32nd Annual International Conference of the IEEE EMBS
Buenos Aires, Argentina, August 31 - September 4, 2010

Low-Dimensional Neural Features Predict Muscle EMG Signals

Zuley Rivera-Alvidrez, Rachel S. Kalmar, Stephen I. Ryu and Krishna V. Shenoy, *Senior Member, IEEE*

***Abstract*— Understanding the relationship between neural activity in motor cortex and muscle activity during movements is important both for basic science and for the design of neural prostheses. While there has been significant work in decoding muscle EMG from neural data, decoders often require many parameters which make the analysis susceptible to overfitting, which reduces generalizability and makes the results difficult to interpret. To address this issue, we recorded simultaneous neural activity from the motor cortices (M1/PMd) of rhesus monkeys performing an arm-reaching task while recording EMG from arm muscles. In this work, we focused on relating the mean neural activity (averaged across reach trials to one target) to the corresponding mean EMG. We reduced the dimensionality of the neural data and found that the curvature of the low-dimensional (low-D) neural activity could be used as a signature of muscle activity. Using this signature, and without directly fitting EMG data to the neural activity, we derived neural axes based on reaches to only one reach target (<5% of the data) that could explain EMG for reaches across multiple targets (average R^2 = 0.65). Our results suggest that cortical population activity is tightly related to muscle EMG measurements, predicting a lag between cortical activity and movement generation of 47.5 ms. Furthermore, our ability to predict EMG features across different movements suggests that there are fundamental axes or directions in the low-D neural space along which the neural population activity moves to activate particular muscles.**

I. INTRODUCTION

Understanding the link between neural activity in motor cortex and limb movement has been the goal of the study of motor control for many decades. A substantial amount of work has been done linking the activity of sequentially-recorded single-neurons to virtually every kind of movement parameter, including velocity, acceleration, force, muscle EMG and many more (for review see [1]). Microelectrode arrays have made it possible to simultaneously record the activity of hundreds of cells, enabling the field of neural prostheses (for an overview see [2]), and pushing neuroscience further away from the study of single cells towards a population view of the system. Because populations of cells drive muscles, this may be a promising framework in which to relate neural activity to movement.

This work was supported by CONACYT, DARE, SGF, BioX, Christopher and Dana Reeve Foundation, NIH RO1-NS054283, and NIH Director's Pioneer Award IDP1OD006409.

Zuley Rivera-Alvidrez is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305. zuley@stanford.edu

Rachel Kalmar is with the Neurosciences Program, Stanford University, Stanford, CA 94305. kalmar@stanford.edu

Stephen Ryu is with the Department of Neurosugery, Palo Alto Medical Foundation, and the Department of Electrical Engineering, Stanford University, Stanford, CA 94305. seoulman@stanford.edu

Krishna Shenoy is with the Departments of Electrical Engineering and Bioengineering, and the Neurosciences Program, Stanford University, Stanford, CA 94305. shenoy@stanford.edu

Analyzing the neural activity from hundreds of cells poses, however, a number of challenges. While successful decoders have been implemented to predict kinematics (for an example see [3]) and muscle EMG (for an example see [4]), these decoders require many parameters, thus making them susceptible to overfitting. Overfitting affects the ability to predict behavior across different movement conditions and contexts that were not included in the training of the decoder. Furthermore, it makes it difficult to interpret the neural-to-behavior relationship found (e.g., the meaning of the decoder coefficients).

Part of the problem is due to the fact that decoders find a mapping between the high-dimensional neural data (one dimension per neuron) and the signal being decoded, thereby requiring at least one parameter per cell. Some of the features of the high-dimensional data, however, might be noise or unrepresentative of the neural population as a whole [5]. Recently, methods to analyze simultaneously recorded data have been proposed based on the idea that a small number of dimensions are enough to explain the meaningful features of the neural population [5]. The resulting state-space models (see Fig 1a), can thus provide us with a way to denoise, visualize (if 3-D or less), and think of the evolution of the neural population (neural trajectory) in a few meaningful dimensions.

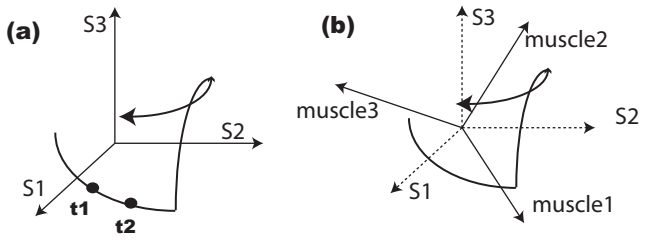

Fig. 1. Low-dimensional neural state-space. (a) shows the hypothetical evolution of the neural population activity (neural trajectory) in the state-space defined by the axes S1, S2, S3. *t1* and *t2* are two different time steps corresponding to different points in the neural trajectory. (b) shows an alternative space defined by behaviorally-relevant axes.

While this state-space may bring us closer to a useful representation of the neural population, thinking of the trajectories in terms of the orthogonal axes in state-space (S1, S2 and S3 in Fig. 1a), is not necessarily intuitive. This is because the state-space axes are likely to be a combination of the behavioral signals explained by the neural data without clearly relating to any one behavioral signal in particular. Figure 1b shows a potentially more useful and intuitive representation of the data in terms of a set of behaviorally-

related axes. A neural trajectory moving along an axis, say muscle1, will activate or deactivate muscle1, whereas one moving along muscle2 will change the activity of muscle2. Moreover, in this idealized view the axis corresponding to a particular muscle reflects the true relationship of the neural data to that muscle. True in the sense of generalizability: it relates to the muscle in the same way across all tasks and behavioral contexts.

Unfortunately, even if we assume that such fundamental axes exist in the data, finding them may well not be trivial. Part of the problem is not knowing the adequacy of our behavioral measurements. While corticomotoneuronal anatomical connections [6] suggest that the the neural activity in motor cortex is linked to muscle activity, measuring muscle activity is difficult. EMG recordings are known to be noisy and vulnerable to movement artifacts. Furthermore, the exact location and orientation of the electrodes is likely to result in somewhat different signals. In our experiments with state-of-the-art EMG measurements, mean EMG recordings of the same muscle for nominally the same movement, correlate by only 0.80 (R^2) on average, below the ideal correlation of 1. Thus, an individual EMG recording is unlikely to fully characterize muscle activity. Additionally, cortical activity may undergo significant transformations by the neural circuits in the spinal cord, which may further weaken the relationship between the cortical signals and the EMG measurements.

In this work we attempt to find general neural-to-muscle relationships (i.e., the axes in Fig. 1b). While our ultimate goal is to relate neural-to-muscle data on a single-trial basis, for this work we concentrated on relating mean neural data to mean EMG muscle data. Importantly, in doing so we wanted to avoid directly fitting the EMG data to the low-dimensional (low-D) neural trajectories as much as possible because, 1) as explained previously, we do not know how tight the relationship between the neural activity and an individual EMG recording is, and 2) any possibility of overfitting will reduce the interpretability of our results. Instead, we assumed that there were signatures of muscles in the neural activity [7], [8]. And from these, we could predict muscle EMG features across different movements *without ever fitting the EMG to the neural data*. Our results make predictions about the mechanism of movement generation and have positive implications for finding general neural-to-muscle relationships that can be harnessed by prosthetic systems.

II. METHODS

A. Behavioral Task

All animal protocols were approved by the Stanford University Institutional Animal Care and Use Committee. We trained two rhesus macaque monkeys to perform a direct (non-delayed) center-out reach task to a grid of targets on a vertical screen. For monkey L, 28 targets were located along concentric rings at radial distances of 45, 70, 95, and 120 mm and angles 40, 85, 130, 175, 220, 310, and 355. For monkey I, 24 targets were located at radial distances of 50, 80, and 110 mm and angles 30, 70, 110, 150, 190, 230, 310, and 350 degrees. Simultaneous neural (microelectrode array, 96 channels, Blackrock Inc, Salt Lake City, UT) and behavioral activity were recorded, as described elsewhere [3]. In addition, we recorded simultaneous muscle EMG (penetrating hookwire electrode) for one muscle per dataset, as described in [9].

For monkey L, the array was placed in the M1/PMd forearm region, and we recorded EMG from a wrist flexor (nominally flexor carpi ulnaris) and a wrist extensor (nominally extensor carpi radialis) muscle. The array for monkey I was placed in the M1/PMd arm region, and we recorded EMG from biceps, trapezious, and deltoid. Because our analysis required the EMG to have a clear point of maximum activity, we only analyzed biceps and trapezious datasets due to the tendency of deltoid EMG to have sustained high activity.

B. EMG and Neural Processing

The neural activity was spike-sorted using a PCA-based spike sorting algorithm [10]. We did not distinguish between single and multi-neuron activity in our population; for the rest of this manuscript, we will use the term "unit" to refer to either case. The neural activity was aligned to movement onset (computed as the time in which 15% of the peak hand speed was reached), and we analyzed a time window starting 350 ms before movement onset and ending 350 ms after movement onset. To select units for our analyses, we computed the standard deviation (SD) of the mean firing rates across target conditions as a function of time. At the time of maximum SD, we required that the units had reliable tuning to our target conditions (ANOVA, $P<.0001$). This criteria was the same as [11], except that we did not exclude units based on plan activity. In addition, we also enforced a modulation depth criterium as the difference between the highest and the lowest mean firing rates (across conditions, at the time of maximum SD) of 10 spikes/s. Between 70-140 units per dataset satisfied this criteria. This criteria was aimed at emphasizing the behaviorally-relevant dimensions of the neural data by selecting only the units that showed modulation by the reach task.

We computed firing rate estimates by low-pass filtering (Gaussian, 25 ms SD) the spike trains, as illustrated in Figure 2a and b. In order to obtain a more compact representation of the neural population, we used dimensionality reduction. The details are explained in the following section. Briefly, for each dataset we picked the target associated with the maximum peak EMG activity (training target, see next section). We used factor analysis to reduce the dimensionality of the neural data corresponding to reaches to that target down to 15 dimensions (Fig. 2c). To work in an X-dimensional space, we picked the top X dimensions of the 15-D factor analysis model (for justification please see next section). Figure 2d shows two 3-D single-trial neural trajectories for two hypothetical trials. We averaged the single-trial trajectories to obtain a mean neural trajectory (Fig. 2e). The neural data for the other targets were processed in the same way, except

that instead of finding a separate factor analysis space per target, we projected the single-trial neural data onto the space found for the training target. The single trial neural data was averaged together for each target to obtain a mean low-D neural trajectory per target in the same low-D space.

The EMG activity was differentiated to remove any DC offset and rectified as in [11]. As with the neural data, we aligned the EMG to movement onset, and we analyzed a time window starting 350 ms before movement onset and ending 350 ms after movement onset. The rectified and differentiated EMG was low-pass filtered (Gaussian, 25 ms SD), as illustrated in Figure 2f and 2g for two hypothetical muscles. We averaged the single-trial EMG recordings for each target to obtain mean EMG activities per target (Fig. 2h).

For some analyses, we wanted a muscle activity representation that we could more easily compare to the low-D neural trajectories. We computed "muscle population trajectories" as follows. We used two mean EMG signals (wrist flexor and wrist extensor for monkey L, biceps and trapezius for monkey I) to create a two-dimensional matrix (2 x number of timesteps) per target. For the training target, we used PCA to change the axes from individual muscles to principal components (Fig. 2i). The axes were normalized to have unit variance. Figure 2j shows a hypothetical muscle trajectory. Note that this process was aimed at replicating the steps used to find the low-D neural trajectories. The sole exception is that we used PCA instead of factor analysis since we were not interested in denoising our mean EMG data because the operation of calculating the mean EMG had already effectively removed the noise from individual trials' EMGs. While we could have used factor analysis to model the noise (see next section) of single trials before computing the mean, the EMG recordings for different muscles were not recorded simultaneously (we recorded one muscle per dataset), thereby making it difficult to accurately match single-trial data from two EMG recordings which corresponded to slightly different movements.

C. Dimensionality Reduction

The purpose of this study was to find general neuron-to-muscle relationships without directly fitting EMG traces to the neural data. While there may be many signals explained by the neural population activity, we reasoned that not all signals were likely to be equally strong at all times and for all types of movements. For a particular muscle, we picked the movement (reach target) associated with the highest peak EMG for that muscle and used that as the training target for that muscle. We reasoned that training targets might correspond to those neural trajectories in Fig. 1b, that had the highest excursions on the sets of axes relevant to each muscle. For each EMG, we used the neural data associated with its training target to find the low-D state-space in which to characterize the neural activity.

To reduce the dimensionality of the data, we used factor analysis with orthogonal axes ordered by explained data covariance, as described in [5]. While we cannot go into the details of the factor analysis probabilistic model here, in simple terms, it is based on the assumption that the high-dimensional data can be described as a combination of a lower-dimensional set of denoised signals (factors) and noise. An important parameter is therefore the dimensionality of the data, or how many factors are needed to represent the part of the data that is not noise. The dimensionality of comparable neural data has been previously estimated to be around 8-12 [5], so 15 was arbitrarily picked to be outside this range. Note that it may be that the neural signal relevant to the muscle EMG that we care about could be fully explained by the top few factors (say X), while the remaining factors represent other, unknown signals that are also explained by the neural data. Because these other signals are not noise, it would be inaccurate to define the factor analysis model with only X factors, as factor analysis assumes that everything it cannot explain is noise. Instead, to work in an X-dimensional space, we used the top X factors of the full 15-D factor analysis model, which presumably contains all the explanatory dimensions of the neural data [5]. Finally, because our methods (such as curvature, defined in the next section) are sensitive to the scale of the factors, and we have no real way of knowing which factors carry EMG-relevant information, we chose to scale our axes such that the factors would have equal variance.

Once the low-D space was found using the neural data for the training target, the low-D axes and their scale were fixed. To obtain low-D trajectories for the other targets, we projected the neural data associated with those targets onto the fixed low-D space.

D. Neural Curvature and Neural Curvature Axis

We computed a curvature metric of the neural trajectories at each time point t by taking a time window 10 ms before t and 10 ms after t (see Figure 3a, inset), and computed the ratio of pathlength to end-point displacement. Note that "turning points" highlighted in the trajectory corresponded to peaks in neural curvature (Fig. 3b).

We wondered whether the axis related to the curvature peak, or trajectory turning point, might be related to muscle activity (see Results). We observed that while the whole neural trajectory lies in a multi-dimensional space (3-D in Figure 3c), the neural data tends to be two-dimensional locally in a small time window. We defined the neural curvature axis (NCA) at time t as an axis that is perpendicular to the neural trajectory at time t (perpendicular to the black line segment in the Fig. 3c, inset) and lies in the same plane as the neural trajectory segment defined by [t-10ms, t+10ms]. Picking an axis that is perpendicular to the neural trajectory at the reference time point ensured that the projection of the neural trajectory onto the axis would have a derivative equal to zero at that point, capturing the turn or point of high curvature in the neural trajectory. Picking an axis that lies in the same local plane as the neural data ensured that the axis would be relevant (explained local neural variance) to the trajectory locally. The plane was found by computing

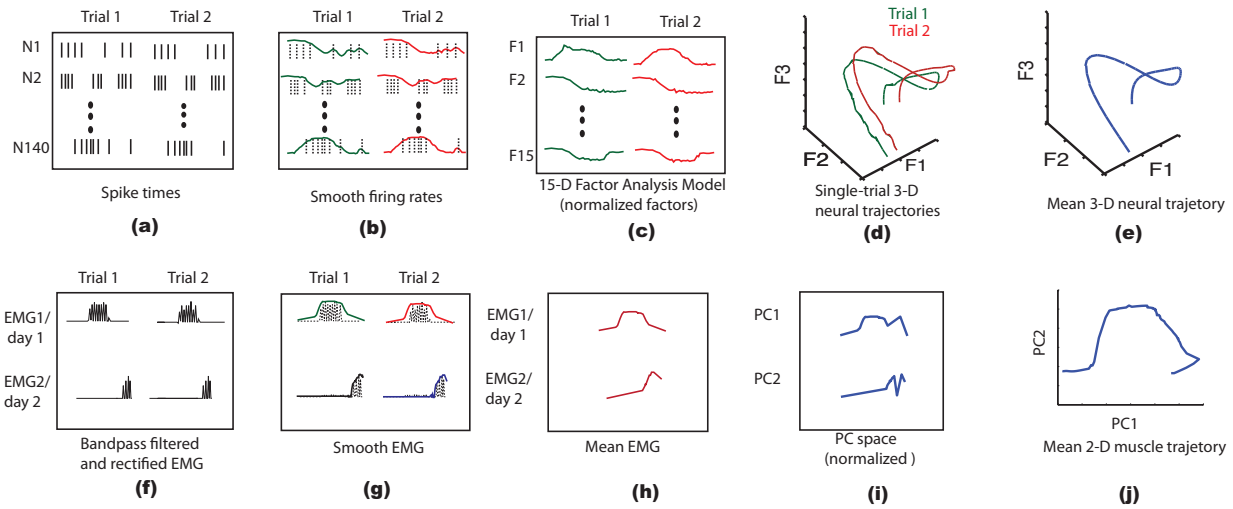

Fig. 2. Illustration of neural and EMG processing. The neural activity for the training target was processed as follows. The spike times (a) were low-pass filtered with a 25 ms SD Gaussian to obtain smooth firing rates (b). The firing rates were used to estimate a low-dimensional (15-D) factor analysis model (c). To obtain single-trial 3-D neural population trajectories, we used the top 3 factors of the 15-D model (d). The single trials for each reach target were averaged together to produce a mean low-D neural trajectory (e). The processing of the neural data corresponding to targets other than the training target was the same except that instead of computing a separate factor analysis model, we projected the data onto the 15-D axes found for the training target (c). EMG data for various muscles was recorded in different days (one muscle per day). The raw EMG data was bandpass filtered, amplified, differentiated to remove the DC offset and rectified (f). Smooth EMG envelopes were obtained by convolving the EMG data with a 25 ms Gaussian (g). The EMG data corresponding to different reach trials was averaged together to compute mean EMG for each target and each muscle (h). For some analyses, we computed muscle EMG trajectories by applying PCA to the mean EMGs for the training target and normalizing the axes (i and j).

the top two principal components of the neural trajectory segment.

We computed the NCA using the mean neural activity for the training target and projected the mean neural trajectories for all the targets onto this single NCA. The result of this process was a 1-D time series per target (see Figure 3d), which we could directly compare to EMG. Because we believed that the NCA was related to the peak EMG (see Results), for an NCA computed at time t, we assumed a lag equal to the difference between the time of the peak EMG (for the training target) and t (labeled assumed lag in Fig. 3d). This lag is presumably related to the biological latency between cortical neural activity and EMG generation (see Discussion). As a measure of the similarity between the NCA projections and the mean EMG, we computed the correlation coefficient for the concatenated 1-dimensional neural projections across all targets, and the similarly concatenated mean EMG traces, time-shifted by the corresponding lag.

III. RESULTS

A. *Neural Curvature Provides a Signature of Muscle Activity*

We wanted to find a general relationship between the neural population trajectories and individual muscles without directly fitting EMG traces to the neural data. In order to do this, we found a low-D space in which to characterize the neural data based on movements to reach targets for which each muscle EMG dataset had the highest level of activity (training targets, see Methods). We reasoned that, for these movements, the neural activity related to each muscle more strongly than to other signals. Using these training targets, we then wanted to find a signature of muscle activity in the neural data that we could use to find a neural-to-muscle

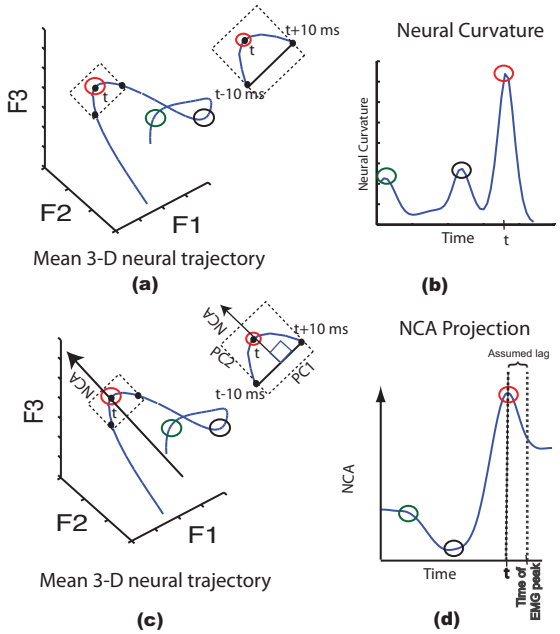

Fig. 3. Illustration of neural curvature and neural curvature axis (NCA) computation. (a) 3-D neural trajectories (circles corresponding to turns). The curvature at time t was computed in a window 10 ms before and after t, as the ratio of pathlength (blue segment in panel inset) to end-point displacement (black segment). (b) The resulting curvature plotted as a function of time. The neural curvature axis (NCA) at time t was computed in the same window (c, inset) as the axis perpendicular to the black segment and lying in the same plane as the neural trajectory segment (PCA, top 2 PCs). The neural trajectories were projected onto the NCA to obtain 1-D timeseries per target (d).

relationship. Using muscle EMG recordings, we explored this problem as follows. Figure 4a shows two mean EMG recordings corresponding to a wrist extensor (left) and a wrist flexor (right) muscle. These two muscles can be used to define the axes of a 2-D muscle trajectory (Fig. 4b, left). Alternatively, we can instead define the axes of the space as the normalized principal components (PC1, PC2, normalized, see Methods) of the 2-D space of mean EMG, as in Fig. 4c. Like the neural trajectory axes (see S1, S2 and S3 in Fig. 1a, for an example), the principal components (PC1 and PC2) are each a linear combination of the signals in the data (flexor and extensor mean EMG signals in this case) and may not necessarily look like any one particular signal. The corresponding PC muscle trajectory is shown in Fig. 4d, left. Note that sharp turning points designated by the circles can be identified in both trajectories. In order to study these turning points, we computed a measure of curvature (see Methods) along the muscle trajectories, which we could plot as a function of time. Figure 4b, right, and Fig. 4d, left, show the curvature as a function of time for the two trajectories. Notice that even though the axes of each 2-D space are quite different (Figs. 4a, 4c), the corresponding curvature traces (Figs. 4b, 4d, right) are very similar with peaks in curvature occuring at the time of peak extensor EMG activity (blue circles) and peak flexor EMG activity (red circles).

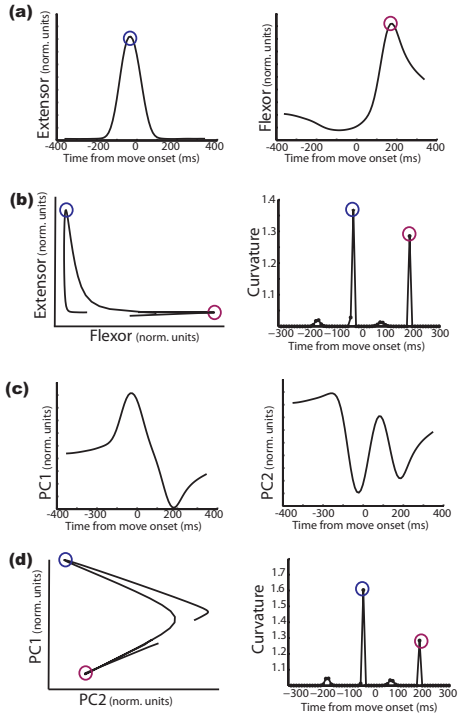

Fig. 4. Curvature provides a signature of muscle activity. (a) Mean EMG data for an extensor (left) and a flexor (right) muscle. (b) 2-D muscle trajectories in the space defined by the EMG axes (left), and the corresponding curvature along the trajectory as a function of time (right). (c) Normalized principal components of the 2-D space defined by the EMG data. (d) 2-D space defined by the normalized PC axes (left), and curvature of the resulting trajectory (right). The blue circles indicate the time of extensor EMG peak (blue), and flexor EMG peak (red).

We wondered whether curvature might also serve as a signature of specific muscle activity in the neural trajectories. Figure 5a shows the mean EMG activity used to construct the muscle population trajectories for the set of training targets. Note that biceps and trapezius had the same training target. Below them (Fig. 5b) are plots of the curvature as a function of time for the resulting muscle trajectories. Curvature peaks occurred at the time of peak EMG for individual muscles, as well as at other times (Fig. 5b, right, first peak). When analyzing the low-D neural trajectories, we first had to pick a dimensionality. For each dataset, we computed the curvature of the neural trajectories for the training targets at each dimensionality (2-15). We then compared the resulting curvature plots to the muscle curvature plots shown in Fig. 5b. For each dataset, we picked the lowest dimensionality of the neural data that had curvature peaks that roughly matched the number of muscle EMG large curvature peaks in Fig. 5b. Fig. 5c shows the curvature of the corresponding neural trajectories in 2-D for monkey L and 7-D for monkey I. For monkey I, the curvature traces of neural trajectories at dimensionalities 2-6 had only one of the two peaks, while the curvature of the trajectories at dimensionalites 7-15 looked almost exactly like the curvature trace shown. Note the remarkable similarity between the curvature of the muscle trajectories and the curvature of the neural trajectories for those targets. The peaks in neural curvature led the peaks of the EMG curvature with a lag of 47.5 ms (SD = 9.65 ms).

Additional analyses (data not shown) showed that individual neural curvature peak times for all targets (even those not trained on) can predict individual EMG peak times, further supporting a relationship between the curvature peaks and the EMG peaks.

B. Neural Axes Projections Correlate with EMG Signals

We reasoned that the relation between the neural curvature peaks and the peak muscle EMG might reflect an excursion of the neural trajectory along the axis related to the muscle in question (see Fig. 1b). With this in mind, we devised a simple method to compute the neural curvature axis (NCA, see Methods) in the low-D space aimed both at capturing neural variance in a local time window and the turn in the neural trajectory centered at the time of peak curvature associated with peak EMG (e.g. the curvature peaks with circles in Fig. 5c). In order to test generalization, for each EMG, we used the mean neural activity in the low-D space for the training target only (see Methods) to compute a unique NCA. We then projected the neural trajectories for all targets onto this axis to obtain a 1-D time-series per target (NCA projections), which we could then compare directly to the corresponding mean EMG recordings.

Figure 6a shows the NCA projections (left column) for the NCA computed at the time of peak neural curvature associated with the wrist flexor muscle (curvature peak designated by the red circle in Fig. 5c, left). The wrist flexor EMG mean signals for all targets are also shown (Fig. 6a, right). Note that the overall shape of the EMG, which goes down in the first part of the movement and peaks during the second part, is well predicted, especially for the training

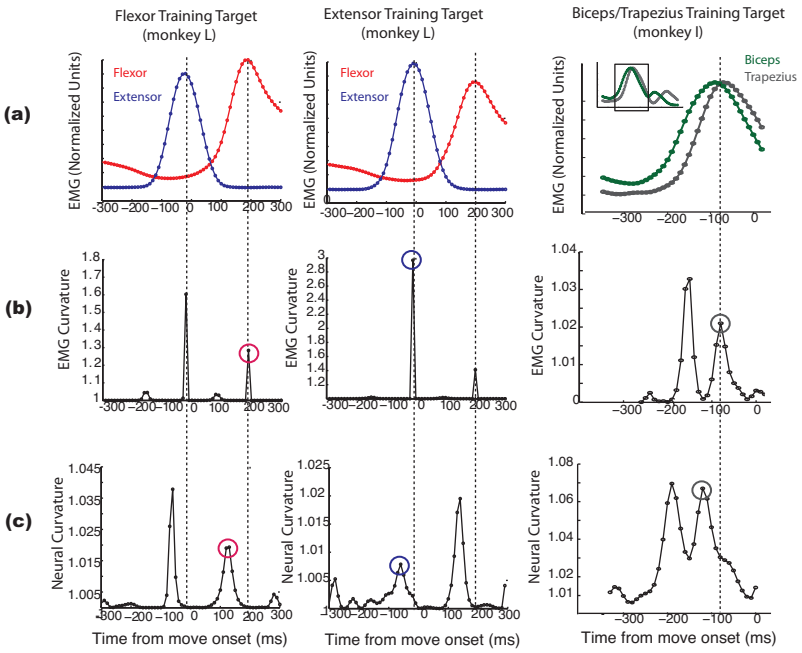

5. Neural and EMG curvature are similar. (a) Mean EMG activity for the wrist flexor training target (target associated with max peak flexor EMG, see Methods) on the left, the wrist extensor training target (center), and the training target for biceps and trapezius (right). (b) Curvature as a function of time for the muscle trajectories constructed (see Methods) using the EMG signals in panel a. (c) Curvature as a function of time corresponding to the 2-D neural trajectories for monkey L (left and center) and to the 7-D trajectories for monkey I (right) for the same targets. The circles indicate the neural curvature points associated with wrist flexor (red), wrist extensor (blue), and biceps/trapezius (gray).

target (black trace). The correlation coefficient (R^2) between the NCA projections and the flexor EMG mean signals was 0.71. Figure 6b shows the same plots for the wrist extensor muscle (curvature peak designated by blue circle in Fig. 5c, center). The NCA projections, like the extensor EMG mean signals, do not go down below pre-movement levels and the R^2 value found was 0.60. Figure 6c, left, shows the NCA projections found using the curvature peak designated by the gray circle in Fig. 5c, right. The NCA projections looked remarkably similar (R^2 = 0.73) to trapezius EMG (Fig. 6c, right), which was the most variable EMG across target conditions. The NCA projection for the training target (black trace) was similar to the mean EMG for that target, as were the overall differences across targets in initial EMG phase, timing and magnitude of primary and secondary peaks.

C. Fine Time Resolution of Muscle Activity

Both biceps and trapezius had the same training target, and for that target, they had remarkably similar EMG activity (Fig. 5a, right). Furthermore, the same neural curvature peak was associated with both muscles (see Figs. 5b, 5c, right). We wondered whether the 20 ms time difference between the biceps and the trapezius EMG peak could be used to dissociate the neural axes linked to those muscles.

As shown in the previous section, the NCA computed at the time of the neural curvature peak was predictive of trapezius EMG (R^2 = 0.73). It was less predictive of biceps EMG (R^2 = 0.29). We asked if computing the NCA 20 ms before the curvature peak would result in projections predictive of biceps EMG activity. Figure 6d shows the resulting NCA projections (left) along with the biceps EMG means (right). These NCA projections were more similar to biceps EMG (R^2 = 0.55) than trapezius EMG (R^2 = 0.49). Computing the NCA at other times varying from 150 ms to 50 ms after the curvature peak showed that the NCA that was computed 20 ms before the curvature peak was indeed the most predictive of biceps (highest R^2) compared to all the NCA axes computed in that time window. In order to assess whether these results were unique to this data dimensionality (7-D), we found the NCA that was most predictive of biceps across all dimensionalities and computed the median value. The median time for which the NCA was most predictive of biceps EMG was indeed 20 ms (SD = 5.5ms) before the curvature peak. For this NCA, the median R^2 value was 0.57 when compared to biceps EMG, and 0.35 when compared to trapezius EMG. We repeated the analyses using the trapezius EMG and found that the median time for which the NCA was most predictive of trapezius was indeed the time of the neural curvature peak (SD = 10ms). For this NCA, the median R^2 value was 0.59 when compared to trapezius and 0.21 when compared to biceps. Thus, using neural data for a training target associated with remarkably similar EMG activity for both biceps and trapezius, we were able to find neural axes predictive of each muscle based on a 20 ms time difference between the EMG peaks.

IV. DISCUSSION

Our goal was to find a general relationship between neural population activity and muscle activity that could be useful for prosthetic systems as well as be interpreted scientifically. While EMG decoders that use linear filters have been proposed, the performance of such decoders tends to drop sharply when testing on a task different from the training task [4]. Thus, traditional decoders are highly sensitive to the characteristics of the training dataset. To address this issue, one possibility is to train on the complete space of all possible movements, task conditions, and behavioral contexts. This, however, may be impractical. Ideally, we would like methods that can extract a general neuron-to-muscle relationship using as little training data as possible. In this

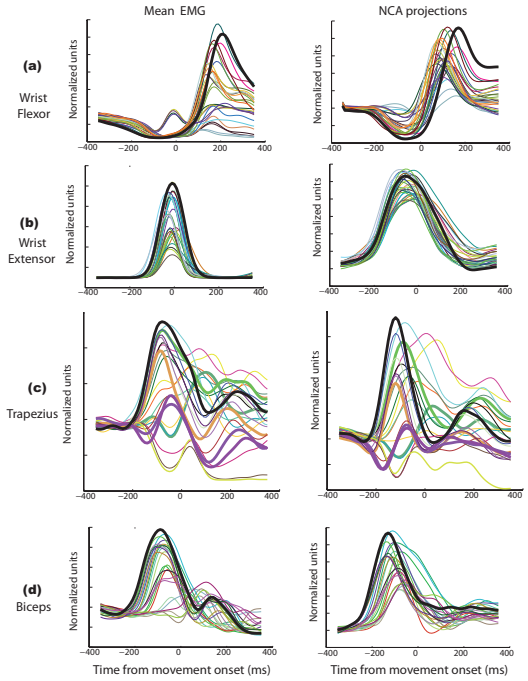

Fig. 6. NCA projections predict EMG features. Left column shows the NCA projections (see text for details) for all targets (superimposed and shown in different colors) associated with wrist flexor (a), wrist extensor (b), trapezius (c), and biceps (d). The correlations R^2 between the NAC projections and the mean EMG actvity for the same muscles (right column) were 0.71 (extensor), 0.6 (flexor), 0.73 (trapezius), and 0.55 (biceps).

work, we were able to extract neuron-to-muscle relationships based on the average movement to the training target (<5% of the data). Because the EMG data was not directly fit to the neural activity, the relationship found was unlikely to overfit to the training data. Indeed, we found that our axes were predictive of EMG features across different movements (R^2 = 0.55-0.73). The strength of such relationships is high, especially considering that an individual EMG recording is unlikely to fully characterize muscle activity. Remarkably, in the case of biceps and trapezius, we were able to dissociate two muscles based on training data that was far from ideal (same training target, nearly the same EMG activity for that target; see Fig. 5a, right) using only a 20 ms difference in the timing of their EMG features. Notably, while for this work we concentrated on predicting mean EMG activity, preliminary results show that the relationships found can also predict differences in single-trial EMG, and more work is underway to directly compare our single-trial performance to the performance of other decoding techniques.

In addition, our results make predictions about the mechanisms of movement generation. Specifically, our results predict a delay between cortical signal generation and muscle activity of 47.5 ms (SD = 9.65 ms). This lag is longer than the estimated delay between corticomotoneuronal single-cell firing and postspike EMG facilitation (7.7 ms - 13.3 ms) [12], but in agreement with lag estimates based on optimal correlations between a population of M1 cells and EMG [13]. Furthermore, our ability to estimate axes in the low-D neural space that we could use to predict EMG features in a general way (across target conditions) suggest that there are fundamental directions in the neural space along which the neural population moves to activate particular muscles. The fact that the population included a mixture of single neurons and multi-units of unidentified type suggests a relationship between cortical population activity and muscle activity that is not limited to corticomotoneuronal cells, and that is surprisingly tight considering the important intervening spinal cord circuitry.

Overall, our ability to relate neural activity to muscle activity with few parameters and using very little training data makes this a promising framework in which to study motor control both from a scientific perspective and for the development of prosthetic systems with highly generalizable performance.

V. ACKNOWLEDGMENTS

We thank M. Churchland for helpful discussions and help with EMG recordings, A. Afshar for helpful discussions, C. Chestek for assistance in animal training, M. Risch for help with EMG recordings, expert surgical assistance and veterinary care, S. Eisensee for administrative assistance, and D. Haven for computer assistance.

REFERENCES

[1] S.H. Scott, "The role of primary motor cortex in goal-directed movements: insight from neurophysiological studies on non-human primates," *Curr Opin Nuurobiol*, vol. 13, pp. 671-677, 2003.

[2] A. B. Schwartz, "Cortical neural prosthetics," *Ann. Rev. Neurosci*, vol. 27, pp. 487-507, 2004.

[3] B.M. Yu, C. Kemere, G. Santhanam, A. Afshar, S.I. Ryu, T.H. Meng, M. Sahani M, K.V. Shenoy, "Mixture of trajectory models for neural decoding of goal-directed movements," *J Neurophysiol*, vol. 97, pp. 3763-3780, 2007.

[4] E.A. Pohlmeyer, S.A. Solla, E.J. Perreault, L.E. Miller, "Prediction of upper limb muscle activity from motor cortical discharge during reaching," *J Neural Eng.*, vol. 4, pp. 369-379, 2007.

[5] B.M. Yu, J.P.Cunningham, G. Santhanam, S.I. Ryu, K.V. Shenoy, M. Sahani, "Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity", *J Neurophysiol*, vol. 102, pp. 614-635, 2009.

[6] R.P. Dum, P.L. Strick, "The origin of corticospinal projections from the premotor areas in the frontal lobe," *J Neurosci*, vol. 11, pp. 667-669, 1991.

[7] L.E. Sergio, C. Hamel-Pâquet, J.F. Kalaska, "Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks," *J Neurophysiol*, vol. 94, pp. 2353-2378, 2005.

[8] D.M. Griffin, H.M. Hudson, A. Belhaj-Saïf, B.J. McKiernan, P.D. Cheney, "Do corticomotoneuronal cells predict target muscle EMG activity?," *J Neurophysiol*, vol. 99, pp. 1169-1186, 2007.

[9] M.M. Churchland, G. Santhanam, K.V. Shenoy, "Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach," *J Neurophysiol*, vol. 96, pp. 3130-3146, 2006.

[10] M. Sahan, "Latent Variable Models for Neural Data Analysis" (PhD thesis). Pasadena, CA: California Institute of Technology, 1999.

[11] M.M. Churchland, K.V. Shenoy, "Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex," *J Neurophysiol*, vol. 97, pp. 4235-4257, 2007.

[12] B.J. McKiernan, J.K. Marcario, J.H. Karrer, P.D. Cheney, "Corticomotoneuronal postspike effects in shoulder, elbow, wrist, digit, and intrinsic hand muscles during a reach and prehension task," *J Neurophysiol*, vol. 80, pp. 1961-1980, 1998.

[13] M.M. Morrow, L.E. Miller, "Prediction of muscle activity by populations of sequentially recorded primary motor cortex neurons," *J Neurophysiol*, vol. 89, pp. 2279-2288, 2003.