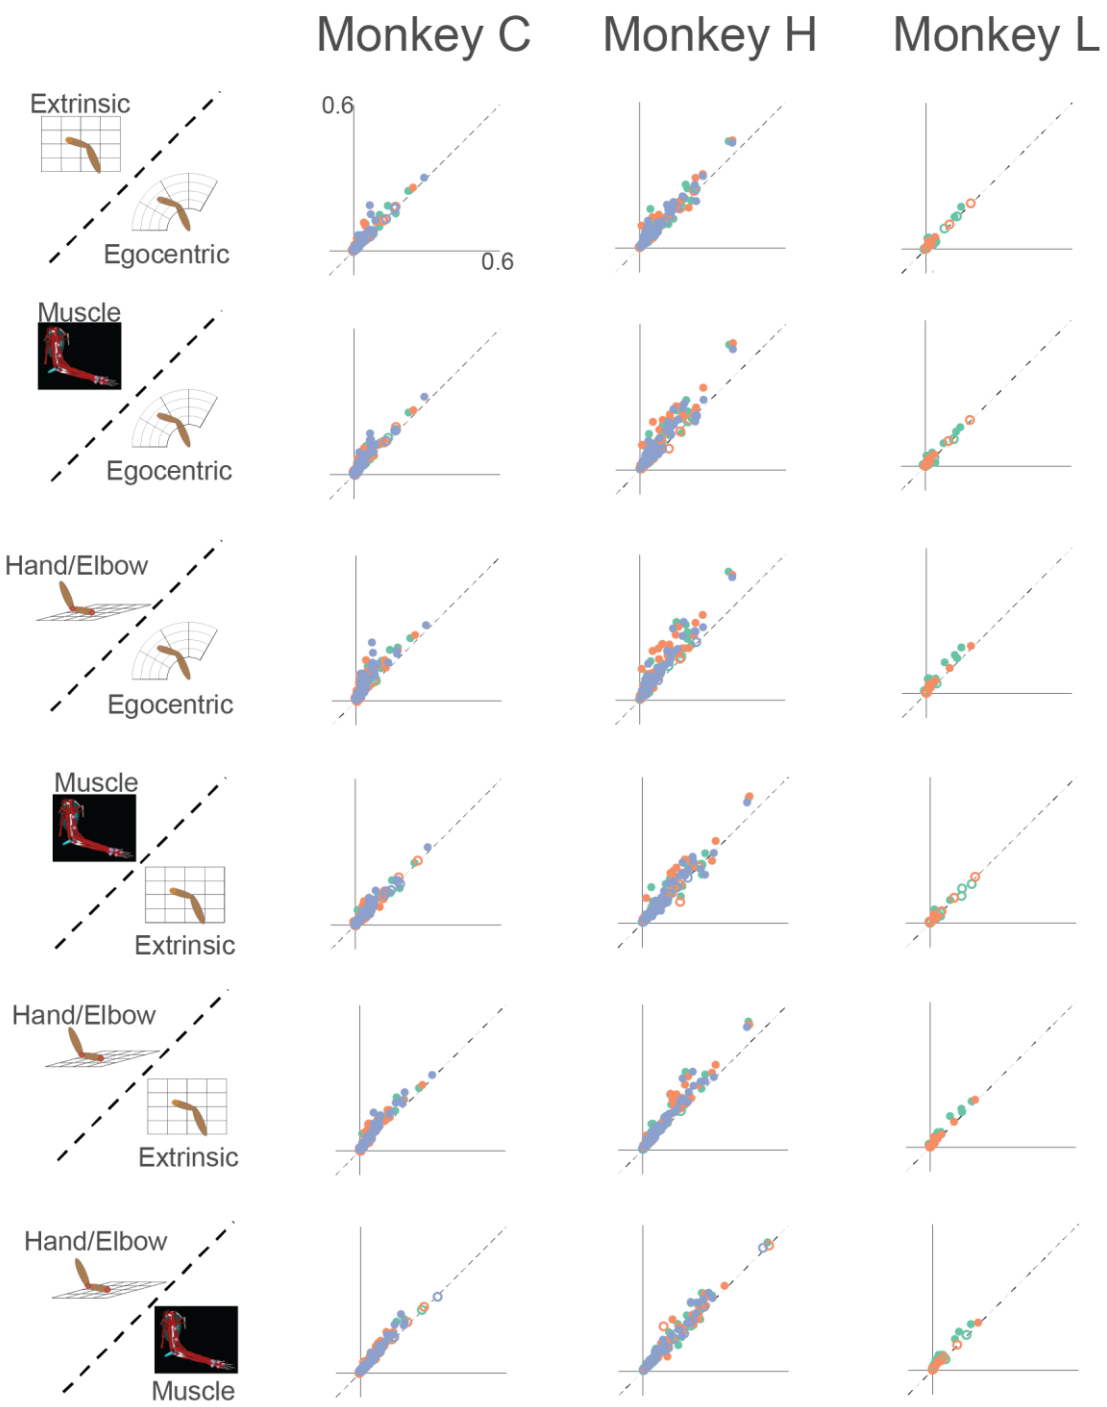

Figure S1: All pairwise comparisons of model pseudo-R2 values. Each column corresponds to one of the monkeys, and each row corresponds to the comparison between two models, indicated by the comparison label on the right. Dots in the scatter plots correspond to neurons, with pseudo-R2 values for one model plotted against those for the other model. Different colors correspond to different sessions, and filled circles correspond to neurons for which one of the two models is significantly better than the other one (i.e. significantly different from the dashed line of unity). Open circles correspond to neurons for which neither model is significantly better than the other one.

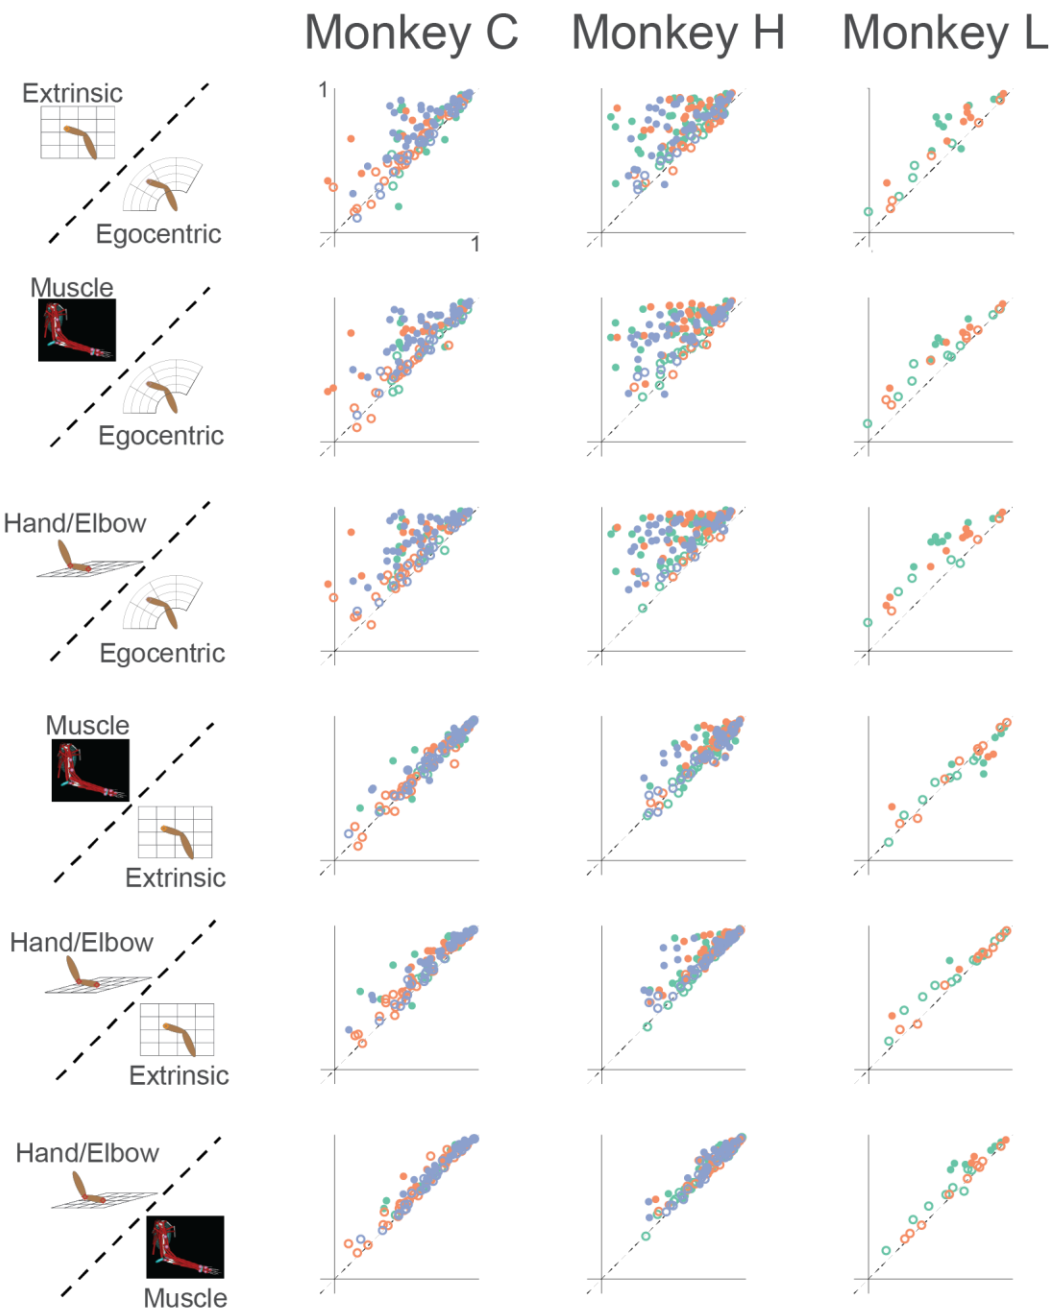

Figure S2: All pairwise comparisons between models of tuning curve shape correlation. Same format as in Figure S1.

An K, Takahashi K, Harrigan T, and Chao E. Determination of muscle orientations and moment arms. *Journal of biomechanical engineering* 106: 280-282, 1984.

Azizi E, Brainerd EL, and Roberts TJ. Variable gearing in pennate muscles. *Proceedings of the National Academy of Sciences* 105: 1745-1750, 2008.

Bell CC. An efference copy which is modified by reafferent input. *Science* 214: 450-453, 1981.

Bosco G, and Poppele RE. Proprioception from a spinocerebellar perspective. *Physiol Rev* 81: 539-568, 2001.

Bosco G, Poppele RE, and Eian J. Reference frames for spinal proprioception: limb endpoint based or joint-level based? *J Neurophysiol* 83: 2931-2945, 2000.

Bosco G, Rankin A, and Poppele R. Representation of passive hindlimb postures in cat spinocerebellar activity. *J Neurophysiol* 76: 715-726, 1996.

Cameron AC, and Windmeijer FA. An R-squared measure of goodness of fit for some common nonlinear regression models. *Journal of econometrics* 77: 329-342, 1997.

Cameron AC, and Windmeijer FA. R-squared measures for count data regression models with applications to health-care utilization. *Journal of Business & Economic Statistics* 14: 209-220, 1996.

Caminiti R, Johnson PB, and Urbano A. Making arm movements within different parts of space: Dynamic aspects in the primate motor cortex. *J Neurosci* 10: 2039-2058, 1990.

Chan SS, and Moran DW. Computational model of a primate arm: from hand position to joint angles, joint torques and muscle forces. *Journal of neural engineering* 3: 327, 2006.

Chapman C, Spidalieri G, and Lamarre Y. Discharge properties of area 5 neurones during arm movements triggered by sensory stimuli in the monkey. *Brain Research* 309: 63-77, 1984.

Chowdhury RH, Tresch MC, and Miller LE. Musculoskeletal geometry accounts for apparent extrinsic representation of paw position in dorsal spinocerebellar tract. *Journal of Neurophysiology* 2017.

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

Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, and Shenoy KV. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron* 68: 387-400, 2010.

Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJC, Velliste M, Boninger ML, and Schwartz AB. High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet* 381: 557-564, 2013.

Cunningham JP, and Byron MY. Dimensionality reduction for large-scale neural recordings. *Nature neuroscience* 17: 1500-1509, 2014.

Dadarlat MC, O'Doherty JE, and Sabes PN. A learning-based approach to artificial sensory feedback leads to optimal integration. *Nat Neurosci* 18: 138-144, 2015.

Dekleva BM, Kording KP, and Miller LE. Single reach plans in dorsal premotor cortex during a two-target task. *Nature communications* 9: 3556, 2018.

Deroy O, and Auvray M. Reading the world through the skin and ears: A new perspective on sensory substitution. *Frontiers in psychology* 3: 457, 2012.

Dijkerman HC, and De Haan EH. Somatosensory processing subserving perception and action: Dissociations, interactions, and integration. *Behavioral and brain sciences* 30: 224-230, 2007.

Elli GV, Benetti S, and Collignon O. Is There a Future for Sensory Substitution Outside Academic Laboratories? 27: 271, 2014.

Elsayed GF, Lara AH, Kaufman MT, Churchland MM, and Cunningham JP. Reorganization between preparatory and movement population responses in motor cortex. *Nature communications* 7: 13239, 2016.

Ernst D. Get CI and p-values for cross validated performance measures. Cross Validated: 2017.

Ethier C, Oby ER, Bauman MJ, and Miller LE. Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature* 485: 368-371, 2012.

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

Felleman DJ, and Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1: 1-47, 1991.

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

Fetz EE, Cheney PD, Mewes K, and Palmer S. Control of forelimb muscle activity by populations of corticomotoneuronal and rubromotoneuronal cells. In: *Progress in Brain Research:Afferent Control of Posture and Locomotion*, edited by Allum JHJ, and Hulliger M. Amsterdam: Elvsevier Science, 1989, p. 437-449.

Flesher SN, Collinger JL, Foldes ST, Weiss JM, Downey JE, Tyler-Kabara EC, Bensmaia SJ, Schwartz AB, Boninger ML, and Gaunt RA. Intracortical microstimulation of human somatosensory cortex. *Science translational medicine* 8: 361ra141-361ra141, 2016.

Fuentes CT, and Bastian AJ. Where Is Your Arm? Variations in Proprioception Across Space and Tasks. *Journal of Neurophysiology* 103: 164-171, 2010.

Gallego JA, Perich MG, Miller LE, and Solla SA. Neural Manifolds for the Control of Movement. *Neuron* 94: 978-984, 2017.

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

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

Ghez C, and Sainburg R. Proprioceptive control of interjoint coordination. *CanJPhysiolPharmacol* 73: 273-284, 1995.

Goodwin GM, McCloskey DI, and Matthews PB. Proprioceptive illusions induced by muscle vibration: contribution by muscle spindles to perception? *Science* 175: 1382-1384, 1972.

Gordon J, Ghilardi MF, and Ghez C. Impairments of reaching movements in patients without proprioception. I. Spatial Errors. *Journal of Neurophysiology* 73: 347-360, 1995.

Haftel VK, Bichler EK, Nichols TR, Pinter MJ, and Cope TC. Movement reduces the dynamic response of muscle spindle afferents and motoneuron synaptic potentials in rat. *Journal of neurophysiology* 91: 2164-2171, 2004.

Heath C, Hore J, and Phillips C. Inputs from low threshold muscle and cutaneous afferents of hand and forearm to areas 3a and 3b of baboon's cerebral cortex. *The Journal of physiology* 257: 199-227, 1976.

Heinzl H, and Mittlböck M. Pseudo R-squared measures for Poisson regression models with over-or underdispersion. *Computational statistics & data analysis* 44: 253-271, 2003.

Hubel DH, and Wiesel TN. Receptive fields of single neurones in the cat's striate cortex. *J Physiol* 148: 574-591, 1959.

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

Hyvärinen J, and Poranen A. Receptive field integration and submodality convergence in the hand area of the post-central gyrus of the alert monkey. *The Journal of Physiology* 283: 539, 1978.

Jennings VA, Lamour Y, Solis H, and Fromm C. Somatosensory cortex activity related to position and force. *J Neurophysiol* 49: 1216-1229, 1983.

Kaas JH, Nelson RJ, Sur M, Lin CS, and Merzenich MM. Multiple representations of the body within the primary somatosensory cortex of primates. *Science* 204: 521-523, 1979.

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

Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, and Shenoy KV. Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat Commun* 6: 2015.

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

Krizhevsky A, Sutskever I, and Hinton GE. Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*2012, p. 1097-1105.

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

Loeb GE, Hoffer JA, and Marks WB. Activity of spindle afferents from cat anterior thigh muscles. III. Effects of external stimuli. *J Neurophysiol* 54: 578-591, 1985.

London BM, and Miller LE. Responses of somatosensory area 2 neurons to actively and passively generated limb movements. *J Neurophysiol* 109: 1505-1513, 2013.

London BM, Ruiz-Torres R, Slutzky MW, and Miller LE. Designing stimulation patterns for an afferent BMI: Representation of kinetics in somatosensory cortex. In: *Engineering in Medicine and Biology Society,EMBC, 2011 Annual International Conference of the IEEE*2011, p. 7521-7524.

Lucas A, Tomlinson T, Rohani N, Chowdhury RH, Solla SA, Katsaggelos AK, and Miller LE. Deep Neural Networks for Modeling Neural Spiking in S1 Cortex. *Frontiers in Systems Neuroscience* 13: 13, 2019.

McFadden D. *Quantitative methods for analyzing travel behavior of individuals: some recent developments*. Institute of Transportation Studies, University of California Berkeley, 1977.

Mishkin M, and Ungerleider LG. Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. *Behavioural brain research* 6: 57-77, 1982.

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

Morrow MM, Jordan LR, and Miller LE. Direct comparison of the task-dependent discharge of M1 in hand space and muscle space. *J Neurophysiol* 97: 1786-1798, 2007.

Mountcastle VB, Lynch JC, Georgopoulos A, Sakata H, and Acuna C. Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *J Neurophysiol* 38: 871-908, 1975.

Mussa-Ivaldi FA. Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci Lett* 91: 106-111, 1988.

Nadeau C, and Bengio Y. Inference for the Generalization Error. *Machine Learning* 52: 239-281, 2003.

Nelson RJ. Activity of monkey primary somatosensory cortical neurons changes prior to active movement. *Brain Research* 406: 402-407, 1987.

Oby ER, Ethier C, and Miller LE. Movement representation in the primary motor cortex and its contribution to generalizable EMG predictions. *J Neurophysiol* 109: 666-678, 2012.

Omrani M, Murnaghan CD, Pruszynski JA, and Scott SH. Distributed task-specific processing of somatosensory feedback for voluntary motor control. *Elife* 5: e13141, 2016.

Padberg J, Cooke DF, Cerkevich CM, Kaas JH, and Krubitzer L. Cortical connections of area 2 and posterior parietal area 5 in macaque monkeys. *Journal of Comparative Neurology* 2018.

Perich MG, Gallego JA, and Miller LE. A Neural Population Mechanism for Rapid Learning. *Neuron* 100: 964-976.e967, 2018.

Phillips CG, Powell TP, and Wiesendanger M. Projection from low-threshold muscle afferents of hand and forearm to area 3a of baboon's cortex. *J Physiol* 217: 419-446, 1971.

Pons TP, Garraghty PE, Cusick CG, and Kaas JH. The somatotopic organization of area 2 in macaque monkeys. *J Comp Neurol* 241: 445-466, 1985.

Prochazka A, and Wand P. Independence of fusimotor and skeletomotor systems during voluntary movement. In: *Muscle receptors and movement*, edited by A. Taylor A. ProchazkaMacMillan London, 1981.

Prochazka A, Westerman R, and Ziccone S. Discharges of single hindlimb afferents in the freely moving cat. *Journal of Neurophysiology* 39: 1090-1104, 1976.

Proske U, and Gandevia SC. The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force. *Physiological Reviews* 92: 1651-1697, 2012.

Proske U, and Stuart G. The initial burst of impulses in responses of toad muscle spindles during stretch. *The Journal of physiology* 368: 1-17, 1985.

Prud'homme MJL, and Kalaska JF. Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements. *J Neurophysiol* 72: 2280-2301, 1994.

Rack PM, and Westbury DR. Elastic properties of the cat soleus tendon and their functional importance. *The Journal of Physiology* 347: 479-495, 1984.

Richardson AG, Weigand PK, Sritharan SY, and Lucas TH. A chronic neural interface to the macaque dorsal column nuclei. *Journal of Neurophysiology* 115: 2255-2264, 2016.

Romo R, Hernandez A, Zainos A, and Salinas E. Somatosensory discrimination based on cortical microstimulation. *Nature* 392: 387-390, 1998.

Russo AA, Bittner SR, Perkins SM, Seely JS, London BM, Lara AH, Miri A, Marshall NJ, Kohn A, Jessell TM, Abbott LF, Cunningham JP, and Churchland MM. Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron* 2018.

Sainburg RL, Ghilardi MF, Poizner H, and Ghez C. Control of limb dynamics in normal subjects and patients without proprioception. *JNeurophysiol* 73: 820-835, 1995.

Sainburg RL, Poizner H, and Ghez C. Loss of proprioception produces deficits in interjoint coordination. *J Neurophysiol* 70: 2136-2147, 1993.

Salas MA, Bashford L, Kellis S, Jafari M, Jo H, Kramer D, Shanfield K, Pejsa K, Lee B, Liu CY, and Andersen RA. Proprioceptive and cutaneous sensations in humans elicited by intracortical microstimulation. *eLife* 7: e32904, 2018.

Sanes JN, Mauritz KH, Evarts EV, Dalakas MC, and Chu A. Motor deficits in patients with large-fiber sensory neuropathy. *Proc Natl Acad Sci U S A* 81: 979-982, 1984.

Sedda A, and Scarpina F. Dorsal and ventral streams across sensory modalities. *Neuroscience Bulletin* 28: 291-300, 2012.

Seelke AM, Padberg JJ, Disbrow E, Purnell SM, Recanzone G, and Krubitzer L. Topographic maps within Brodmann's area 5 of macaque monkeys. *Cerebral Cortex* 22: 1834-1850, 2011.

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

Stevenson IH, Cherian A, London BM, Sachs NA, Lindberg E, Reimer J, Slutzky MW, Hatsopoulos NG, Miller LE, and Kording KP. Statistical assessment of the stability of neural movement representations. *J Neurophysiol* 106: 764-774, 2011.

Suresh AK, Winberry J, Versteeg C, Chowdhury RH, Tomlinson T, Rosenow JM, Miller LE, and Bensmaia SJ. Methodological considerations for a chronic neural interface with the cuneate nucleus of macaques. *Journal of Neurophysiology* jn. 00436.02017, 2017.

Sussillo D, Churchland MM, Kaufman MT, and Shenoy KV. A neural network that finds a naturalistic solution for the production of muscle activity. *Nat Neurosci* 18: 1025-1033, 2015.

Tabot GA, Dammann JF, Berg JA, Tenore FV, Boback JL, Vogelstein RJ, and Bensmaia SJ. Restoring the sense of touch with a prosthetic hand through a brain interface. *Proceedings of the National Academy of Sciences* 110: 18279-18284, 2013.

Tomlinson T, and Miller LE. Toward a proprioceptive neural interface that mimics natural cortical activity. In: *Progress in Motor Control: Theories and Translations*, edited by Laczko J, and Latash MLSpringer, 2016.

Truccolo W, Eden UT, Fellows MR, Donoghue JP, and Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol* 93: 1074-1089, 2005.

Weber DJ, London BM, Hokanson JA, Ayers CA, Gaunt RA, Ruiz-Torres R, Zaaimi B, and Miller LE. Limb-state information encoded by peripheral and central somatosensory neurons: Implications for an afferent interface. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on* 19: 501-513, 2011.

Weber DJ, Stein RB, Everaert DG, and Prochazka A. Decoding sensory feedback from firing rates of afferent ensembles recorded in cat dorsal root ganglia in normal locomotion. *IEEE Trans Neural Syst Rehabil Eng* 14: 240-243, 2006.

Yamada H, Yaguchi H, Tomatsu S, Takei T, Oya T, and Seki K. Representation of Afferent Signals from Forearm Muscle and Cutaneous Nerves in the Primary Somatosensory Cortex of the Macaque Monkey. *PLOS ONE* 11: e0163948, 2016.

Young D, Willett F, Memberg WD, Murphy BA, Rezaii P, Walter B, Sweet JA, Miller J, Shenoy KV, and Hochberg L. Closed-loop cortical control of virtual reach and posture using Cartesian and joint velocity commands. *Journal of neural engineering* 2018.

Zaaimi B, Ruiz-Torres R, Solla SA, and Miller LE. Multi-electrode stimulation in somatosensory cortex increases probability of detection. *Journal of Neural Engineering* 10: 056013, 2013.