

Long-Term Stability of Motor Cortical Activity: Implications for Brain Machine Interfaces and Optimal Feedback Control

Robert D. Flint,^{1*} Michael R. Scheid,^{1*} Zachary A. Wright,¹ Sara A. Solla,^{2,4} and Marc W. Slutzky^{1,2,3,5}

¹Department of Neurology, ²Department of Physiology, ³Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, Illinois 60611, ⁴Department of Physics and Astronomy, Northwestern University, Evanston, Illinois 60208, and ⁵Rehabilitation Institute of Chicago, Chicago, Illinois 60611

The human motor system is capable of remarkably precise control of movements—consider the skill of professional baseball pitchers or surgeons. This precise control relies upon stable representations of movements in the brain. Here, we investigated the stability of cortical activity at multiple spatial and temporal scales by recording local field potentials (LFPs) and action potentials (multiunit spikes, MSPs) while two monkeys controlled a cursor either with their hand or directly from the brain using a brain–machine interface. LFPs and some MSPs were remarkably stable over time periods ranging from 3 d to over 3 years; overall, LFPs were significantly more stable than spikes. We then assessed whether the stability of all neural activity, or just a subset of activity, was necessary to achieve stable behavior. We showed that projections of neural activity into the subspace relevant to the task (the “task-relevant space”) were significantly more stable than were projections into the task-irrelevant (or “task-null”) space. This provides cortical evidence in support of the minimum intervention principle, which proposes that optimal feedback control (OFC) allows the brain to tightly control only activity in the task-relevant space while allowing activity in the task-irrelevant space to vary substantially from trial to trial. We found that the brain appears capable of maintaining stable movement representations for extremely long periods of time, particularly so for neural activity in the task-relevant space, which agrees with OFC predictions.

Key words: brain–machine interface; LFPs; minimum intervention; motor cortex; optimal feedback control; stability

Significance Statement

It is unknown whether cortical signals are stable for more than a few weeks. Here, we demonstrate that motor cortical signals can exhibit high stability over several years. This result is particularly important to brain–machine interfaces because it could enable stable performance with infrequent recalibration. Although we can maintain movement accuracy over time, movement components that are unrelated to the goals of a task (such as elbow position during reaching) often vary from trial to trial. This is consistent with the minimum intervention principle of optimal feedback control. We provide evidence that the motor cortex acts according to this principle: cortical activity is more stable in the task-relevant space and more variable in the task-irrelevant space.

Introduction

Animals perform complex learned movements countless times with little error. This consistent performance implies some level of stability within movement representations in the motor cortex. However, several experiments have also shown substantial

variability in individual neuronal representations (Padoa-Schioppa et al., 2004; Carmena et al., 2005; Rokni et al., 2007). Here, we face an unanswered question of fundamental importance: at what spatial and temporal scales do stable representations exist? More specifically, do neurons have a constant relationship with a given movement or is there some other mechanism through which highly variable neuronal signals produce consistent movements?

The stability of motor cortical signals is also an important issue in the design of brain–machine interfaces (BMIs), which translate brain signals into controls for external devices such as prosthetic arms. Stable BMI performance depends, at least partly, on the stability of the input signals. However, it remains unclear whether movement representations at the single-neuron level are stable over long time periods (months to years). Prior studies have found conflicting evidence of single-unit stability, with some showing variability (Carmena et al., 2005; Rokni et al.,

Received June 18, 2015; revised Feb. 17, 2016; accepted Feb. 20, 2016.

Author contributions: R.D.F., S.A.S., and M.W.S. designed research; R.D.F., M.R.S., and Z.A.W. performed research; R.D.F., M.R.S., S.A.S., and M.W.S. analyzed data; R.D.F., M.R.S., S.A.S., and M.W.S. wrote the paper.

This work was supported by the National Institutes of Health (Grant K08 NS060223 to M.W.S. and Grants R25 GM79300 and T32 HD057845 to M.R.S.) and the Defense Advanced Research Projects Agency (Grant N66001121-4023 to M.W.S.). We thank Lee Miller for the use of recording equipment.

The authors declare no competing financial interests.

*R.D.F. and M.R.S. contributed equally to this work.

Correspondence should be addressed to Marc W. Slutzky, MD, PhD, Department of Neurology, Northwestern University, 303 E. Superior Ave., Lurie 8-121, Chicago, IL 60611. E-mail: mslutzky@northwestern.edu.

DOI:10.1523/JNEUROSCI.2339-15.2016

Copyright © 2016 the authors 0270-6474/16/363623-10\$15.00/0

2007) and others showing stable properties over the course of a few hours (Stevenson et al., 2011), days (Chestek et al., 2007), or weeks (Ganguly and Carmena, 2009). Recently, numerous studies have used unsorted threshold crossings (multiunit spikes, or MSPs) in both scientific and neuroprosthetic experiments (Fraser et al., 2009; Chestek et al., 2011; Flint et al., 2013; Todorova et al., 2014), yet the level of stability of these signals has not been investigated. Local field potentials (LFPs) have long been postulated to be more stable than spikes (Pesaran et al., 2002). Indirect evidence of LFP stability over months to a year was suggested by stable ensemble decoding performance both during reaching (Chao et al., 2010) and in an online BMI (Flint et al., 2013). However, it is still desirable to obtain direct evidence of this stability in terms of the information that each individual signal provides.

Here, we examine directly the stability of LFPs and MSPs in the primary motor cortex (M1) of monkeys controlling a cursor using either arm movements (hand control) or BMI control over a period of months to years. Because the BMI translates M1 activity directly into cursor movement, it removes any variability that could be introduced along the efferent pathway between motor cortex and muscles. Therefore, we can assess directly the stability of M1 signals with respect to cursor movement. Overall, we find that LFPs, and to a lesser extent MSPs, maintain stable representations of cursor movement over very long time periods.

How can this finding be reconciled with reports of motor cortical variability? The framework of optimal feedback control (OFC) provides one potential explanation. OFC postulates that the brain controls movement in a given task by correcting only those errors that interfere with task objectives—the minimum intervention principle (MIP). This is achieved by limiting variability in a “task-relevant” space (the space relevant to objectives) and ignoring variability in a task-irrelevant or “task-null” space (Scholz and Schöner, 1999; Todorov and Jordan, 2002). Here, we investigated the differential stability of both LFPs and MSPs in task-relevant and task-null spaces. We found that, over the long term, stability is extremely high in the task-relevant space and lower in the task-null space. These findings provide important and novel cortical population-level evidence supporting OFC in the neural control of movement.

Materials and Methods

Experimental methods. All procedures were approved by the Northwestern University Institutional Animal Care and Use Committee. The behavioral training, surgical implantation, and recording methods for the current study were described in detail in Flint et al. (2013). Briefly, two adult rhesus macaques, one male and one female (Monkeys C and M), were each trained to use a two-link manipulandum to move a cursor (1-cm-diameter circle) within a rectangular planar workspace (20 × 20 cm) to sequential targets in a random target pursuit task (Fig. 1A). The choice of a random target task was motivated by our desire to sample more of the kinematic workspace than would be sampled by a center-out task. In each monkey, we surgically implanted a 96-electrode silicon array (Blackrock Microsystems) into the primary motor cortex contralateral to the arm used to control the cursor. We recorded LFPs and unsorted threshold crossings (which we term here MSPs) during the reaching behavior and built Wiener cascade decoders consisting of a linear Wiener filter cascaded with a third-order polynomial (Westwick et al., 2006) to relate end-point 2D velocity to MSP or LFP signals. LFP decoders were built using features extracted from the continuous LFP signal (Flint et al., 2012a). For each electrode, we calculated the local motor potential (LMP) by averaging the time domain signal in 50 ms bins (Schalk et al., 2007) and the spectral power in 5 frequency bands (0–4, 7–20, 70–110, 130–200, and 200–300 Hz) by using a fast Fourier transform on Hanning-windowed, 256 ms segments. These features, binned at 50 ms

intervals, were used to train a Wiener cascade BMI decoder. MSP decoders were built from spike counts binned at 50 ms. LFP and MSP decoders for BMI control were built on early epochs (10 min blocks) of hand-controlled reaching behavior and then held static until the end of all data collection for the entire study (hundreds of days). Monkeys C and M used each static decoder for BMI control typically two or three times per week. On most days, the monkeys had epochs of both hand control and one or more types of BMI control (Flint et al., 2013). The monkeys' arms used in hand control were not restrained during BMI control, but the manipulandum was removed. Their arm movements during BMI control were not stereotyped, differed from those when under hand control, and tended to attenuate slightly in amplitude over days.

The dataset used in this study includes BMI control data that were presented in a previous study (Flint et al., 2013), as well as hand control data. Days 180–380 (Monkey C) and 50–250 (Monkey M) of hand control were recorded concurrently with the BMI control study.

Stability of single feature decoders. We used two measures of stability in this study (see “Stability of ensemble tuning” section for details on the second measure). The first approach is based on single feature decoders (SFDs), which had been used earlier to evaluate single-unit stability over 1–2 d (Chestek et al., 2007). SFDs are constructed by treating each feature of each input signal as an independent input and calculating the ability of that single feature to predict the output. Here, we built SFDs using the data from each recording epoch over a period of 6 months for each feature (frequency band power or LMP) used in LFP control and for each electrode used in MSP control. SFDs were trained separately for hand control epochs, LFP control epochs, and MSP control epochs. After each SFD was trained, we calculated SFD prediction accuracy using that SFD to predict the actual velocity in five recording epochs from the end of the study's data collection period (which lasted months to years) and averaged the accuracy over those five epochs. Prediction accuracy was defined as the coefficient of determination (R^2) between the actual cursor velocity (either hand- or brain-controlled) and the velocity predicted by the SFD. Testing epochs were not used for SFD training and training epochs were never used for prediction accuracy testing. In the Results, we report SFD R^2 values for each day of the study. To calculate the prediction accuracy for a day, we computed the average R^2 obtained by using the SFDs trained on all epochs from that day. Examples of single-feature predictions are shown in Figure 1, C and D. Figure 2A shows SFD prediction accuracy for each day as a single column; for example, the leftmost column in Figure 2A shows the SFD prediction accuracy for all features on day 1 of the study. During hand control, SFD R^2 values measured the accuracy of predictions made by a single feature. During BMI control, the R^2 value for each SFD estimated the contribution of that feature to the BMI (population-level) decoder.

The entire hand control dataset spanned 1040 d for Monkey C and 800 d for Monkey M. Figure 2A shows an example color plot of R^2 values (for y -velocity) with N rows and D columns, where N is the number of included features and D is the number of days in the study. Each column represents the average performance for all SFDs trained on that day; the performance of each SFD is obtained as an average over the five testing epochs. We computed the correlation coefficient between each pair of columns in the $N \times D$ matrix of Figure 2A, producing the $D \times D$ color map of correlations shown in Figure 2B. The SFD stability index (SI_{SFD}) was then calculated as the mean of the correlation map across rows (or equivalently across columns), as seen in Figure 2C (see also Fig. 3C,F); this is similar to the measure in Ganguly and Carmena (2009). Overall, the color map (Fig. 2B) provides detailed information about the stability of the individual features and the SI_{SFD} provides a summary of stability over all features over time.

Features included in the analysis. In this study, we analyzed the performance of only those LFP features that were included in the decoder for LFP-based BMI control. When analyzing MSP stability, we included only those channels that corresponded to included LFP features so that stability indices calculated for both LFPs and MSPs would be based on signals from the same electrodes. As the multielectrode arrays aged, they developed leakage (or shunting) between some of the electrodes, likely due to insulation degradation (Barrese et al., 2013). This shunting manifested itself as an unusually high correlation in the signals recorded from pairs of electrodes, usually those that were physically adjacent on the array. In our analysis for

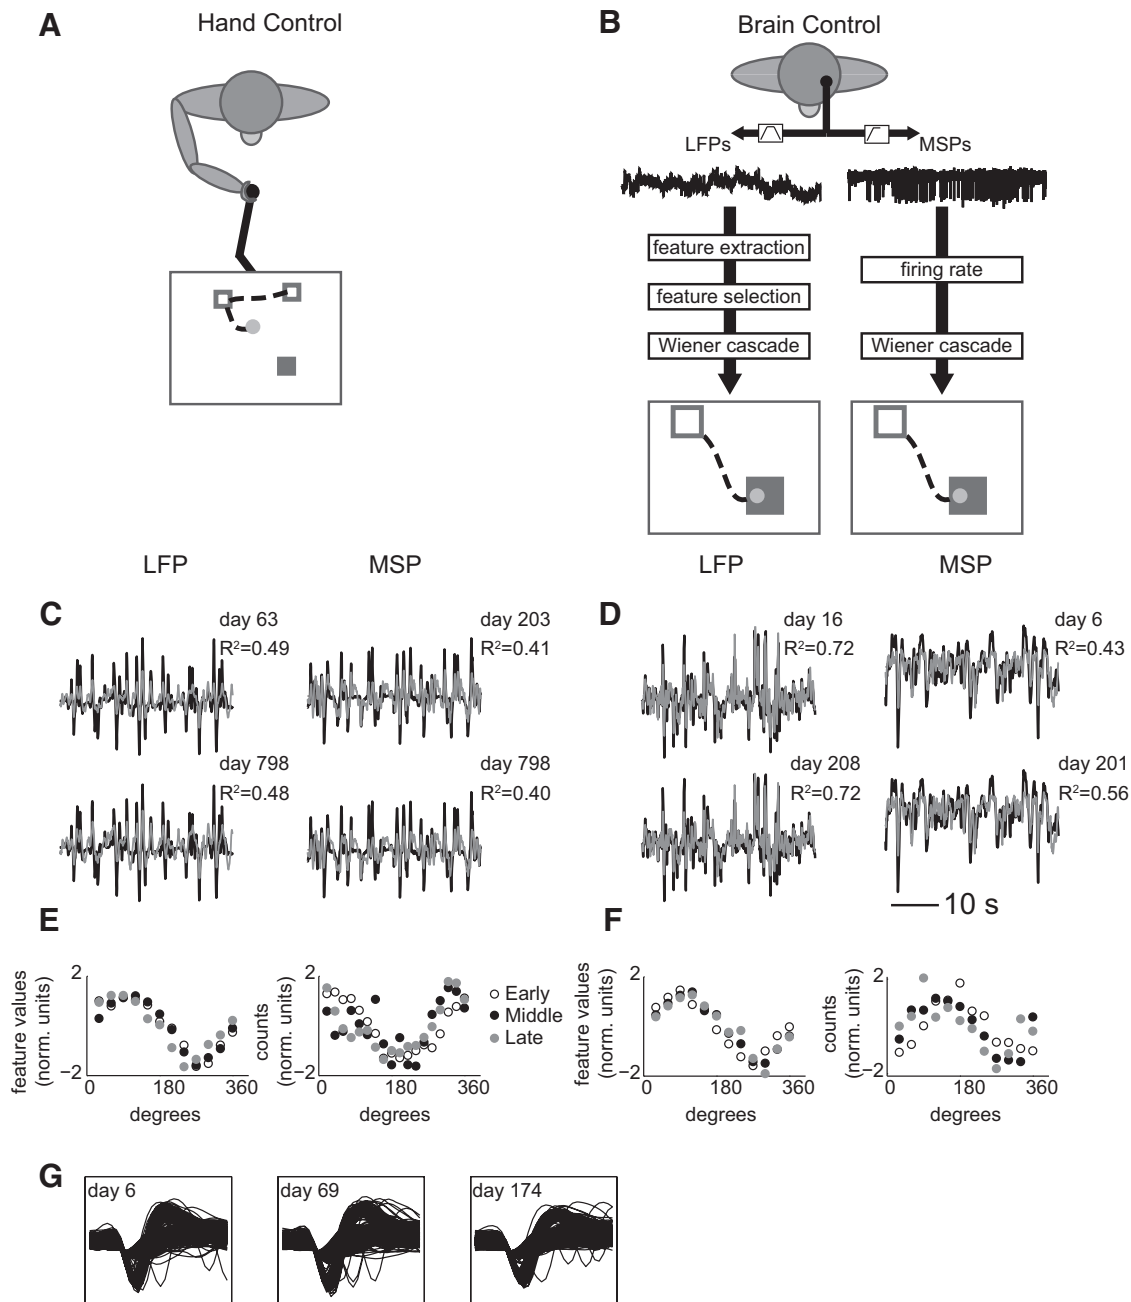


Figure 1. Random target pursuit behavior in hand control (**A**) and BMI control (**B**). BMI control was performed with either LFP features (**B**, left) or MSP firing rates (**B**, right). **C, D**, Brief examples of decoding actual Y velocity on the last day using single feature decoders built from 2 prior days (labeled “day” to the right of each trace) for hand (**C**) and brain (**D**) control. R^2 refers to the fit between the actual velocity (black) and the predicted velocity (gray). **E, F**, Example tuning curves in hand (**E**) and brain (**F**) control. As in **C** and **D**, LFP curves are on the left and MSP curves are on the right. Each plot shows three example curves as open circles or filled gray or black circles. For display, we normalized each tuning curve by subtracting its mean across bins and dividing by the SD. **G**, Example MSP waveforms from one representative electrode during days 6, 69, and 174 of recording.

this study, we removed from consideration all channels that were shunted with any other channel.

Stability of ensemble tuning. The second measure of stability used here was based on the directional tuning of individual features, as in Ganguly and Carmena (2009). Because movement velocities were in random directions, we evaluated movement direction every 100 ms and grouped these directions into 12 bins of 30° each. In each 10 min epoch, we computed the mean of each spike rate and LFP feature value over all movements that corresponded to each binned direction. This resulted in a tuning curve with 12 components for each feature/spike in each epoch. We computed the stability of the ensemble of tuning curves over all features/spikes used in BMI control (as in Fig. 3 of Ganguly and Carmena, 2009). In other words, we calculated the correlation coefficient

between the ensemble of tuning curves of all features from each day (averaged over all epochs in that day) and the ensemble of curves from all other days to produce a color map of correlations over days (as for SFDs in Fig. 2B). The ensemble tuning stability index (SI_{ET}) was then calculated as the mean of these correlations across days (Fig. 3, bottom).

To provide an intuition about the meaning of a given value of SI_{SFD} or SI_{ET} , we performed an analysis using simulated data. We started with an actual ensemble tuning curve (the first column from Fig. 4K), added noise with SD σ to n of N total features, and computed the SI value for each value of σ (ranging from 0 to 1) and n (ranging from 2 to N). We then plotted these SI values in a color plot (Fig. 5). We also calculated the σ that was actually observed in our results from the single features (Fig. 5, histogram overlay). Therefore, Figure 5 demonstrates how many features

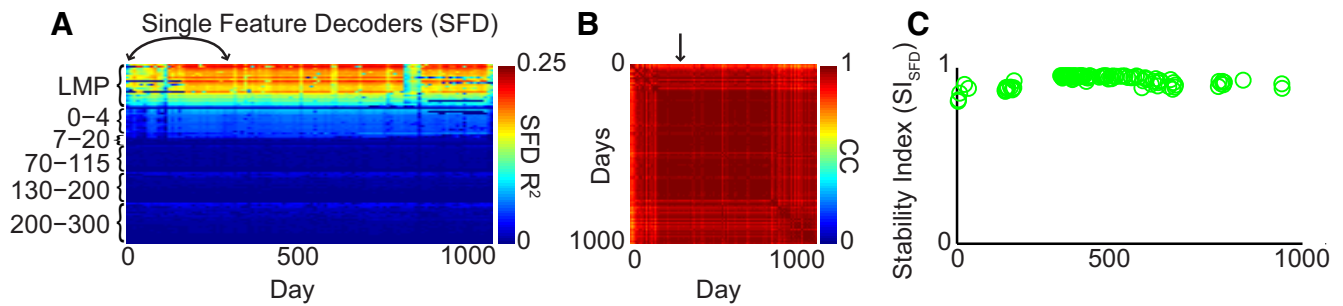


Figure 2. Calculation of SI_{SFD} , the mean pairwise correlation of LFP (or MSP) single-feature performance over all recording days. **A**, LFP SFD map for Monkey C during hand control. Color shows the R^2 value between actual and decoded movements. Features (rows) are sorted within each feature type in descending order of average SFD performance. **B**, Correlation map for LFP SFD R^2 performance. Each point in this color map represents the correlation coefficient between 2 d of the study (i.e., two columns of the maps in **A**). For example, the point in the top row of **B** designated with an arrow corresponds to the correlation of the two columns in **A** designated with a double arrow. **C**, SI_{SFD} over time for Monkey C; each circle is the average over all rows in a given column in the ensemble correlation map from **B**.

must vary, and by how much, to arrive at a particular value of the stability index. We repeated the simulation for SI_{SFD} with similar results.

Singular value decomposition analysis. To investigate the variability of neural activity during BMI control in a more principled manner, we performed a singular value decomposition (SVD) of the Wiener filter matrix H . This was an $M \times (10N)$ matrix, where M is the number of outputs (two in this case, for the two components of the velocity vector) and $10N$ is the number of inputs, given by the number N of neurons (or features) multiplied by the 10 time lags (the number of taps in the Wiener filter). At each time lag, the Wiener filter maps N inputs onto two outputs using the transform h , where h is the subset of H that corresponds to a single time lag. We used SVD to decompose h into a task-relevant space (of dimension two) and a task-null space (of dimension $N-2$). Because there were 10 time lags in our decoder, we calculated the SVD 10 times (once for each h) and combined them as described below.

For each time lag, SVD decomposed h into a left orthonormal matrix U , a scaling matrix Σ that contained the singular values along the diagonal, and a right orthonormal matrix V , such that:

$$h = U\Sigma V^*$$

Where V^* refers to the transpose of V . Intuitively, U and V can be thought of as rotation matrices and Σ as a scaling matrix. Therefore, SVD can be thought of as a sequence of three geometrical transformations: a rotation in input space (V^*), a scaling along the new axes (Σ), and a rotation in output space (U). This resulted in 10 V matrices, one for each lag, denoted as V^i , $1 \leq i \leq 10$. In each case, the matrix V^i captured the corresponding task-relevant neural space through its first two singular vectors; that is, its first two columns: v_j^i or $j = 1, 2$. The SVD has only two nonzero eigenvalues because the linear filter maps the N dimensional neural space into a 2D velocity space. For each lag i , the corresponding null space, is the one spanned by the remaining $N - 2$ singular vectors, the columns v_j^i for $3 \leq j \leq n$.

We projected each N -dimensional neural vector x , a vector of either firing rates (during MSP-based BMI control) or LFP features (during LFP-based BMI control), onto each of the singular vectors v_j^i . The projections α_j^i were thus defined as follows:

$$\alpha_j^i = x \cdot v_j^i$$

In a similar manner to the SFD analysis, for each BMI control epoch, we built two decoders of the BMI-controlled cursor velocity, one based on the first two projections $\{\alpha_1^i, \alpha_2^i\}$, for $1 \leq i \leq 10$, and one based on the remaining projections $\{\alpha_3^i, \dots, \alpha_N^i\}$, for $1 \leq i \leq 10$. We evaluated the performance of these decoders, one for each epoch, based on their ability to predict the BMI-controlled cursor velocity in the last five epochs—the same metric as in the SFD analysis. We measured the R^2 performance of each of these decoders and averaged this performance over all epochs in a given day. We performed this analysis in congruent control modes: only on MSPs for MSP-based BMI control and only on LFP features for LFP-

based BMI control. We used the coefficient of variation (CV) of the performance over all days to quantify variability.

Results

We recorded LFPs and threshold crossings (MSPs) for an extended time period while two monkeys controlled a computer cursor using either hand control with a 2D manipulandum (>3 years) or direct brain (BMI) control with LFPs or MSPs (~7 months). Table 1 summarizes the monkeys' behavioral performance during the hand control and BMI control tasks.

LFP and spike movement representations are stable during reaching

We first investigated how long neural representations of movement remain stable in the brain during natural motor tasks. SFD performance for all LFP features and MSP units (the ensemble “map”) in Monkey C is shown in Figure 3, **A** and **B**. During hand control, both signal types were highly stable in Monkey M (Fig. 3*D,E*), whereas MSPs were less stable than LFPs in Monkey C. This was true both on a single-feature level (each row in Fig. 3*A,D*) and for the ensemble as a whole (Fig. 3*C,F*). LFPs were significantly more stable than MSPs during hand control for Monkey C. The mean (\pm SD) of the SI_{SFD} over all days was 0.92 ± 0.05 for LFPs versus 0.77 ± 0.10 for MSP ($p = 10^{-33}$, t test). LFPs were also more stable than MSPs for Monkey M, by a small but statistically significant amount (SI_{SFD} 0.98 ± 0.02 vs 0.95 ± 0.02 ; $p = 10^{-31}$).

Ensemble tuning (ET), or directional tuning, of all LFP features and MSP units during hand control is shown in Figure 3, **G**, **H**, **J**, and **K**. The SI_{ET} for all LFP features and MSP units is shown in Figure 3*I* (Monkey C) and Figure 3*L* (Monkey M). LFPs were significantly more stable than MSPs in both monkeys during hand control. The mean (\pm SD) of SI_{ET} over all days for LFPs was 0.75 ± 0.20 and 0.38 ± 0.14 for MSPs for Monkey C ($p = 10^{-30}$). The mean of SI_{ET} over all days for LFPs was 0.79 ± 0.19 and 0.45 ± 0.08 for MSPs for Monkey M ($p = 10^{-30}$).

LFP and spike movement representations are stable during BMI control

We next analyzed the stability of each signal type during BMI control (Fig. 4). Similar to our results during hand control, we found that SI_{SFD} values indicated stable movement representations in both LFPs and MSPs (Fig. 4*C,F*). For Monkey C, LFPs were significantly more stable than MSPs (SI_{SFD} 0.94 ± 0.03 and 0.80 ± 0.09 , respectively; $p = 10^{-19}$). Monkey M exhibited very stable signals, with an

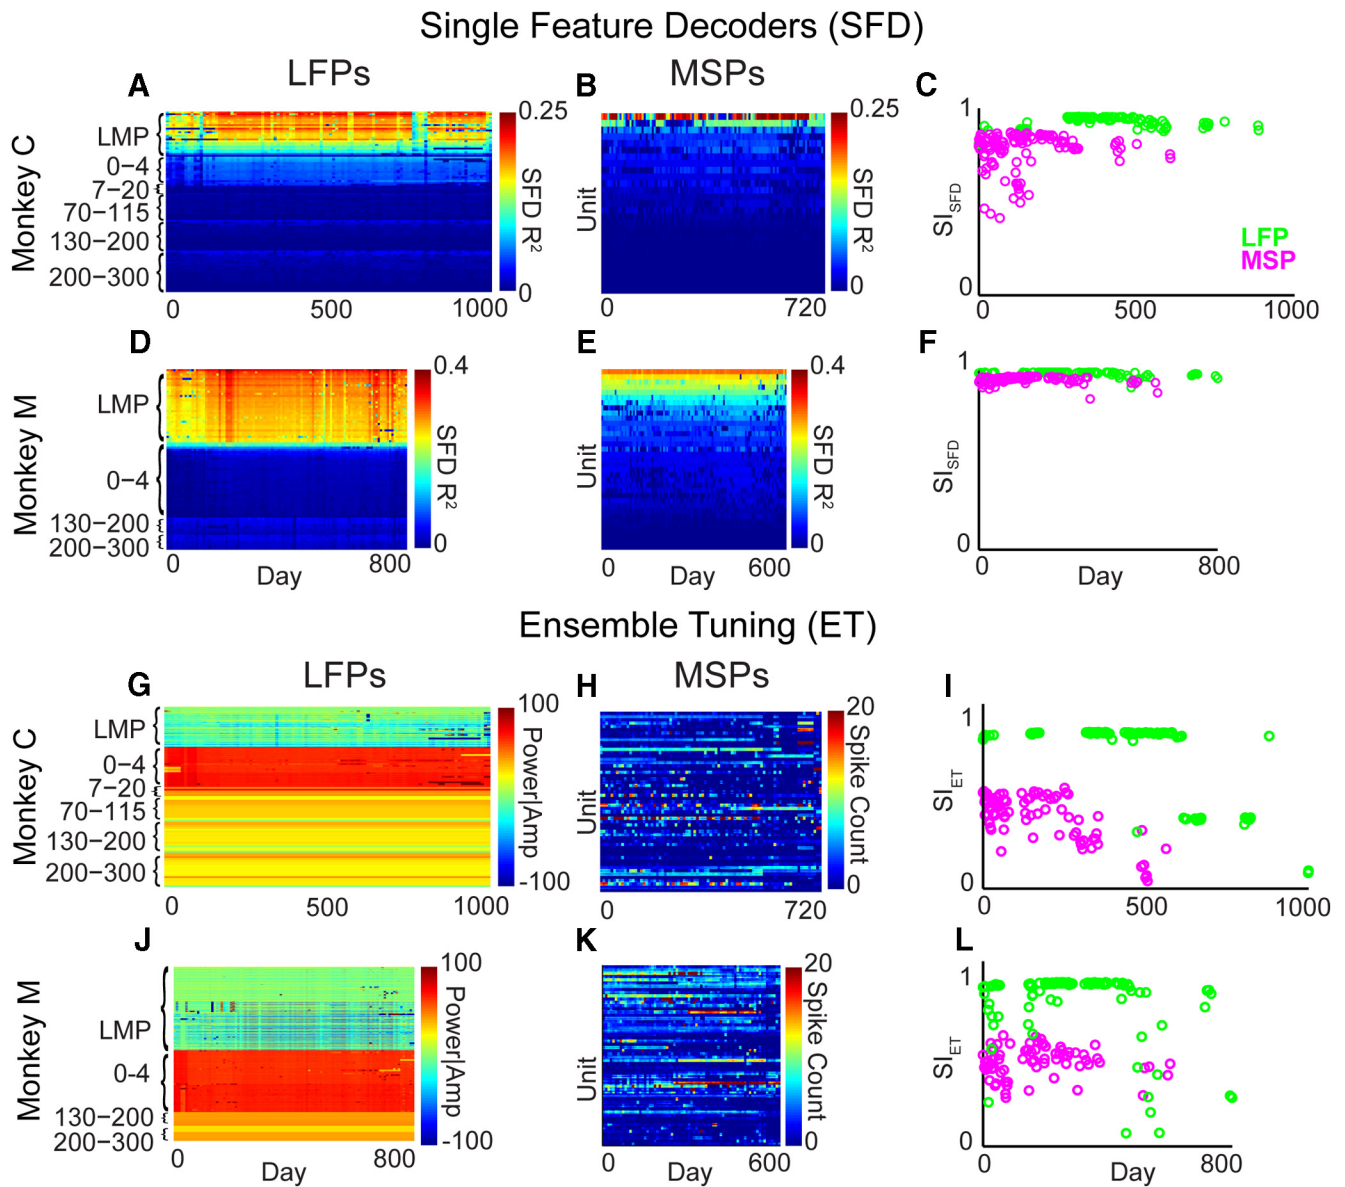


Figure 3. Stability of LFPs and MSPs during hand control. **A**, LFP SFD performance map for Monkey C. Rows are sorted within feature type by descending average SFD performance. **B**, MSP SFD performance map for Monkey C. **C**, SI_{SFD} values reflect high stability for LFPs (green circles) and MSPs (magenta circles) during hand control in Monkey C. **D**, **E**, Same as **A** and **B** but for Monkey M. **F**, Stability of LFPs was also slightly greater than that of MSPs in Monkey M. **G**, LFP ET map for Monkey C sorted by feature type. **H**, MSP ET map for Monkey C. **I**, Using SI_{ET}, LFPs were still highly stable and MSPs less stable. **J**, **K**, Same as **G** and **H** but for Monkey M. **L**, SI_{ET} showing that the stability of LFPs was also greater than that of MSPs in Monkey M.

SI_{SFD} of LFPs (0.99 ± 0.001) that was slightly but significantly greater than that of MSPs (0.98 ± 0.01 ; $p = 10^{-33}$). Using ensemble tuning, a subpopulation of spikes was still stable over time (Fig. 4H,K), but the population was notably less stable than LFPs. For Monkey C, SI_{ET} was 0.97 ± 0.01 and 0.51 ± 0.01 for LFPs and MSPs, respectively ($p = 10^{-81}$). For Monkey M, SI_{ET} was 0.97 ± 0.03 and 0.57 ± 0.02 ($p = 10^{-50}$).

To help clarify the meaning of SI_{ET} values, we computed the SI_{ET} as we introduced different levels of noise into a simulated MSP dataset (Fig. 5). As an example, a value of 0.5 for SI_{ET} means that if half the population (~30 features) is changing, then those signals are varying with an SD (σ) of ~0.11. Alternatively, if the average σ were ~0.2, an SI_{ET} of 0.5 would mean that ~20% of the population is changing. For reference, we have superimposed actual σ values taken from our data (Fig. 5, histogram overlay). Similar color maps and histograms were found for SI_{SFD} for MSPs as well as for LFPs (data not shown).

LFPs are highly stable in the time domain and high-frequency bands

The above results reflect the aggregate of multiple-frequency bands and time domain (LMP) LFP features. The high-gamma (70–300 Hz) band and LMP features contained most of the information about movement and contributed most to our decoders (Flint et al., 2012b). Therefore, we investigated the stability of the LMP and the three high gamma bands (γ_1 : 70–115 Hz, γ_2 : 130–200 Hz, and γ_3 : 200–300 Hz) that were used in the BMI decoders by the monkeys during LFP-based BMI control (Fig. 6). As anticipated from the monkeys' stable BMI performance, the LMP and high-gamma features were highly stable. This mirrored the stability of the overall population (cf. mean values over time in Fig. 4C). Interestingly, the γ_2 and γ_3 bands were even more stable than the γ_1 band for Monkey C. This could suggest that high-gamma signals, in particular, might provide highly stable signals for studying neurophysiology and for controlling BMIs.

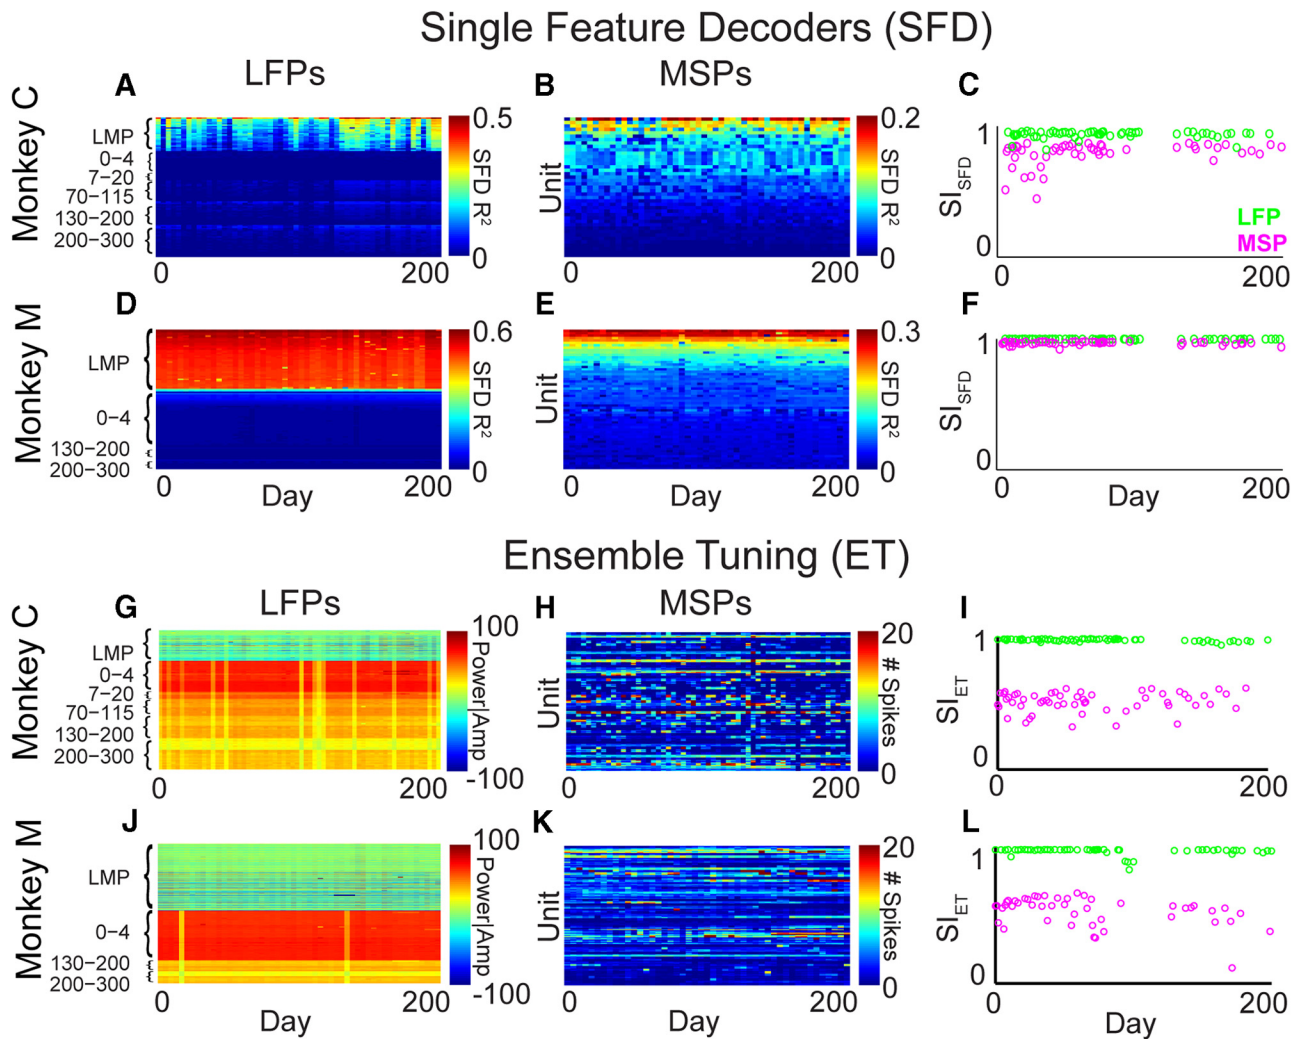


Figure 4. Stability of LFPs and MSPs during BMI control. **A**, Monkey C's LFP SFD performance (R^2) during LFP BMI control. **B**, Monkey C's MSP SFD performance during MSP BMI control. **C**, SI_{SFD} values reflect high stability for LFPs (green circles) and MSPs (magenta circles) in Monkey C. **D**, **E**, Same as **A** and **B** but for Monkey M. **F**, LFP stability was also higher than MSP stability in Monkey M. **G**, LFP ET map for Monkey C sorted by feature type. **H**, MSP ET map for Monkey C. **I**, Using SI_{ET} , LFPs remained highly stable, whereas MSPs were less stable. **J**–**L**, Same as **G**–**I** but for Monkey M. SI_{ET} shows that the stability of LFPs was noticeably higher than that of MSPs during BMI control.




Figure 5. Simulated data show changes in SI_{ET} with number and variability of features. The color map shows a surface of SI_{ET} values as we systematically changed the number of simulated features that were influenced by noise (y -axis), and the SD (σ) of the noise (x -axis). The histogram overlay shows the σ values for the actual single feature ensemble tuning (Monkey M, MSP control, taken from Fig. 4K).

LFP stability increases during early BMI decoder learning

The prior analyses examined the stability of LFPs during long-term BMI use. To gain insight into what happens during the early learning of a new LFP-based BMI decoder and to facilitate com-

Table 1. Summary statistics of reaching behavior (median \pm interquartile range) in each epoch

	Reaches/epoch	Peak speed (cm/s)
Monkey C		
Hand control	659 \pm 137	27 \pm 13
LFP BMI control	85 \pm 22	35 \pm 20
MSP BMI control	121 \pm 33	31 \pm 20
Monkey M		
Hand control	561 \pm 92	23 \pm 14
LFP BMI control	71 \pm 17	38 \pm 37
MSP BMI control	164 \pm 58	22 \pm 16

parison with prior studies, we also examined the stability over a short timescale. We selected data from the first 3 d of Monkey C's use of an LFP-based BMI decoder different from the one used in the above analyses. We calculated the SI_{SFD} by dividing the data in those 3 d into 80 epochs of ~ 2 min each and performing the SFD stability analysis as described previously. The LFPs stabilized on day 2, after an accumulated task time of ~ 50 min. This was followed by an extended period of high stability across days 2 and 3 (Fig. 7). This timescale of days for achieving stability is similar

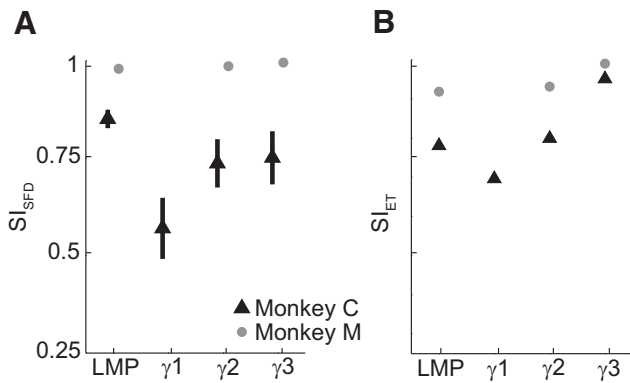


Figure 6. Stability of selected LFP feature types during LFP-based BMI control. **A**, Mean (\pm SD) SI_{SFD} over all days for Monkey C (triangles) and Monkey M (circles). **B**, Mean SI_{ET} for Monkey C (triangles) and M (circles). Some error bars are not visible because they are smaller than the marker sizes. Overall, LMP and high-gamma features were highly stable.




Figure 7. LFP stability increases quickly after the introduction of a new decoder. After the introduction of a new decoder, the SI_{SFD} increased sharply after \sim 45 min of BMI control. Dashed lines denote boundaries between days of recording.

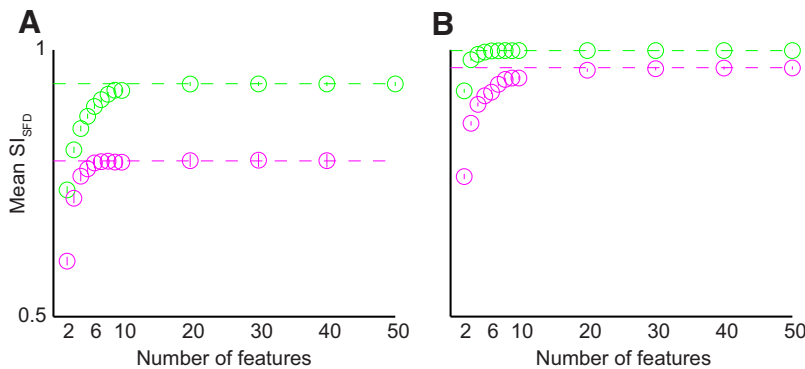


Figure 8. LFPs (green points) and MSPs (magenta points) achieved stability with a small number of features. **A**, **B**, Feature-dropping curves for Monkey C (**A**) and Monkey M (**B**). Each point represents SI_{SFD} values for a feature set of size F (horizontal axis), averaged across days. Vertical lines show SD values, which were very small. For both LFPs and MSPs, the SI_{SFD} indicates remarkable stability even with very few features, approaching its population level (dashed horizontal lines) with the inclusion of as few as 10–20 features.

to the timescale of map stabilization found using small ensembles of spikes to control a BMI (Ganguly and Carmena, 2009).

Stability measures remain high even when using very few features

Because we computed stability by simultaneously using a large number of features with more LFP features than spikes, one possible concern could be that stability was overestimated due to averaging effects. We therefore sought to determine how many features must

be included to approach the level of stability achieved by the entire population. We conducted a feature-dropping analysis to answer this question. We randomly sampled F features from the population of LFP or MSP features, where F varied from 1 to 50. For each group of features, we calculated the SI_{SFD} and repeated this 500 times. We found that a small number of features, <20 in each case, sufficed to achieve the SI_{SFD} found for the entire population (Fig. 8). Therefore, our metric was unlikely to overestimate stability due to the large number of features in the ensemble.

Cortical stability is higher in the task-relevant space than in the null space

The above analyses showed that LFPs and MSPs were stable overall, yet some neurons and features were more stable than others (Fig. 3B,E). To investigate these differences in a principled manner, we examined the projections of LFP features and spikes into two neural subspaces: task-relevant and task-null subspaces. We hypothesized that, if the cortex used a strategy like the MIP of OFC to control the BMI, then the activity in the task-null space (or neural null space) should be much less stable than that in the task-relevant space.

The mean R^2 of the $\alpha_{1,2}$ decoders (task-relevant space) was much higher than that of the $\alpha_{3,\dots,N}$ decoders (Fig. 9C–F, null space). Stability of the decoders based on task-relevant space features was not surprising because, by definition, the task-relevant space should be an excellent predictor of the full Wiener filter output. What was not certain, however, was whether the variability of the predictions would be lower in the task-relevant space than in the task-null space (Fig. 9A). An alternative possibility could have been for the variability to be equally low in both spaces, which would have indicated that all components of cortical activity were controlled with equal precision regardless of task relevance (Fig. 9B).

We found that the task-relevant SFD performance showed low variability, as measured by the CV ($CV = 0.010$ and 1.8×10^{-4} for Monkeys C and M, respectively; see Fig. 9C,D). This would be expected from the stable and accurate BMI performance previously reported by Flint et al. (2013). In contrast, the null space SFD performance had much higher variability ($CV = 0.23$ and 0.0026 , respectively, using spikes). Similar, but smaller, differences were found for LFPs between null space ($CV = 0.060$ and 0.091) and task-relevant space ($CV = 10^{-5}$ and 6.7×10^{-3}) SFD performances (Fig. 9E,F).

Principal components analysis of the LFP features used in the decoders showed that 13 (Monkey C) and 17 (Monkey M) principal components were needed to account for 80% of the variance in these features. This suggested that LFP activity in the null space did indeed covary substantially with that in the task-relevant space. This would explain the higher performance and lower variability of the null space predictions for LFPs than for spikes.

We next tested whether any pair of null space projections provided as stable an input to a decoder as the pair of task relevant projections. This allowed us to determine whether the large difference in the number of α inputs to the decoder (2 vs 83) biased the stability of the decoding. We repeated the decoder calculation

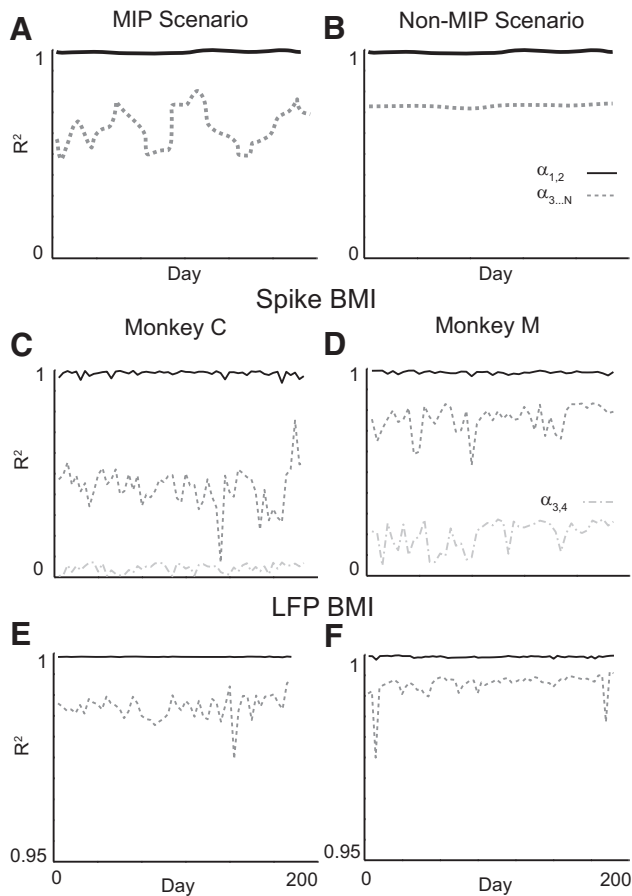


Figure 9. Stability in the task-relevant and task-null neural spaces. **A**, Expected decoding performance if cortical activity adheres to the MIP in the task-relevant (using the $\alpha_{1,2}$ decoder; solid line) and null spaces (using the $\alpha_{3...N}$ decoder; dashed line). Mean null space performance is lower than that in the task-relevant space and null space performance is much more variable. **B**, Expected decoding performance if cortical activity does not adhere to the MIP. Here, null space performance has lower mean but the same variance as performance in the task-relevant space. **C, D**, Actual decoding performance in the task-relevant ($\alpha_{1,2}$), full task-null ($\alpha_{3...N}$), and partial task-null ($\alpha_{3,4}$) spaces using MSP-based BMI data in Monkeys C and M, respectively. The task-relevant space performance was extremely high and stable, whereas the null space performances were much more variable. The same was true for LFP-based BMI (**E, F**), but the null space decoder in this case had much higher performance than for the MSP-based BMI.

using a random subsampling of two α projections in the null space of LFP control in Monkey C. We repeated this random subsampling 200 times to obtain a distribution of CVs and compared it with the CV value obtained using $\{\alpha_1, \alpha_2\}$ as decoder input. The entire distribution of predictions based on 2D null space projections had significantly higher CVs than those based on task-relevant space projections ($p = 0.005$ using the cumulative empirical distribution; Fig. 10). Therefore, the increased variability found in the null space was not due to bias from the higher number of α inputs to the decoder.

Discussion

The long-term stability of individual MSPs and LFPs in the motor cortex had not been established previously. We found LFPs to be highly stable, and MSPs somewhat less stable, over time scales ranging from 3 d to 3 years. In general, LFPs and MSPs behaved stably with respect to cursor movement controlled by either the hand or the brain. LFPs were somewhat more stable than MSPs during hand control and substantially more stable than MSPs during BMI control. In particular, we found that signal stability




Figure 10. Variability in the null space versus the task-relevant space. The CVs of the decoding performance over time for decoders based on two randomly selected null-space projections (black bars) were significantly higher than the CV for a decoder based on the two task-relevant projections (dashed line).

was much higher in the task-relevant space than the task-null space, an observation consistent with the hypothesis of optimal feedback control at the cortical level.

Stability at different spatial scales

Stability of movement representation would seem to be important for overlearned, daily activities such as reaching for food. However, it has been debated whether this high behavioral stability requires highly stable activity at the neuronal level (Chestek et al., 2007; Ganguly and Carmena, 2009; Stevenson et al., 2011) or if stability is a network-level phenomenon that arises despite unstable representations in single neurons (Cohen and Nicolelis, 2004; Carmena et al., 2005; Rokni et al., 2007). Here, we found evidence for stability of both MSPs and LFPs at the single-feature level. Stability was higher at the mesoscopic (LFP) scale than at the MSP scale and a subset of MSPs were much less stable. This could suggest that the brain is capable of different levels of control on mesoscopic (LFPs) and single-neuron scales (Ganguly and Carmena, 2009), although it could also be due to signal averaging at the mesoscopic scale.

Some apparent contradictions about stability in published reports are likely due to differences in methods or analysis. For example, Rokni et al. (2007) found variable preferred directions (PDs) during a movement task, but measurement noise may have played an important role in the apparent amount of PD variation in that study (Stevenson et al., 2011). Two prior studies found lower stability in ensemble-level, offline decoding (Flint et al., 2013; Perge et al., 2014) than we saw previously with online BMI control (Flint et al., 2013) or here using single-feature analyses. In addition, movement representations may change with learning of a task (Cohen and Nicolelis, 2004; Chestek et al., 2007; Mandelblat-Cerf et al., 2009). Ganguly and Carmena (2009) showed that this phenomenon also occurs when learning a BMI task using a small ensemble of neurons while requiring the hand to remain stationary. In our study, as in Chestek et al. (2007), monkeys performed overlearned control tasks. In addition, our BMI decoders were biomimetic (trained on actual reaches) and the monkeys' arms were not required to remain stationary during BMI control. These differences may account for the minimal and rapid early learning required for our BMI task based on both behavioral performance (Flint et al., 2013) and on SI_{SFD} (Fig. 7). These aspects of our experiment also may have contributed to the highly stable movement representations that we found.

One potential concern with our approach could be that our stability indices were calculated across a large number of features (or spikes), so the actual stability of individual signals may be overestimated. To validate our population-level analysis, we repeatedly eliminated randomly selected inputs and recalculated SI_{SFD} (“feature dropping”). We found that the high degree of signal stability observed for LFPs and MSPs was preserved even when using a very low number of features to compute SI_{SFD} (Fig. 8). In some cases, even as few as five to six features provided remarkably high stability and, generally, 20 features sufficed to achieve population-level stability. We found a similar trend for hand control results (data not shown).

Evidence supporting the minimal intervention principle

According to the MIP, the brain minimizes effort by controlling only those movement components that are relevant to the task goals while allowing variability in those movement components that are unrelated to the task goals. This principle accounts for the simultaneous observation of highly accurate control of movements in a particular task with substantial trial-to-trial variability in the way the limb or body moves while achieving the task goal (Todorov, 2004). The principle predicts low variability in the task-relevant space and higher variability in the task-null space, a phenomenon that can be explained by optimal feedback control (Todorov, 2004) among other models (Martin et al., 2009). Substantial evidence for this concept was found in modeling and behavioral studies (Pandy et al., 1990; Scholz and Schöner, 1999; Latash et al., 2001; Todorov and Jordan, 2002; Scott, 2004) and in a myoelectric interface study (Nazarpour et al., 2012), although some evidence against it also exists (Mosier et al., 2005; de Rugy et al., 2012). Cortical electrophysiological evidence supporting the OFC framework, but not the MIP, was presented previously in the context of long-latency stretch responses (Pruszynski and Scott, 2012). In the current study, direct BMI control using signals from primary motor cortex provided a unique opportunity to study cortical control strategies during voluntary, goal-oriented behavior without the intervening levels of the neuromuscular system (spinal cord, muscle activations, joint biomechanical constraints, etc.). Although somatosensory feedback was intact during BMI control, it did not inform the monkeys about their performance in the BMI task. Therefore, we were able to assess directly the relative variabilities in the “neural” null and task spaces.

Neural patterns during BMI control were extremely stable in the task-relevant space and more variable in the task-null space. Although we expected that the task-relevant space would be high performing and stable and that the null space would be lower performing on average, there was no a priori reason that the variability in the null space had to be greater. For example, the brain could have controlled all of the neural signal components with equal precision and this would have resulted in a neural null space with low mean and low variance in performance (Fig. 9B). Instead, we found low mean but high variance in the neural null space performance (Fig. 9A). This finding is consistent with the MIP.

This result was not due to bias from unequal dimensionality of the input space. None of the 2D combinations of neural patterns in the null space was as stable as those in the task-relevant space (Fig. 10). These findings could provide an explanation for the conflicting evidence of neuronal stability presented in some prior studies. More importantly, this result provides evidence in support of the minimal intervention principle in the motor cortex itself. More direct evidence of cortical signals’ consistency with

the OFC framework would require an examination of trial-by-trial variability in the task and null spaces. Because our experiments were not designed to track trial-to-trial variability, further investigation of short-term variability in the neural null and task spaces is warranted.

Implications for decoding motor intent over long time periods

By leveraging the ability to decode limb kinematics (Georgopoulos et al., 1982; Moran and Schwartz, 1999) or kinetics (Evarts, 1968; Fagg et al., 2007; Flint et al., 2014) from cortical activity, BMIs have the potential to restore lost function to individuals with paralysis (Hochberg et al., 2006; Collinger et al., 2013; Ethier et al., 2012). However, single-neuron action potentials are difficult to record over long time scales with silicon-based arrays and those that are recorded over long times can exhibit substantial nonstationarity in their amplitudes and waveforms (Chestek et al., 2011; Simeral et al., 2011). This variability represents a significant roadblock to clinically viable BMIs. These issues have fueled a push toward examining unsorted spikes (MSPs, also called threshold crossings) and LFPs for decoding movement (Pesaran et al., 2002; Mehring et al., 2003; Stark and Abeles, 2007; Jarosiewicz et al., 2008; Zhuang et al., 2010; Bansal et al., 2011; Chestek et al., 2011; Slutzky et al., 2011; Flint et al., 2012a; Flint et al., 2012b; Gilja et al., 2012; Flint et al., 2013; Nuyujukian et al., 2014; Perge et al., 2014; Todorova et al., 2014). The results reported here, and our prior results showing high long-term BMI performance (Flint et al., 2013), indicate that it should be possible to decode intended movement accurately over clinically useful time spans, limited only by recording hardware rather than by inherent instabilities in the neural signals. Although LFPs were more stable than MSPs, this did not translate to higher BMI performance (Flint et al., 2013), possibly because spikes were slightly more informative overall. Adaptive decoding algorithms (Gilja et al., 2012; Orsborn et al., 2014) may provide a way to counteract changes that do occur in neural representations. Further, adaptation to BMI decoders is greatly facilitated when the cortical modulations required to achieve control are within the “intrinsic manifold” of the cortical activity (Sadler et al., 2014). This result was based solely on the analysis of neural activity and did not take into account the decoder. The findings of Sadler et al. (2014), when combined with the results reported here, suggest that a more judicious design of BMIs may make them both easier to learn and more stable to use over long periods of time.

References

- Bansal AK, Vargas-Irwin CE, Truccolo W, Donoghue JP (2011) Relationships among low-frequency local field potentials, spiking activity, and three-dimensional reach and grasp kinematics in primary motor and ventral premotor cortices. *J Neurophysiol* 105:1603–1619. [CrossRef Medline](#)
- Barrese JC, Rao N, Paroo K, Triebwasser C, Vargas-Irwin C, Franquemont L, Donoghue JP (2013) Failure mode analysis of silicon-based intracortical microelectrode arrays in non-human primates. *J Neural Eng* 10:066014. [CrossRef Medline](#)
- Carmena JM, Lebedev MA, Henriquez CS, Nicolelis MA (2005) Stable ensemble performance with single-neuron variability during reaching movements in primates. *J Neurosci* 25:10712–10716. [CrossRef Medline](#)
- Chao ZC, Nagasaka Y, Fujii N (2010) Long-term asynchronous decoding of arm motion using electrocorticographic signals in monkeys. *Front Neuroeng* 3.
- Chestek CA, Batista AP, Santhanam G, Yu BM, Afshar A, Cunningham JP, Gilja V, Ryu SI, Churchland MM, Shenoy KV (2007) Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci* 27:10742–10750. [CrossRef Medline](#)
- Chestek CA, Gilja V, Nuyujukian P, Foster JD, Fan JM, Kaufman MT, Churchland MM, Rivera-Alvidrez Z, Cunningham JP, Ryu SI, Shenoy KV

- (2011) Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *J Neural Eng* 8:045005. [CrossRef Medline](#)
- Cohen D, Nicolelis MA (2004) Reduction of single-neuron firing uncertainty by cortical ensembles during motor skill learning. *J Neurosci* 24:3574–3582. [CrossRef Medline](#)
- Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJ, Velliste M, Boninger ML, Schwartz AB (2013) High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet* 381:557–564. [Medline](#)
- de Rugy A, Loeb GE, Carroll TJ (2012) Muscle coordination is habitual rather than optimal. *J Neurosci* 32:7384–7391. [CrossRef Medline](#)
- Ethier C, Oby ER, Bauman MJ, Miller LE (2012) Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature* 485:368–371. [CrossRef Medline](#)
- Evarts EV (1968) Relation of pyramidal tract activity to force exerted during voluntary movement. *J Neurophysiol* 31:14–27. [Medline](#)
- Fagg AH, Hatsopoulos NG, de Lafuente V, Moxon KA, Nemati S, Rebesco JM, Romo R, Solla SA, Reimer J, Tkach D, Pohlmeier EA, Miller LE (2007) Biomimetic brain machine interfaces for the control of movement. *J Neurosci* 27:11842–11846. [CrossRef Medline](#)
- Flint RD, Ethier C, Oby ER, Miller LE, Slutzky MW (2012a) Local field potentials allow accurate decoding of muscle activity. *J Neurophysiol* 108:18–24. [CrossRef Medline](#)
- Flint RD, Lindberg EW, Jordan LR, Miller LE, Slutzky MW (2012b) Accurate decoding of reaching movements from field potentials in the absence of spikes. *J Neural Eng* 9:046006. [CrossRef Medline](#)
- Flint RD, Wright ZA, Scheid MR, Slutzky MW (2013) Long term, stable brain machine interface performance using local field potentials and multiunit spikes. *J Neural Eng* 10.
- Flint RD, Wang PT, Wright ZA, King CE, Krucoff MO, Schuele SU, Rosenow JM, Hsu FP, Liu CY, Lin JJ, Sazgar M, Millett DE, Shaw SJ, Nenadic Z, Do AH, Slutzky MW (2014) Extracting kinetic information from human motor cortical signals. *Neuroimage* 101:695–703. [CrossRef Medline](#)
- Fraser GW, Chase SM, Whitford A, Schwartz AB (2009) Control of a brain-computer interface without spike sorting. *J Neural Eng* 6:055004. [CrossRef Medline](#)
- Ganguly K, Carmena JM (2009) Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol* 7:e1000153. [CrossRef Medline](#)
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2:1527–1537. [Medline](#)
- Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Churchland MM, Kaufman MT, Kao JC, Ryu SI, Shenoy KV (2012) A high-performance neural prosthesis enabled by control algorithm design. *Nat Neurosci* 15:1752–1757. [CrossRef Medline](#)
- Hochberg LR, Serruya MD, Friebs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442:164–171. [CrossRef Medline](#)
- Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB (2008) Functional network reorganization during learning in a brain-computer interface paradigm. *Proc Natl Acad Sci U S A* 105:19486–19491. [CrossRef Medline](#)
- Latash ML, Scholz JF, Danion F, Schöner G (2001) Structure of motor variability in marginally redundant multifinger force production tasks. *Exp Brain Res* 141:153–165. [CrossRef Medline](#)
- Mandelblat-Cerf Y, Paz R, Vaadia E (2009) Trial-to-trial variability of single cells in motor cortices is dynamically modified during visuomotor adaptation. *J Neurosci* 29:15053–15062. [CrossRef Medline](#)
- Martin V, Scholz JP, Schöner G (2009) Redundancy, self-motion, and motor control. *Neural Comput* 21:1371–1414. [CrossRef Medline](#)
- Mehring C, Rickert J, Vaadia E, Cardoso de Oliveira S, Aertsen A, Rotter S (2003) Inference of hand movements from local field potentials in monkey motor cortex. *Nat Neurosci* 6:1253–1254. [CrossRef Medline](#)
- Moran DW, Schwartz AB (1999) Motor cortical activity during drawing movements: population representation during spiral tracing. *J Neurophysiol* 82:2693–2704. [Medline](#)
- Mosier KM, Scheidt RA, Acosta S, Mussa-Ivaldi FA (2005) Remapping hand movements in a novel geometrical environment. *J Neurophysiol* 94:4362–4372. [CrossRef Medline](#)
- Nazarpour K, Barnard A, Jackson A (2012) Flexible cortical control of task-specific muscle synergies. *J Neurosci* 32:12349–12360. [CrossRef Medline](#)
- Nuyujukian P, Kao JC, Fan JM, Stavisky SD, Ryu SI, Shenoy KV (2014) Performance sustaining intracortical neural prostheses. *J Neural Eng* 11:066003. [CrossRef Medline](#)
- Orsborn AL, Moorman HG, Overduin SA, Shانهchi MM, Dimitrov DF, Carmena JM (2014) Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron* 82:1380–1393. [CrossRef Medline](#)
- Padoa-Schioppa C, Li CS, Bizzi E (2004) Neuronal activity in the supplementary motor area of monkeys adapting to a new dynamic environment. *J Neurophysiol* 91:449–473. [Medline](#)
- Pandy MG, Zajac FE, Sim E, Levine WS (1990) An optimal control model for maximum-height human jumping. *J Biomech* 23:1185–1198. [CrossRef Medline](#)
- Perge JA, Zhang S, Malik WQ, Homer ML, Cash S, Friebs G, Eskandar EN, Donoghue JP, Hochberg LR (2014) Reliability of directional information in unsorted spikes and local field potentials recorded in human motor cortex. *J Neural Eng* 11:046007. [CrossRef Medline](#)
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5:805–811. [CrossRef Medline](#)
- Pruszynski JA, Scott SH (2012) Optimal feedback control and the long-latency stretch response. *Exp Brain Res* 218:341–359. [CrossRef Medline](#)
- Rokni U, Richardson AG, Bizzi E, Seung HS (2007) Motor learning with unstable neural representations. *Neuron* 54:653–666. [CrossRef Medline](#)
- Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP (2014) Neural constraints on learning. *Nature* 512:423–426. [CrossRef Medline](#)
- Schalk G, Kubánek J, Miller KJ, Anderson NR, Leuthardt EC, Ojemann JG, Limbrick D, Moran D, Gerhardt LA, Wolpaw JR (2007) Decoding two-dimensional movement trajectories using electrocorticographic signals in humans. *J Neural Eng* 4:264–275. [CrossRef Medline](#)
- Scholz JP, Schöner G (1999) The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126:289–306. [CrossRef Medline](#)
- Scott SH (2004) Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci* 5:532–546. [Medline](#)
- Simeral JD, Kim SP, Black MJ, Donoghue JP, Hochberg LR (2011) Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array. *J Neural Eng* 8:025027. [CrossRef Medline](#)
- Slutzky MW, Jordan LR, Lindberg EW, Lindsay KE, Miller LE (2011) Decoding the rat forelimb movement direction from epidural and intracortical field potentials. *J Neural Eng* 8:036013. [CrossRef Medline](#)
- Stark E, Abeles M (2007) Predicting movement from multiunit activity. *J Neurosci* 27:8387–8394. [CrossRef Medline](#)
- Stevenson IH, Cherian A, London BM, Sachs NA, Lindberg E, Reimer J, Slutzky MW, Hatsopoulos NG, Miller LE, Kording KP (2011) Statistical assessment of the stability of neural movement representations. *J Neurophysiol* 106:764–774. [CrossRef Medline](#)
- Todorova S, Sadtler P, Batista A, Chase S, Ventura V (2014) To sort or not to sort: the impact of spike-sorting on neural decoding performance. *J Neural Eng* 11:056005. [CrossRef Medline](#)
- Todorov E (2004) Optimality principles in sensorimotor control. *Nat Neurosci* 7:907–915. [CrossRef Medline](#)
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235. [CrossRef Medline](#)
- Westwick DT, Pohlmeier EA, Solla SA, Miller LE, Perreault EJ (2006) Identification of multiple-input systems with highly coupled inputs: application to EMG prediction from multiple intracortical electrodes. *Neural Comput* 18:329–355. [CrossRef Medline](#)
- Zhuang J, Truccolo W, Vargas-Irwin C, Donoghue JP (2010) Decoding 3-D reach and grasp kinematics from high-frequency local field potentials in primate primary motor cortex. *IEEE Trans Biomed Eng* 57:1774–1784. [CrossRef Medline](#)