Edinburgh Research Explorer

Contribution of LFP dynamics to single-neuron spiking variability in motor cortex during movement execution

OPEN ACCESS

Contribution of LFP dynamics to single-neuron spiking variability in motor cortex during movement execution

Michael E. Rule [1]*, Carlos Vargas-Irwin [1], John P. Donoghue [1, 2, 3] and Wilson Truccolo [1, 2, 3]*

[1] Department of Neuroscience, Brown University, Providence, RI, USA, [2] Institute for Brain Science, Brown University, Providence, RI, USA, [3] Department of Veterans Affairs, Center for Neurorestoration and Neurotechnology, Brown University, Providence, RI, USA

Understanding the sources of variability in single-neuron spiking responses is an important open problem for the theory of neural coding. This variability is thought to result primarily from spontaneous collective dynamics in neuronal networks. Here, we investigate how well collective dynamics reflected in motor cortex local field potentials (LFPs) can account for spiking variability during motor behavior. Neural activity was recorded via microelectrode arrays implanted in ventral and dorsal premotor and primary motor cortices of non-human primates performing naturalistic 3-D reaching and grasping actions. Point process models were used to quantify how well LFP features accounted for spiking variability not explained by the measured 3-D reach and grasp kinematics. LFP features included the instantaneous magnitude, phase and analytic-signal components of narrow band-pass filtered ($\delta,\theta,\alpha,\beta$) LFPs, and analytic signal and amplitude envelope features in higher-frequency bands. Multiband LFP features predicted single-neuron spiking (1ms resolution) with substantial accuracy as assessed via ROC analysis. Notably, however, models including both LFP and kinematics features displayed marginal improvement over kinematics-only models. Furthermore, the small predictive information added by LFP features to kinematic models was redundant to information available in fast-timescale (<100 ms) spiking history. Overall, information in multiband LFP features, although predictive of single-neuron spiking during movement execution, was redundant to information available in movement parameters and spiking history. Our findings suggest that, during movement execution, collective dynamics reflected in motor cortex LFPs primarily relate to sensorimotor processes directly controlling movement output, adding little explanatory power to variability not accounted by movement parameters.

Keywords: neural dynamics, neural point processes, generalized linear models, local field potentials, neural variability

Introduction

The variability of neuronal responses at the level of single-neuron spiking is a fundamental problem in neuroscience (Shadlen and Newsome, 1998; Churchland, 2010; Churchland and Abbott, 2012). Neuronal responses in neocortex to repeated stimuli presentation or behavioral tasks show substantial variability. Determining the sources of this variability is particularly important

for understanding encoding of stimuli and behavioral parameters in neuronal ensembles. The issue is also critical for the development of reliable brain machine interfaces for the restoration of movement, communication, and sensory function in people with sensorimotor impairments (e.g., Hochberg et al., 2012). Beyond intrinsic stochasticity due to, for example, thermal noise and synaptic release failure (Faisal et al., 2008), variability in cortical neural responses has been proposed to arise from fluctuations in spontaneous, ongoing neural dynamics (Arieli et al., 1996; Wörgötter et al., 1998; Truccolo et al., 2002; Carandini, 2004). Although often neglected, spontaneous and ongoing neural dynamics are likely to affect how neurons respond to sensory inputs or even how they modulate their activity during behavior. In this way, spontaneous neural dynamics can provide a background of contextual effects which otherwise may appear as spiking variability (Fiser et al., 2004; Hermes et al., 2012; Goris et al., 2014) due to noise.

Local field potential (LFP) oscillations in different frequency bands result, to a large extent, from ongoing collective dynamics, i.e., modes of coordinated or coherent activity in neuronal populations (Nunez and Srinivasan, 2006; Buzsáki et al., 2012). Previous studies have investigated how features in multiband LFP oscillations relate to sensory stimuli and behavior and how decoding based on LFP features compares to decoding based on spiking activity (Belitski et al., 2008; Bansal et al., 2012). Additionally, some studies have examined how well LFP features predict spiking activity (e.g., Haslinger et al., 2006; Montemurro et al., 2008; Rasch et al., 2008; Kayser et al., 2009; Kelly et al., 2010). However, most of these studies have focused on early sensory cortices and most have been conducted during anesthesia, a neural state distinct from alert and active behavior. More importantly, none of the above studies have addressed how well ongoing collective dynamics reflected in LFPs account for single-neuron spiking variability that is not explained by behavioral parameters (e.g., movement kinematics). For example, the result in Bansal et al. (2012) showing LFP and spiking activity are redundant with respect to decoding kinematics does not address the issue of excess variability in single-neuron spiking.

Here, we examined and quantified how well features in multiband LFP oscillations account for single-neuron variability not explained by behavioral parameters in a naturalistic 3-D reach and grasp task performed by rhesus macaques (*Macaca mulatta*). The behavioral task elicited diverse reaching and grasping kinematics, and included reaching to grasp different objects with different styles of grip.

We examine various LFP features, such as the amplitude envelope and phase, of several LFP bands. The frequency bands included low (<2 Hz; delta) frequency components that are common in this task, including primarily motor related responses associated with these continuously performed sequences of reach and grasp actions (Bansal et al., 2011). In addition, these low frequency signals tend to be highly correlated to the neuronal population spike count (e.g., Bansal et al., 2012, **Figure 1**). Other bands included the theta (2–7 Hz), alpha (7–15 Hz), beta (15–30 Hz), gamma (30–60 Hz), high gamma (60–100 Hz), and higher frequency bands (100–200 Hz and 200–400 Hz). Although beta band oscillations dominate motor cortex LFP activity during movement preparation (Jasper and Penfield, 1949; Murthy and

Fetz, 1992; Sanes and Donoghue, 1993; Brovelli et al., 2004; Rubino et al., 2006), they are also characteristic during execution of isometric force tasks and other steady state conditions (Baker et al., 2001) and, less frequently, during the execution of reach and grasp actions (Reimer and Hatsopoulos, 2010). Neocortical LFP activity in higher (>100 Hz) frequency bands is generally considered to reflect fluctuations in multiunit activity resulting from coordinated activity in neuronal populations (Zhuang et al., 2010; Bansal et al., 2012; Buzsáki et al., 2012; Scheffer-Teixeira et al., 2013), and thus is also a reflection of ongoing collective dynamics.

We quantified the amount of variability accounted for by LFP features by fitting point process models (Truccolo et al., 2005, 2010) in which the conditional intensity function ("instantaneous" conditional spiking rate) was modeled as a function of covariates, including the ongoing LFP features mentioned above. To assess the amount of spiking variability explained by various models, we compare the relative predictive power of LFP features, reach and grasp kinematics, intrinsic spiking history, and combinations of these covariates, using receiver operating characteristic (ROC) curve analysis.

Methods

Behavioral Task
Data from three male rhesus macaques were examined in this study, C, R, and S. Data from C have been used previously in Vargas-Irwin et al. (2010). All experimental procedures were conducted as approved by the local Institutional Animal Care and Use Committee (IACUC). This study employed a task termed Free-Reach and Grasp (FRG). In FRG, an experimenter pseudo-randomly swings within the monkey's reach one of various small (3–5 cm) objects of differing shapes attached to a string, and the monkey is rewarded for grasping the object. The FRG task was designed to elicit naturalistic and continuous three-dimensional reach and grasp behaviors that require online motor control (Vargas-Irwin et al., 2010). Reach and grasp actions were organized in blocks within each experimental session. We analyze data from the entire FRG block, and so our data contains a diversity of behavioral conditions, including visually guided reaching (including online corrections as the objects' movement is unpredictable), the grasping and holding of the object, and the period of juice reward (see Vargas-Irwin et al., 2010).

Kinematic Feature Extraction
Arm and hand kinematics were recorded at 240 frames per second using a Vicon motion capture system (Vicon Motion Systems; Oxford Metrics Group) as detailed in Vargas-Irwin et al. (2010). This system employs infrared-reflective markers to track arm and hand positions in real time, and is capable of inferring missing data from briefly occluded markers. For our analysis, we focus on the 3-D kinematics measured at the wrist, as well as the distance between the thumb and forefinger (grasp aperture), as indicators of the kinematics related to reaching (i.e., hand position in space) and grasping. We analyzed kinematics features similar to those used in Hatsopoulos et al. (2007). These are normalized velocity trajectories of both the wrist endpoint and grip aperture in time, combined with zero lag position

FIGURE 1 | Behavioral and neural signals during the free-reach and grasp task. Kinematic, spiking, and LFP data from a single trial of the free reach to grasp task (monkey S, area M1). **(A)** Velocity of 3D wrist endpoint ("X," "Y," "Z") and distance between the thumb and forefinger ("Aperture"). **(B)** Broadband LFP, low-pass filtered at 500 Hz during two reach and grasp movements. **(C)** Instantaneous amplitude envelope (bold trace), via Hilbert transform, from the bandpass filtered LFP (gray trace). This particular example is filtered in the 7–15 Hz range. **(D)** Corresponding instantaneous phase extracted from the same bandpass filtered signal as in **(C)**. **(E)** Spiking population raster during this trial: spikes from 47 units are plotted along the y axis. Neurons show different task-related modulations.

and speed. For comparison, we also analyze position trajectories of these markers over time. The velocities of motion capture markers were estimated using a Savitzky-Golay filter generated by fitting a 5th order polynomial to a discrete differentiation operator at the sampling rate of the kinematics. The polynomial extended 25 samples (10 ms) to either side of the current time-point. The polynomial order was selected such that frequencies higher than 20 Hz were attenuated, so that the resultant velocity trajectories were sufficiently smooth to down-sample at 40 samples-per-second. Velocity trajectories were sampled from the smoothed velocity every 25 ms, starting 100 ms before the current time-point and extending 300 ms into the future. Similarly, a smoothed position estimate was extracted using Savitzky-Golay filters based on a 5th order polynomial fit to a discrete impulse. For the normalized velocity feature set, we followed the steps in Hatsopoulos et al. (2007), i.e., velocities were normalized by the L2 norm of the velocity trajectory. 3-D wrist position and

grip aperture were normalized separately. The average speed and position over the trajectories were added as additional features. A separate feature set based on position trajectories was used for comparison. Position trajectories were taken from the position estimates, sampled at 40 samples-per-second, covering the same period (100 ms before to 300 ms after) as the normalized velocity trajectories.

Neural Recordings

Recordings were made using microelectrode arrays (Blackrock Microsystems), as previously described in Vargas-Irwin et al. (2010). Electrodes were 1.5 mm long, likely targeting layer 5 of motor cortex. Monkey G was implanted with 4×4 mm 96-microelectrode arrays in both M1 and PMv. Two monkeys (R,S) were implanted with one 96-microelectrode array in PMv, and a 3.2×2.4 mm 48-microelectrode array in each of M1 and PMd areas. Electric potential signals were recorded broadband

(analog band-pass filtered to 0.3 Hz – 7.5 kHz; digitized at 16bit and sampled at 30 kilosamples per second) using a Cerebus Data Acquisition System (Blackrock Microsystems). For spike detection, recorded signals were digitally filtered with a 250 Hz fourth-order high-pass Butterworth filter. For each electrode, candidate spikes (action potentials) were identified online via threshold crossing detection in the amplitude of the high-pass filtered signal (Cerebus Data Acquisition System). Preliminary spike sorting was performed by a custom automated spike sorter (Vargas-Irwin and Donoghue, 2007), and verified using the commercial Plexon Offline Sorter (Plexon Inc.). Candidate units to be included in the analysis had a minimum signal-to-noise ratio (SNR) of 3.0 (defined as in Vargas-Irwin and Donoghue, 2007). In addition, we required that (a) the inter-spike-interval (ISI) histogram display a clear refractory period to exclude potential multi-unit clusters; (b) that there be at least 500 inter-spike-interval events smaller than 100 ms within the task data, to provide accurate estimates of spike-history filters; and (c) that units be clearly separated into different clusters in the PCA feature space. Electrodes exhibiting cross-talk or excessive noise were excluded from analysis. For monkey C, LFPs were extracted during recording sessions from the broadband signal using a causal 500 Hz fourth-order low-pass Butterworth filter, and stored at two kilosamples per second. LFP data from monkeys R and S were filtered offline to match this processing. Channels displaying cross-talk or excessive noise were excluded from the analysis.

LFP Feature Extraction

For analysis of band-limited LFP oscillations, LFP signals were filtered using causal (forward only) fourth-order Butterworth low-pass and band-pass filters, with cutoff frequencies 0.3–2 Hz (δ), 2–7 Hz (θ), 7–15 Hz (α), 15–30 Hz (β), 30–60 Hz (γ), 60–100 Hz (high γ), 100–200 Hz (MUA1), 200–400 Hz (MUA2). The 0.3–2 Hz low-pass signal captures slow motor-evoked potentials (MEPs). The two highest-frequency bands are likely to reflect a substantial contribution from multi-unit activity (MUA) to neocortical LFPs, as well as other possible high-frequency source signals. For the narrow delta, theta, alpha, and beta frequency bands, we considered four features: the instantaneous phase and amplitude of the analytic signal, and the real and imaginary component of the analytic signal. The LFP analytic signal was computed from the band-pass filtered LFP using the Hilbert transform. LFP instantaneous phase and amplitude were computed as the complex argument and modulus of the analytic signal, respectively. (The real component of the analytical signal corresponds to the band-pass LFP itself.) For the broader, higher frequency gamma and multi-unit bands, we use only the analytic signal and the amplitude envelope. Feature extraction was performed on the LFP sampled at two kilosamples per second, and decimated to one kilosample per second for neural point process modeling.

Spike Contamination

In this analysis, we predict single-unit spiking from features of the LFP recorded on the same electrode as the isolated unit. Because of this, when predicting neuronal spiking from

LFP features, it is important to prevent action potentials from contaminating the filtered LFP. We elected not to use existing spike removal procedures (e.g., Zanos et al., 2012) because the broadband LFP data were unavailable for monkey C, and because spike-removal methods make implicit assumptions about which features of the LFP relate to the spike waveform as opposed to collective dynamics locked to spiking. Instead, we employed causal filtering to extract LFP features. Although spike contamination can occur as low as 10 Hz (Waldert et al., 2013), causal filtering restricts this contamination to times following a spike, avoiding the confound of predicting spikes from themselves (i.e., via their contamination of the LFP features). Although the discrete Hilbert transform used to compute phase and amplitude features is non-causal, the effective filters created by the composition of the Hilbert transform with the causal band-pass filters remain predominantly causal. As a further precaution, and to guard against imprecision in localizing spike times, we added 1 ms delay to LFP features. Under this approach, the noncausal contribution to the imaginary component of the analytic signal was negligible: less than 0.14% of the impulse response (measured as the percentage of the area under the absolute impulse response) was non-causal. Since causal filters can add amplitude and phase distortions, we addressed this concern by comparing the predictive power of causally filtered LFP and that of zero-phase (non-causal) filtered LFP, which contains no delay. We determined that the choice of causal verses zero-phase filtering did not alter the conclusions of this paper for frequencies below 30 Hz. Zero-phase filtering for higher frequencies resulted in higher predictive power in some cases, which was likely the result of action potential contamination as supported by the recent studies mentioned above.

Intrinsic Spiking History

To assess the extent to which a neuron's own spiking history explains spiking variability, and to compare its predictive power to that of kinematic and LFP features, we included features of spiking history in our modeling (Truccolo et al., 2005). In addition to intrinsic biophysical processes (refractory/recovery period, bursting, etc.), spiking history can potentially also reflect indirect neuronal network dynamics effects. For example, spiking history models are capable of capturing spiking rhythmicity that may arise as a result of oscillatory input. We used raised cosine bases in logarithmically scaled time, covering the past 100 ms of spiking activity, to estimate temporal filters (see Pillow et al., 2008 and Truccolo et al., 2010, for more details). The resulting temporal filters were convolved with the past spiking activity to capture history effects on the spiking probability at a given time. Ten basis functions were used. More recent spiking history, typically related to after-spike refractory and recovery periods, and bursting, was modeled with more localized (finer temporal resolution) basis functions. Longer-term history effects can capture intrinsic rhythmicity and also, implicitly, network dynamics.

Stochastic Neural Point Process Models

We used a generalized linear point process model (Truccolo et al., 2005) to explore the extent to which different covariates

explain spiking variability. The probability of a neuron spiking in a sufficiently small time interval, indexed by t, of duration Δ, can be written as

$$\Pr(Y_t = 1|\lambda_t) = \lambda_t \Delta + o(\Delta), \qquad (1)$$

where Y_t corresponds to the spiking activity at time t, $Y_t=1$ for a spike, 0 otherwise, and λ_t is the conditional intensity function ("instantaneous conditional spiking rate," in spikes per second) of the modeled neuron. The bin size Δ must be chosen small enough such that the probability of two spikes occurring within the same time bin is negligible. Here $\Delta = 1$ ms. We used a regularized maximum likelihood approach to model the logarithm of the conditional intensity function as a linear combination of model features:

$$\ln(\lambda_t) = \mu + A \cdot X_t, \qquad (2)$$

where X_t is the covariate vector at time t, A is a vector of model parameters, and μ is a parameter related to background activity level. X_t can refer to LFP features at time t, past and future kinematics, convolutions of intrinsic spiking history up to but not including time t with temporal filters, or combinations of these covariates. For example, for a given Hilbert-transform of an LFP band, $z(t)$, the feature vector

$$X_t = \big(|z_t|,\ \mathrm{Re}(z_t),\ \mathrm{Im}(z_t),\ \cos(\mathrm{Arg}(z)),\ \sin(\mathrm{Arg}(z))\big) \qquad (3)$$

corresponds to a model with cosine tuning to a preferred Hilbert phase θ_0, as well as amplitude envelope and analytic signal features, i.e.,

$$\ln(\lambda_t) = \mu + a_1|z_t| + a_2\,\mathrm{Re}(z_t) + a_3\,\mathrm{Im}(z_t) + a_4\cos(\theta_0 - \mathrm{Arg}(z)) \qquad (4)$$

Model Fitting

Model estimation is solved by finding parameters A and μ that maximize the L2-regularized log-likelihood of the observed spiking activity (Truccolo et al., 2005):

$$\underset{A,\mu}{\mathrm{argmax}}[\ln(\Pr(Y|X, A, \mu))] = \frac{1}{T}\sum_{t=1}^{T}[Y_t\ln(\lambda_t\Delta) - \lambda_t\Delta] - \alpha\|A\|^2, \qquad (5)$$

where α is a penalty or regularization parameter. The log-likelihood is normalized by the number of samples T so that the strength of regularization does not depend on the amount of data. The parameter μ is not penalized. All features are z-scored prior to model fitting to ensure that all features are zero mean and of comparable scale, which ensures that the L2 penalty is applied equally to all features and improves numerical accuracy.

We used a gradient descent approach for the minimization of the negative log-likelihood under L2 regularization. Models were fit under a two-tier cross-validation scheme. An outer level of 10-fold cross-validation ensures that results are not overfit. An inner level of cross-validation selects the regularization parameter α. Ten values of the regularization parameter α, base-10 logarithmically spaced between 1e-9 and 1e2 inclusive,

as well as $\alpha = 0$, were tested. On each of the 10 outer-level cross-validations, 90% of the data were taken as training data, and 10% were reserved for testing. The training data were split randomly into two equal groups. For each group, models were generated for each value of the regularization parameter α. The value of the regularization parameter that led to the best generalization (in terms of predictive power, see below) in this internal cross-validation step was selected for fitting a model on all of the training data. This model was then validated on the remaining 10% of the data that had been withheld for testing. This two-tier cross-validation procedure was repeated 10 times, such that all of the available data was used for model validation and assessing predictive power. To confirm that L2 regularization sufficiently prevented over-fitting when adding LFP features to the kinematics model, we shuffled LFP features in 100 ms blocks with respect to the spiking activity. We found that adding these non-informative features to the kinematics and kinematics-history combined models reduces the predictive power very little, by at most 0.03, and with the population mean decrease ranging from 0.001 and 0.006 across all sessions. This difference is too small to alter the conclusions of our study. In preliminary analysis, we also explored L1 regularization and also a mixed L1-L2 regularization (fitted via elastic net, Friedman et al., 2010). We found that L2 regularization outperformed these alternatives, both in terms of computational time and predictive power under generalization. Additionally, we found that L2 regularized GLM models outperformed simpler approaches, such as naive Bayes and spike-triggered average and covariance analysis.

Assessment of Predictive Power

Model performance was evaluated using the area under the ROC convex hull (AUC) measured on testing data, using the model (Equation 2) to compute the conditional spiking probability, $\Pr(Y_t=1|X_t, A, \mu) \approx \lambda_t(X_t, A, \mu)\Delta$, from observed covariates (Truccolo et al., 2010). ROC analysis assesses predictive power in the context of binary (in this case spike train) sequences (Fawcet, 2006). We report predictive power (PP) as $2 \times AUC - 1$, which ranges from 0 (no predictive power) to 1 (complete prediction of single-neuron spikes in 1 ms time bins). Note that this predictive power measure is based on both true and false positive rates, since is derived from the ROC analysis.

Results

We are interested in how well collective neural dynamics, as reflected by features in ongoing and evoked multi-band LFP signals, can explain motor cortex single-neuron spiking variability not accounted for by motor behavioral covariates such as reach and grasp kinematics. We first demonstrate that the examined LFP features can predict single-neuron spiking in motor cortex, then we assess the extent to which this predictive power compares and is redundant to information available in 3-D kinematics. We also assess the extent to which intrinsic spiking history, i.e., temporal dynamics or correlation in the modeled spiking activity itself, adds predictive power to kinematics, and evaluate whether LFP features remain predictive conditioned on kinematics and intrinsic spiking history.
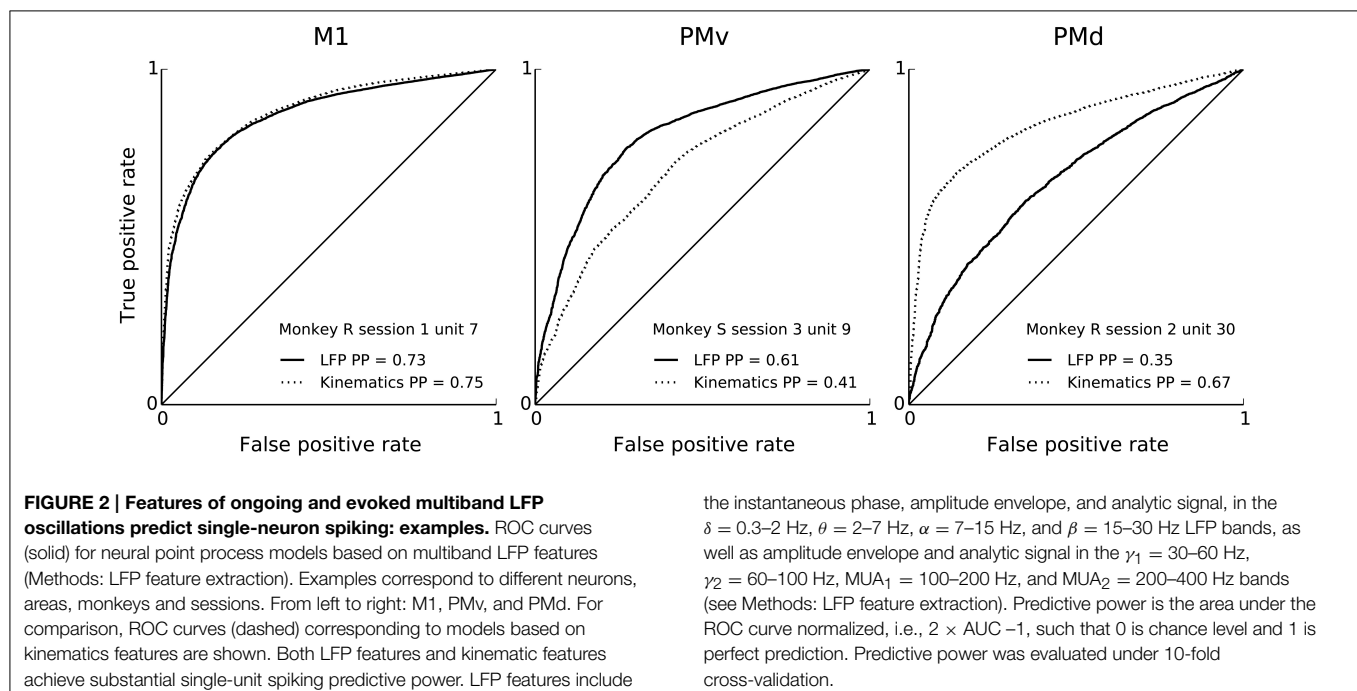
Datasets from seven experimental sessions were used in these analyses: two each from monkey C and R, and three from monkey S. Sessions from monkeys R and S were collected within a week of each-other. The two sessions from monkey C were collected 3 months apart. Between three and nine reach and grasp blocks, averaging 140 s long, were collected on each session. Each session included 15–42 successful free-reach-to-grasp trials or reaches in each block. This yielded 7–17 min of FRG task data, averaging 10 min of data per session. A detailed statistical description of the kinematics and examples of kinematic trajectories in these experimental blocks can be found in Vargas-Irwin et al. (2010) and Bansal et al. (2012), who reported some of the data from monkey C in this task. An example of kinematics trajectories with the corresponding neuronal ensemble spike raster is shown in **Figure 1**. For each array in each session, between 19 and 83 well-isolated units were identified for analysis (mean = 52, σ = 16). For a given monkey and area, some of the neurons are thought to be the same across sessions, for this reason we do not combine sessions when we perform statistical significance tests.
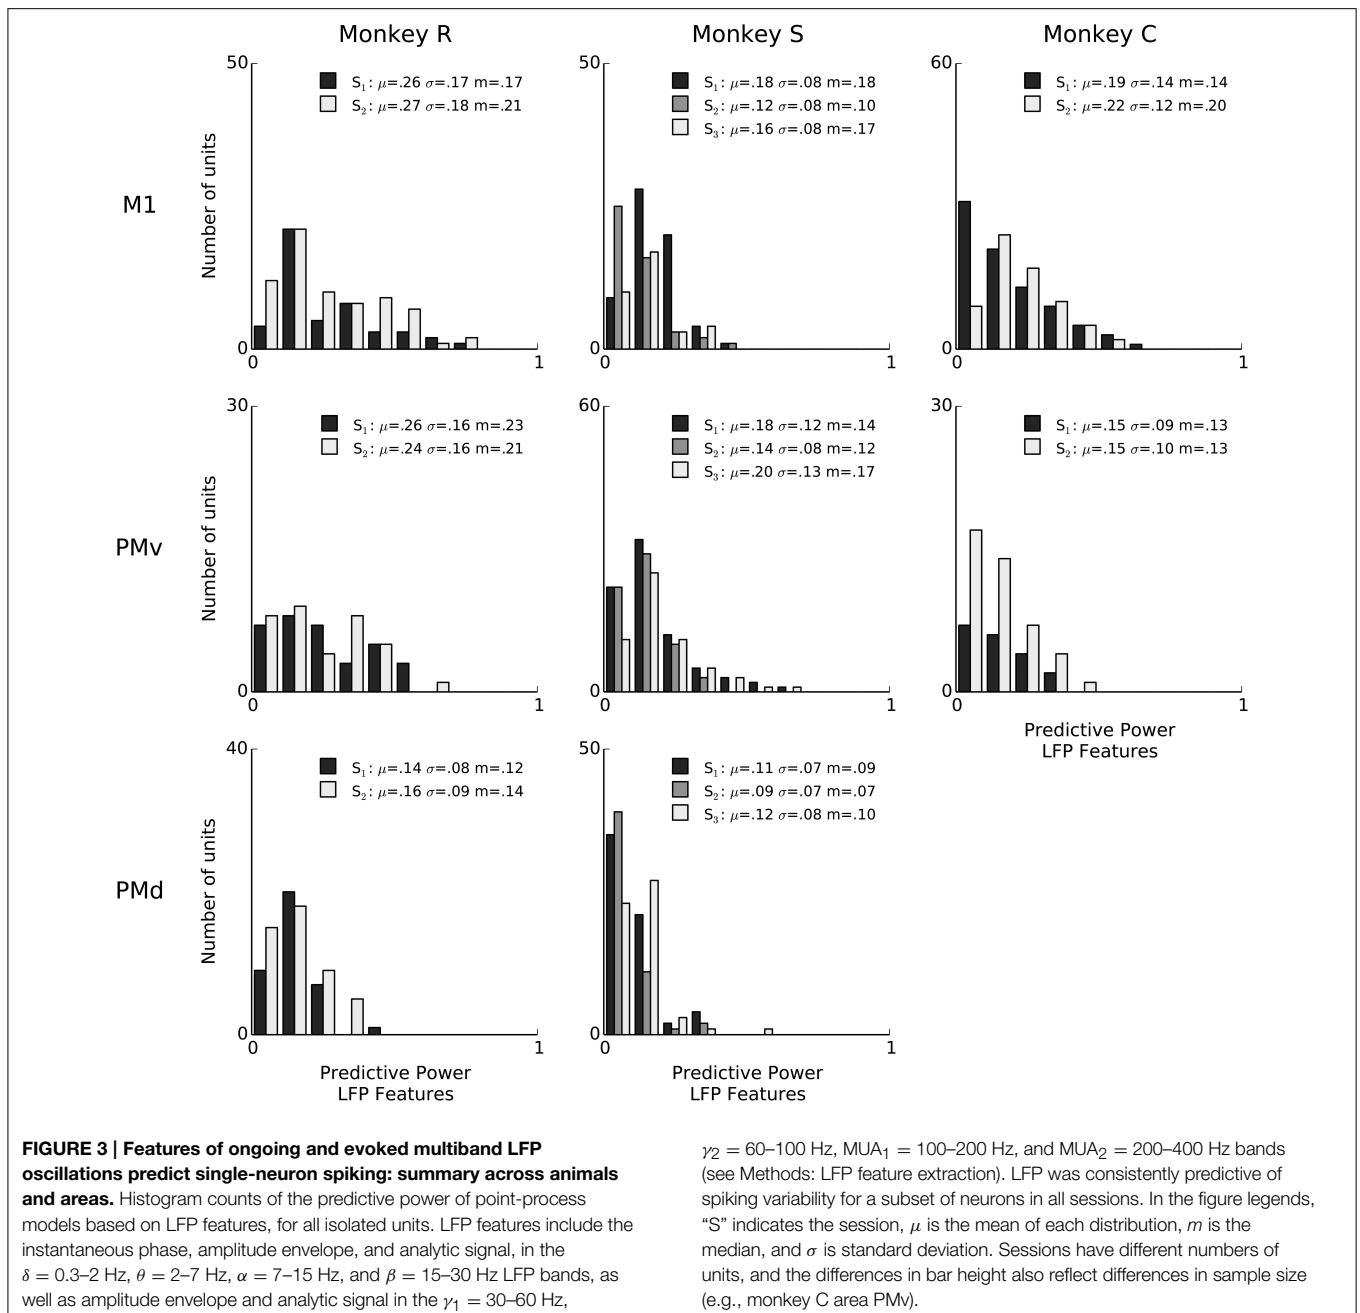
Features of LFP Oscillations Predict Single-Neuron Spiking with Substantial Power

First, we evaluated the ability for multiband LFP features to predict single-neuron spiking. We fit a regularized generalized linear model to predict single-unit spiking (1 ms time resolution) from multiband LFP features. Spiking probability at any given 1 ms time interval was modeled as a function (Equations 1–2, Methods) of features of ongoing LFP activity. LFP features included instantaneous phase and amplitude envelope as well as the analytic signal, extracted via a Hilbert transform, from four narrow LFP bands (Methods), δ (0.3–2 Hz, motor related potentials), θ (2–7 Hz), α (7–15 Hz), β (15–30 Hz), as well

as the amplitude envelop and analytic signal for four broader, higher frequency bands: γ_1 (30–60 Hz), γ_2 (60–100 Hz), and two multi-unit activity (MUA) bands MUA_1 (100–200 Hz), and MUA_2 (200–400 Hz). **Figure 1** illustrates the signal processing steps involved in the computation of the instantaneous phase and magnitude for a single frequency band.

We report the extent to which a model explains spiking variability in terms of "predictive power" (PP). Predictive power is the normalized area under the receiver operator characteristic (ROC) curve such that 0 corresponds to chance level and 1 to perfect prediction of spike times at 1ms resolution (see Methods: Model fitting; Methods: Assessment of predictive power). **Figure 2** shows PPs obtained from three example neurons from different monkeys and areas. For illustration and comparison, we also show the corresponding PPs based on a model including only kinematics features related to lagged 3-D velocity and position (similar to "pathlets" in Hatsopoulos et al., 2007) and grasp aperture (Methods). The examples show a case (**Figure 2**, left) in which LFP features and kinematics both explained a substantial fraction of spiking variability (PP = 0.73, 0.75, respectively) and two other examples in which LFP features did better and worse than kinematics, respectively. Overall, we found that LFP features were typically predictive, and that for some neurons LFP accounted for a substantial (i.e., PP > 0.5) fraction of variability. We performed a permutation test to assess chance level LFP predictive power by shuffling LFP features in 100 ms blocks relative to spiking, and found that the 95% chance predictive power ranged from 0.03 to 0.06 across sessions. Across sessions and areas, between 85% and 100% of units showed LFP predictive power higher than this chance level. As shown in **Figure 3**, high predictive power from LFP was consistent across all monkeys, motor cortical areas (PMv, PMd, and M1)

FIGURE 2 | Features of ongoing and evoked multiband LFP oscillations predict single-neuron spiking: examples. ROC curves (solid) for neural point process models based on multiband LFP features (Methods: LFP feature extraction). Examples correspond to different neurons, areas, monkeys and sessions. From left to right: M1, PMv, and PMd. For comparison, ROC curves (dashed) corresponding to models based on kinematics features are shown. Both LFP features and kinematic features achieve substantial single-unit spiking predictive power. LFP features include

the instantaneous phase, amplitude envelope, and analytic signal, in the δ = 0.3–2 Hz, θ = 2–7 Hz, α = 7–15 Hz, and β = 15–30 Hz LFP bands, as well as amplitude envelope and analytic signal in the γ_1 = 30–60 Hz, γ_2 = 60–100 Hz, MUA_1 = 100–200 Hz, and MUA_2 = 200–400 Hz bands (see Methods: LFP feature extraction). Predictive power is the area under the ROC curve normalized, i.e., 2 × AUC −1, such that 0 is chance level and 1 is perfect prediction. Predictive power was evaluated under 10-fold cross-validation.

FIGURE 3 | Features of ongoing and evoked multiband LFP oscillations predict single-neuron spiking: summary across animals and areas. Histogram counts of the predictive power of point-process models based on LFP features, for all isolated units. LFP features include the instantaneous phase, amplitude envelope, and analytic signal, in the $\delta = 0.3-2$ Hz, $\theta = 2-7$ Hz, $\alpha = 7-15$ Hz, and $\beta = 15-30$ Hz LFP bands, as well as amplitude envelope and analytic signal in the $\gamma_1 = 30-60$ Hz,

$\gamma_2 = 60-100$ Hz, $MUA_1 = 100-200$ Hz, and $MUA_2 = 200-400$ Hz bands (see Methods: LFP feature extraction). LFP was consistently predictive of spiking variability for a subset of neurons in all sessions. In the figure legends, "S" indicates the session, μ is the mean of each distribution, m is the median, and σ is standard deviation. Sessions have different numbers of units, and the differences in bar height also reflect differences in sample size (e.g., monkey C area PMv).
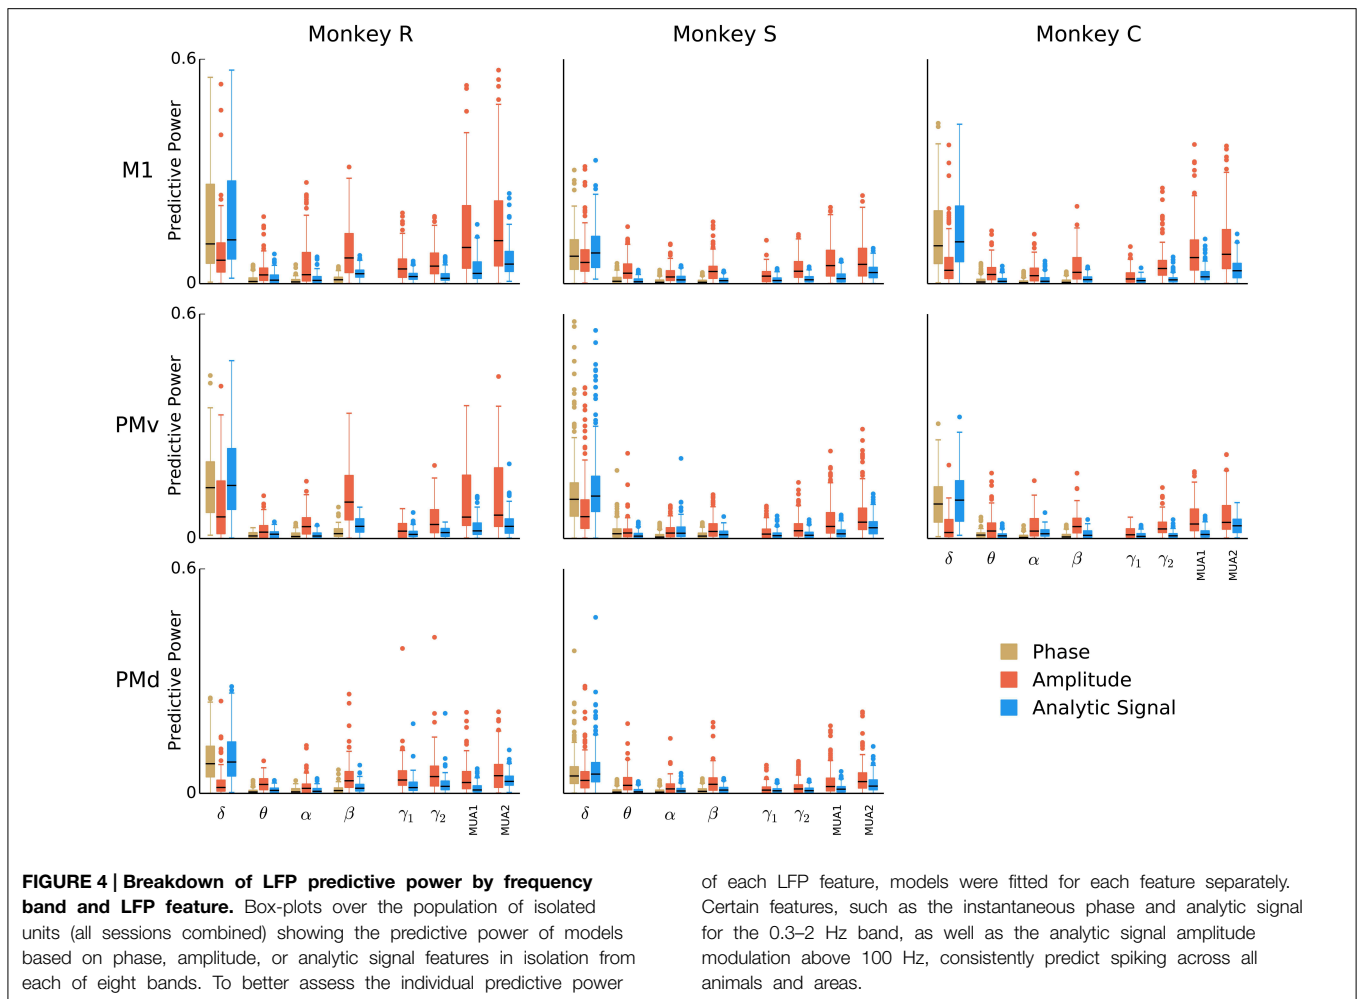
and sessions. This finding demonstrates that collective dynamics reflected in ongoing and evoked LFP oscillations can account for a substantial fraction of single-neuron spiking variability.

LFP Features Contributing to Prediction of Single-Neuron Spiking

We examined whether some of the multiband LFP features contribute more to prediction of single neuron spiking than others. Analysis based on estimated model coefficients is complicated because of the nonlinear (multiplicative) interactions between different features (amplitude, phase, or analytical signal) in different frequency bands. Instead, we

performed an analysis based on fitting a single model for each feature separately and assessing how well each separate model and feature predicted spiking. This allowed easy visualization of the predictive power of each individual LFP feature.

This analysis revealed some trends common to all animals and motor cortex areas, but also some variations (**Figure 4**). Consistently across animals and areas, low-frequency local field potentials (δ, 0.3–2 Hz) showed predictive power in the time domain signals and phases, but not amplitude envelopes. Additionally, the amplitude envelope in the multi-unit (100–200 and 200–400 Hz) bands was predictive, more-so than the signal. The low-frequency <2 Hz analytic signal was the most predictive

FIGURE 4 | Breakdown of LFP predictive power by frequency band and LFP feature. Box-plots over the population of isolated units (all sessions combined) showing the predictive power of models based on phase, amplitude, or analytic signal features in isolation from each of eight bands. To better assess the individual predictive power of each LFP feature, models were fitted for each feature separately. Certain features, such as the instantaneous phase and analytic signal for the 0.3–2 Hz band, as well as the analytic signal amplitude modulation above 100 Hz, consistently predict spiking across all animals and areas.

for 49% of units (486 units out of 991), and the instantaneous amplitude envelope and phase the most predictive for 13% and 8% of units (133 and 82 units out of 991), respectively. The amplitude envelope in the 200–400 Hz band was the most predictive for 14% of units (142 units out of 991). Features from intermediate 2–100 Hz bands generally performed poorly, with the exception of beta (15–30 Hz) amplitude, which although less predictive than the aforementioned features, was still amongst the top 4 most predictive features for 23% of units (227 units out of 991). The finding that LFP amplitude was predictive for the beta-frequency LFP was strongest in monkey R for areas M1 and PMv. To understand in more detail how the predictive power in beta amplitude varies across monkeys and areas, we examined the distribution of the model parameter weights for beta amplitude (in the case of the amplitude-only model, the parameter matrix A in Equation (2) is simply a single scalar parameter). Model weights for beta amplitude in monkey R areas M1 and PMv were more negative (mean $\pm$ 2SD for M1 and PMv were -0.18 ± 0.4 and -0.23 ± 0.4) than those from other monkeys and areas (-0.04 ± 0.16), indicating that a reduction in beta amplitude is typically associated with an increase in firing rate.

Predictive Power of Kinematics During Naturalistic Reach and Grasp Movements

We next quantified the predictive power of motor behavior, specifically kinematic features of the 3-D reach and grasp movements. We found that kinematics trajectories can also predict single-neuron spiking with substantial accuracy, at times achieving predictive power levels around 0.8 (e.g., **Figure 5**, monkey R area M1). However, similarly to LFP features, there was considerable diversity in the extent to which kinematics predicted spiking, with some units being predicted poorly. Similar results were obtained by using position trajectories (i.e., position at multiple time lags with respect to spiking; Methods). The 95% chance level predictive power for kinematics ranged from 0.03 to 0.08, as assessed by shuffling kinematics features in 100ms blocks relative to spiking. Across sessions and areas, between 49 and 100% of units showed LFP predictive power higher than this chance level. These effects were consistent across all animals, sessions, and motor areas (**Figure 5**), with mean predictive power ranging between 0.16 to 0.36 across sessions and areas. The fact that the task kinematics predict single-unit spiking variability confirms that we are recording from motor
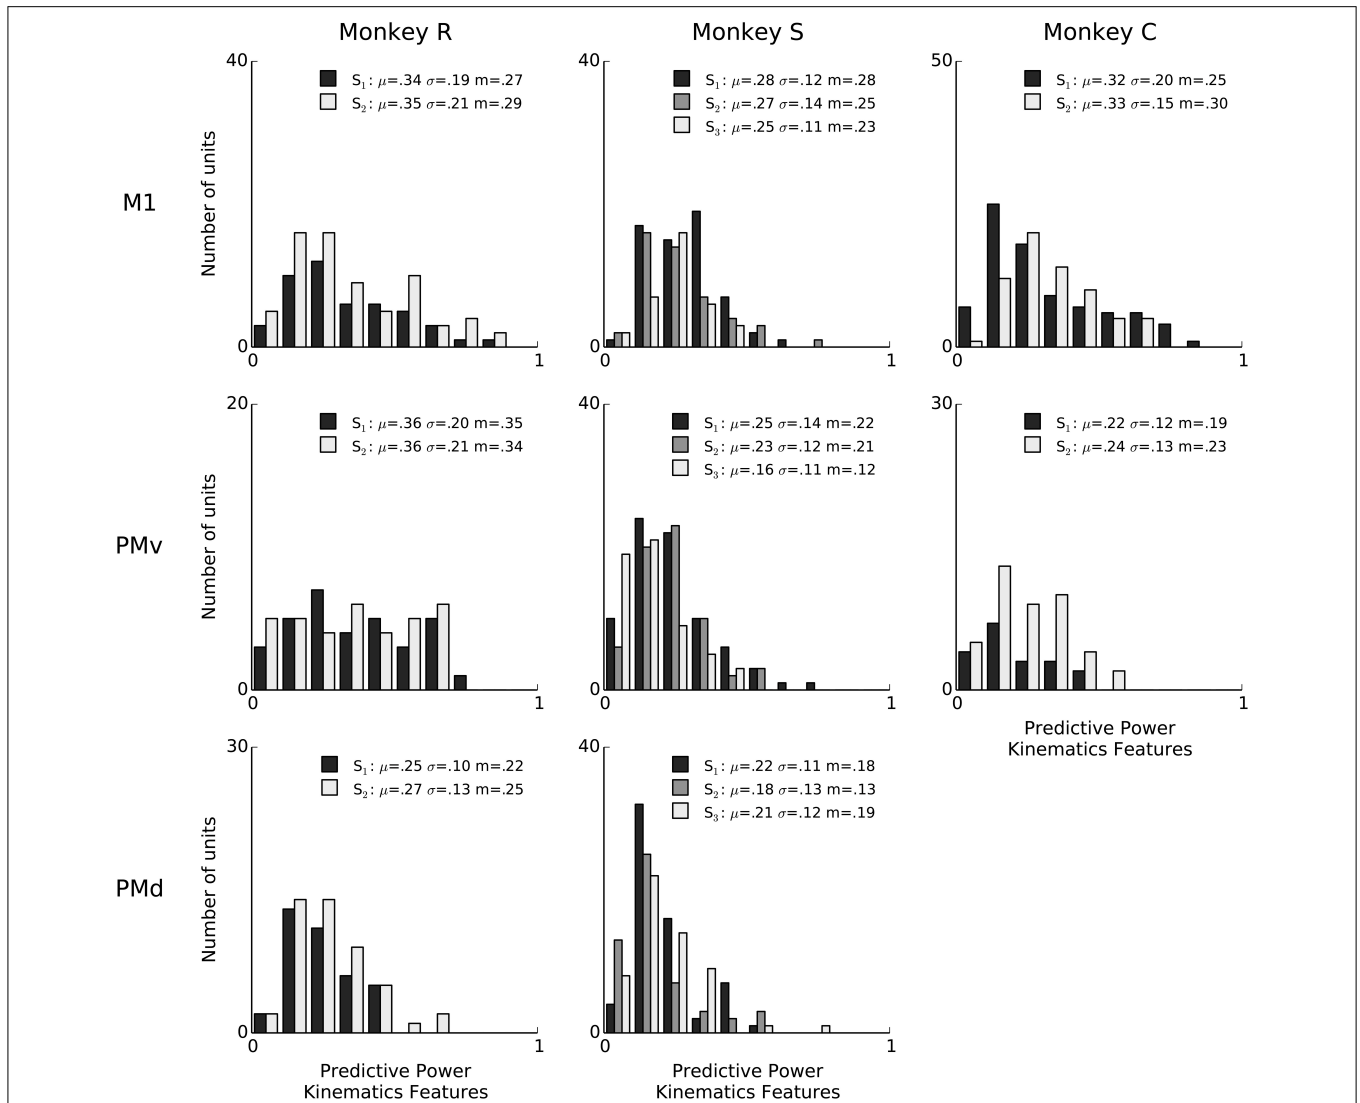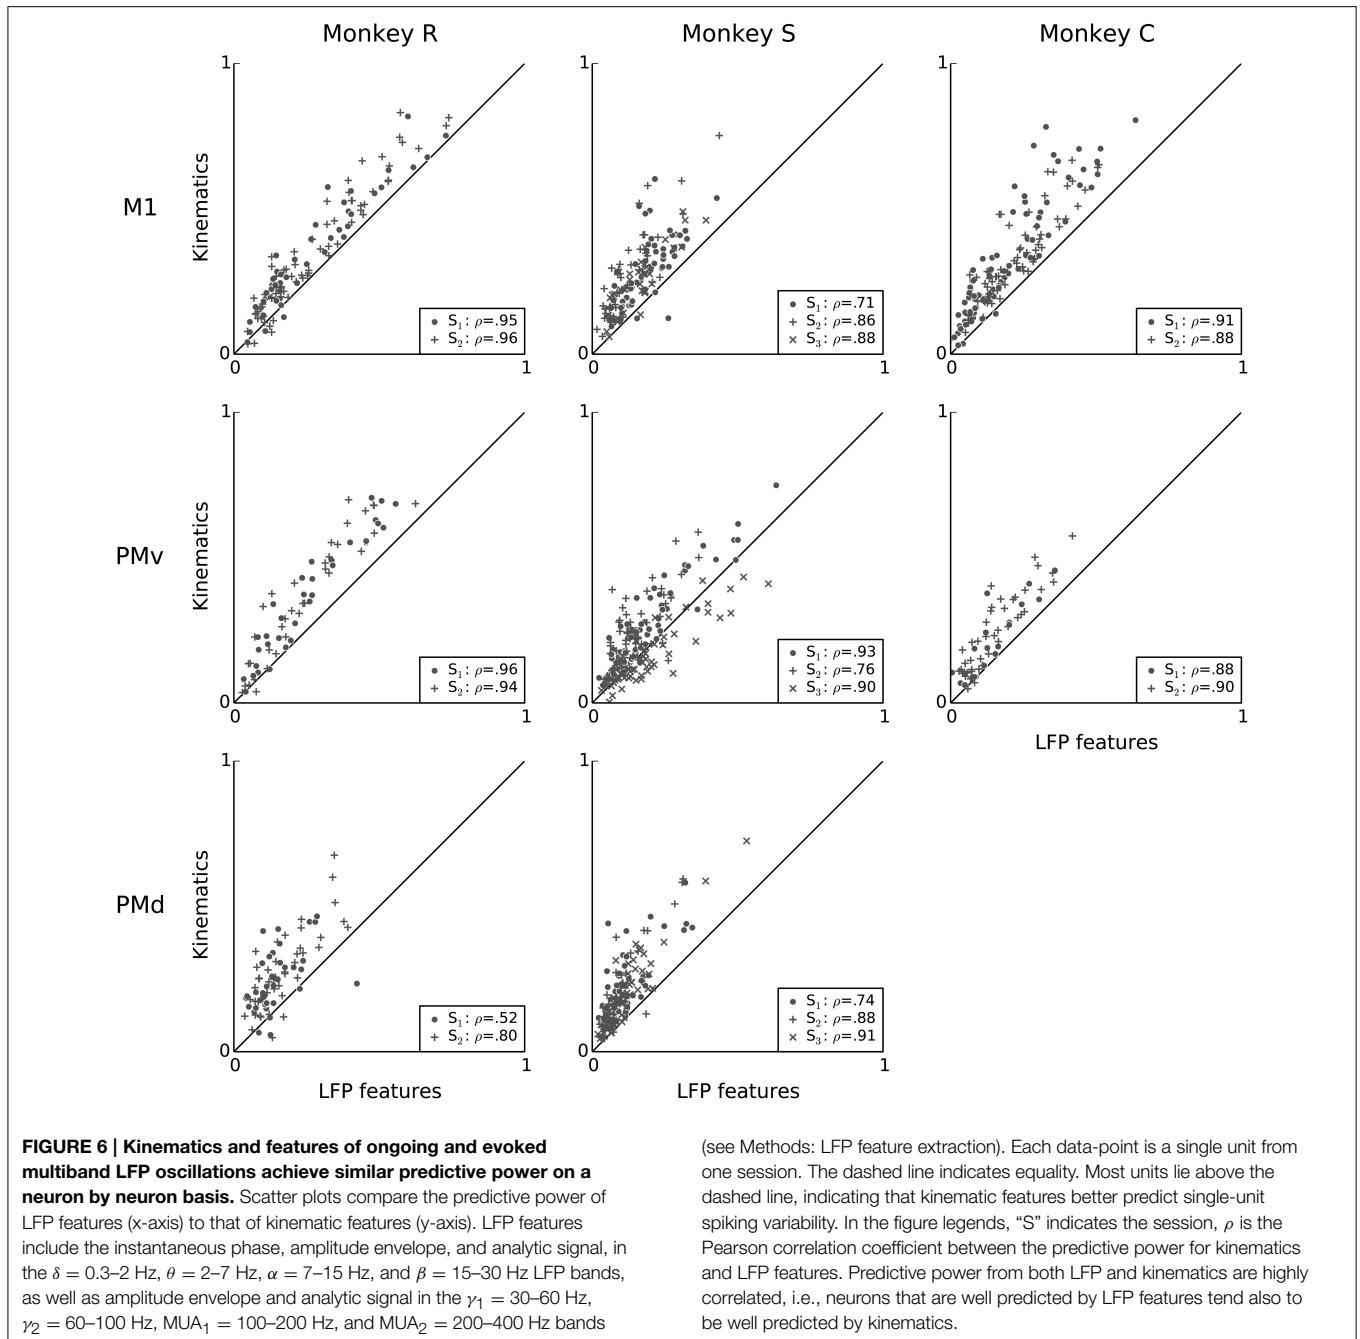
FIGURE 5 | Kinematic features predictive power for single-neuron spiking: summary across animals and areas. Histogram counts of the predictive power of point-process models based on kinematics features, for all isolated units. Kinematics features are normalized velocity trajectories of wrist endpoint and grip aperture, extending from 100 ms in the past to 300 ms in the future, sampled every 25 ms, as well as the average speed and zero-lag position for wrist endpoint and grip aperture (Methods: Kinematic feature extraction). In the figure legends, "S" indicates the session, μ is the mean of each distribution, m is the median, and σ is standard deviation. Sessions have different numbers of units, and the differences in bar height also reflect differences in sample size (e.g., monkey C area PMv).

cortex populations that exhibit task-modulation and tuning to motor output.

Figure 6 directly compares the predictive power of LFP and kinematics features. Overall, the predictive power of LFP features was typically less than that of kinematics during this free reach and grasp task: the difference between the predictive power of models based on kinematics and LFP features ranged from −0.20 to 0.45. With exception of monkey S area PMv, units for which LFP features explain more variability than kinematics are rare. The mean difference in predictive power within each session ranged from −0.04 to 0.14, and the median difference from −0.02

to 0.12, with all (session, area) pairs except monkey S area PMv session 3 displaying significantly better median predictive power for kinematics. (Wilcoxon signed-rank test with $p < 0.05$, corrected for multiple comparisons using Bonferroni correction for 19 (session, area) pairs.) Furthermore, across all monkeys and areas, predictive power of LFP was highly correlated with that of kinematics: the Pearson correlation coefficient between the predictive power of kinematics and LFP features ranged from 0.52 to 0.96, with a mean of 0.86 and a median of 0.88. This raises the possibility that LFP and kinematics shared some common effect, which we address below.

FIGURE 6 | Kinematics and features of ongoing and evoked multiband LFP oscillations achieve similar predictive power on a neuron by neuron basis. Scatter plots compare the predictive power of LFP features (x-axis) to that of kinematic features (y-axis). LFP features include the instantaneous phase, amplitude envelope, and analytic signal, in the $\delta = 0.3\text{--}2$ Hz, $\theta = 2\text{--}7$ Hz, $\alpha = 7\text{--}15$ Hz, and $\beta = 15\text{--}30$ Hz LFP bands, as well as amplitude envelope and analytic signal in the $\gamma_1 = 30\text{--}60$ Hz, $\gamma_2 = 60\text{--}100$ Hz, $MUA_1 = 100\text{--}200$ Hz, and $MUA_2 = 200\text{--}400$ Hz bands

(see Methods: LFP feature extraction). Each data-point is a single unit from one session. The dashed line indicates equality. Most units lie above the dashed line, indicating that kinematic features better predict single-unit spiking variability. In the figure legends, "S" indicates the session, ρ is the Pearson correlation coefficient between the predictive power for kinematics and LFP features. Predictive power from both LFP and kinematics are highly correlated, i.e., neurons that are well predicted by LFP features tend also to be well predicted by kinematics.
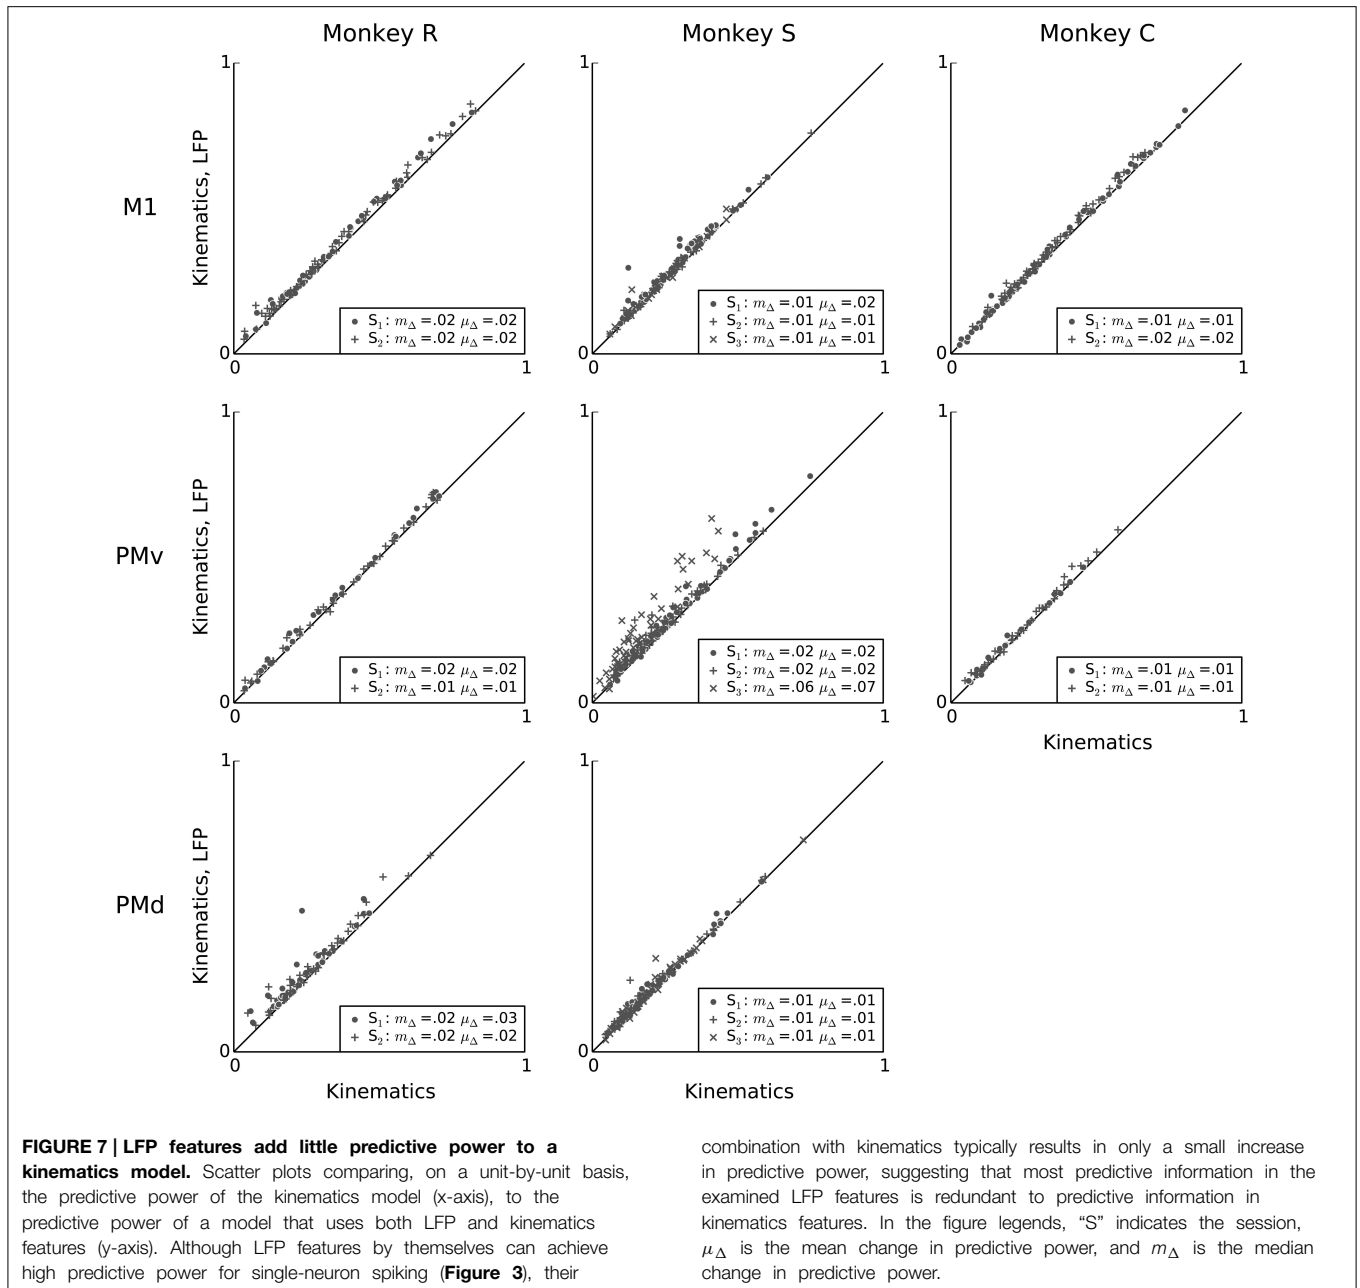
LFP Features are Mostly Redundant to Kinematics When Explaining Single-neuron Spiking Variability

In order to determine whether LFP features can account for single-neuron spiking variability not explained by kinematics, we asked whether the predictive information carried in the examined LFP features about single-neuron spiking variability was redundant to the predictive information carried in kinematic features. To assess redundancy we compare the predictive power of point process models based only on kinematics features to models that included both kinematics and LFP features. We

used L2 regularization to control for overfitting to the training data due to the larger number of parameters in the models that combined both kinematics and LFP features (Methods). We found that forgoing L2 regularization led to overfitting, in which the larger number of parameters in the combined LFP-kinematics model generalized less well to the evaluation data. Tests using shuffled LFP features confirmed that the L2 regularization approach adequately prevented overfitting (Methods).

Figure 7 compares, on a unit-by-unit basis, the relative predictive powers of kinematics and LFP features. The analysis reveals that, with a few exceptions (e.g., some neurons in PMv
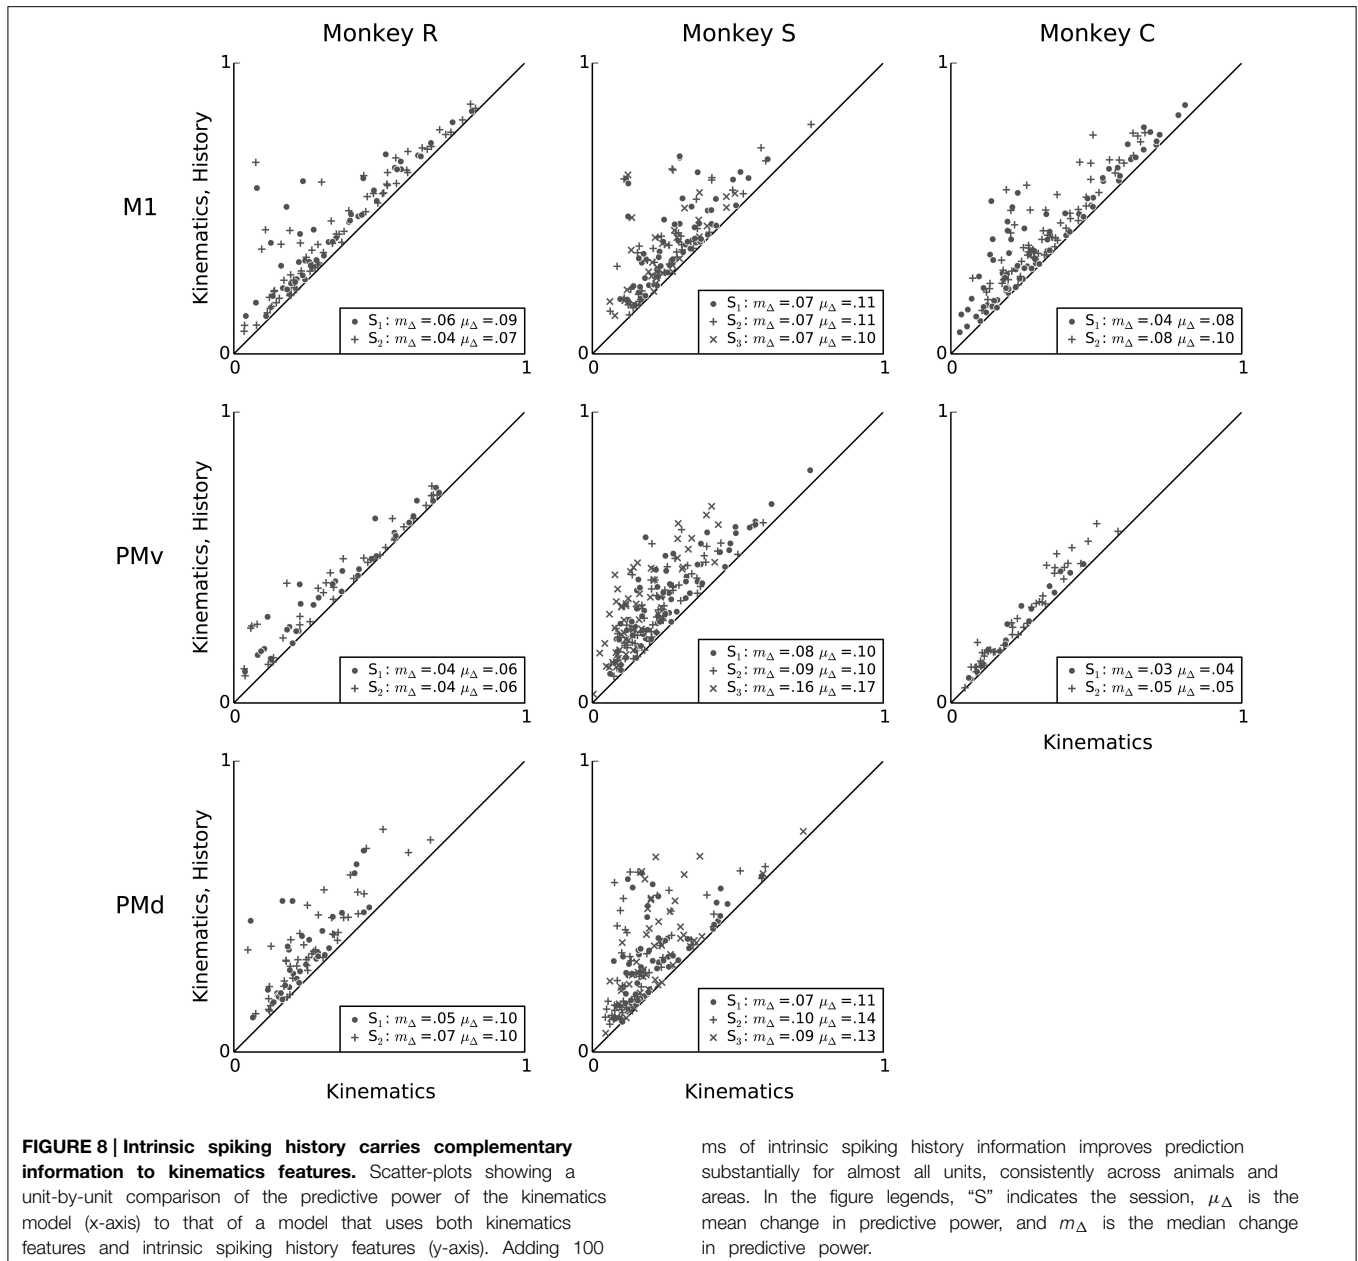
FIGURE 7 | LFP features add little predictive power to a kinematics model. Scatter plots comparing, on a unit-by-unit basis, the predictive power of the kinematics model (x-axis), to the predictive power of a model that uses both LFP and kinematics features (y-axis). Although LFP features by themselves can achieve high predictive power for single-neuron spiking (**Figure 3**), their combination with kinematics typically results in only a small increase in predictive power, suggesting that most predictive information in the examined LFP features is redundant to predictive information in kinematics features. In the figure legends, "S" indicates the session, μ_Δ is the mean change in predictive power, and m_Δ is the median change in predictive power.

in monkey S), LFP features added little predictive power to kinematics. This finding suggests that although LFP features were able to account for a substantial fraction of spiking variability, this information was highly redundant to information available in the examined kinematics features. This finding confirms the conjecture raised earlier, based on the high correlation between LFP and kinematics predictive power (**Figure 6**), that the information available in these two signals was redundant in terms of prediction of single-neuron spiking activity. Nevertheless, for each session, after adding LFP features, the mean change in predictive power was positive, ranging from 0.008 to 0.07, and the median change in predictive power ranged from 0.006 to 0.06. This median increase was statistically significant for all sessions.

(Wilcoxon signed-rank test, $p < 0.05$ corrected for multiple comparisons using Bonferroni correction for 19 (session, area) pairs).

Intrinsic Spiking History Adds Substantial Power to Kinematics in the Prediction of Single-neuron Spiking Variability
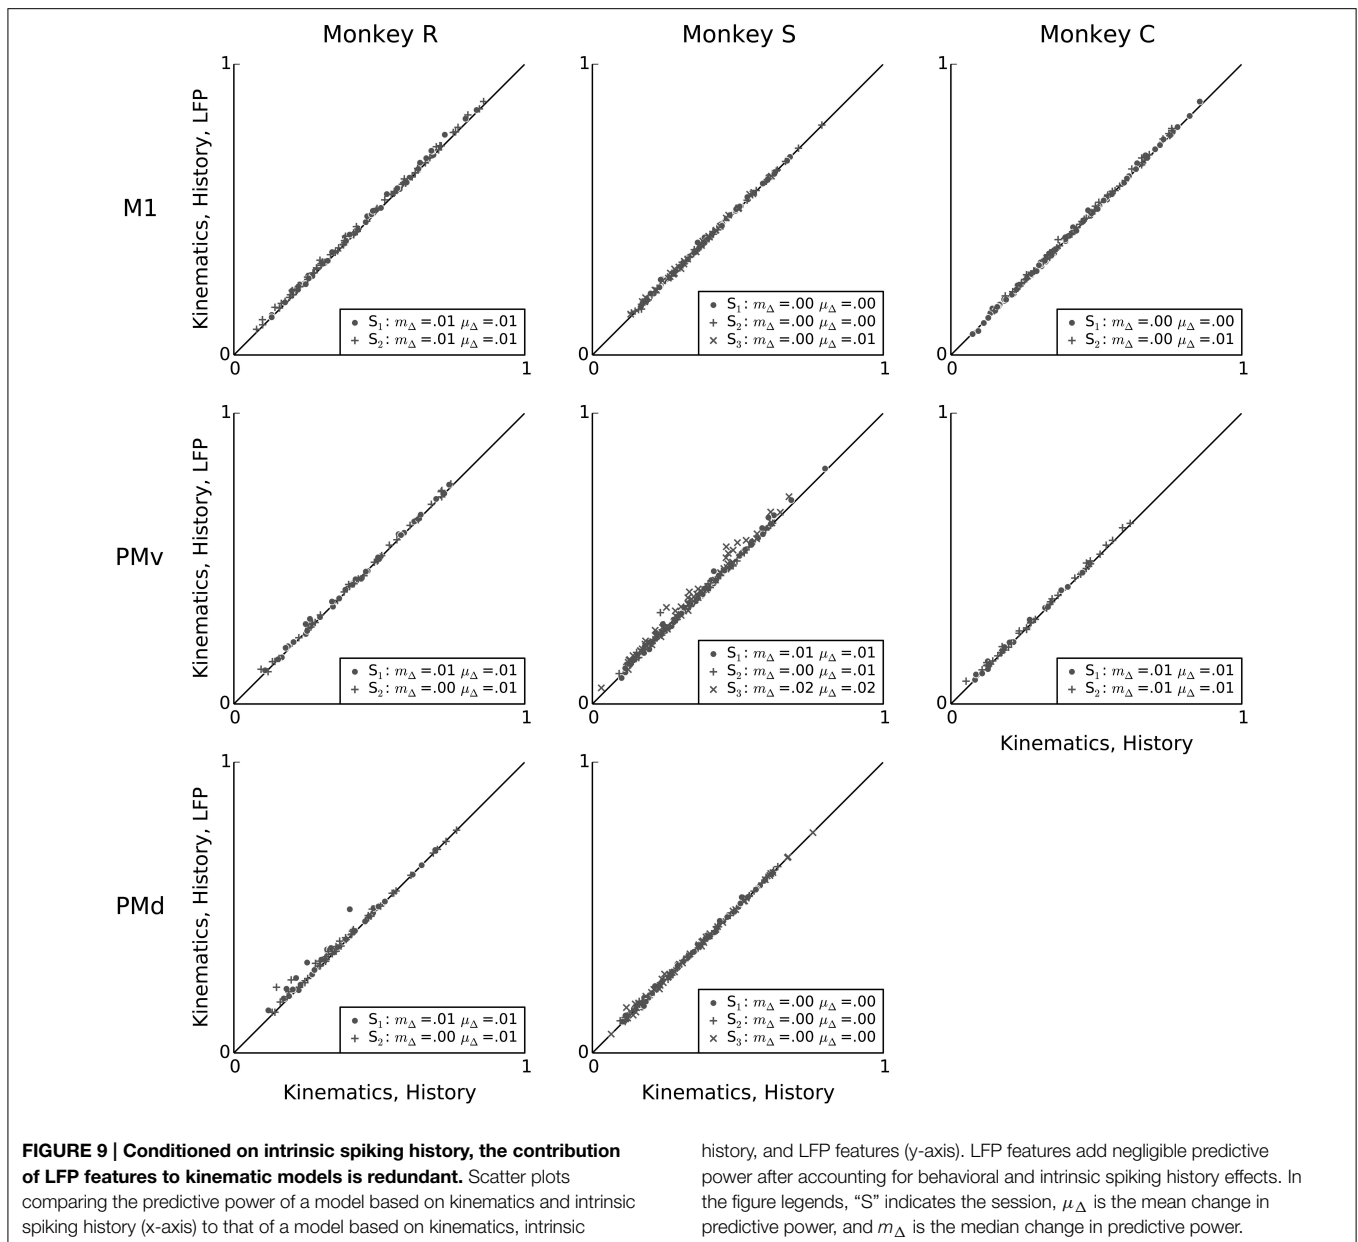
The above findings suggest that the examined features of LFP collective dynamics during movement reflect primarily sensorimotor processes related to motor representations and computations associated with the measured reach and grasp kinematics. Nevertheless, these LFP features accounted for a small, but statistically significant, fraction of neural variability not

FIGURE 8 | Intrinsic spiking history carries complementary information to kinematics features. Scatter-plots showing a unit-by-unit comparison of the predictive power of the kinematics model (x-axis) to that of a model that uses both kinematics features and intrinsic spiking history features (y-axis). Adding 100 ms of intrinsic spiking history information improves prediction substantially for almost all units, consistently across animals and areas. In the figure legends, "S" indicates the session, μ_Δ is the mean change in predictive power, and m_Δ is the median change in predictive power.

accounted by the examined kinematics features. To investigate the potential sources of this additional predictive power, we take a detour in this section and consider first the predictive power of a neuron's own spiking history. Here we focused on the preceding 100 ms spiking history, which can capture fast intrinsic biophysical processes such as refractory and recovery periods after an action potential, and also bursting dynamics, which are common in certain types of motor cortex neurons (Chen and Fetz, 2005). In addition, temporal autocorrelations within single-neuron's spiking activity can be induced, for example, by both intrinsic rhythmicity and rhythmicity due to ongoing neuronal network dynamics affecting spiking. We used temporal filters to capture the effects of intrinsic spiking history. Temporal filters were estimated with semi-parametric models using raised cosine functions (Pillow et al., 2008; Truccolo et al., 2010; Methods). Ten logarithmically-spaced raised cosine functions on the past 100 ms were used.

When information about a single neuron's own spiking history is added to the model, there is a significant improvement in predictive power compared to a model including only kinematics features (**Figure 8**). Within each session and area, the mean increase in predictive power ranged from 0.04 to 0.17. The median increase in predictive power ranged from 0.03 to 0.16, and was statistically significant for all sessions and motor areas (Wilcoxon signed-rank test, $p < 0.05$ with Bonferroni correction for 19 (session, array) multiple tests).

FIGURE 9 | Conditioned on intrinsic spiking history, the contribution of LFP features to kinematic models is redundant. Scatter plots comparing the predictive power of a model based on kinematics and intrinsic spiking history (x-axis) to that of a model based on kinematics, intrinsic history, and LFP features (y-axis). LFP features add negligible predictive power after accounting for behavioral and intrinsic spiking history effects. In the figure legends, "S" indicates the session, μ_Δ is the mean change in predictive power, and m_Δ is the median change in predictive power.

This result demonstrates that fast-timescale spiking history can explain variability in single-neuron spiking that is not redundant to variability examined by the kinematic features in this motor task.

Conditioned on Spiking History, Contribution of LFP Features to Kinematic Models is Further Reduced

Having demonstrated that intrinsic spiking history adds predictive power to kinematic models, we finally assessed whether LFP features can account for variability in single-neuron spiking not accounted for by kinematics and intrinsic history features. **Figure 9** shows that, across monkeys and motor areas, adding LFP features to models based on kinematics and intrinsic

spiking history leads to no substantial improvement in predictive power. Across sessions and areas, the mean change in predictive power when adding LFP features to a model containing both kinematics features and intrinsic spiking history features ranged from 0.002 to 0.02, and the median change ranged from 0.001 to 0.02. Nevertheless, this median improvement was statistically significant for all but one session (monkey C area PMv session 1) (Wilcoxon signed-rank test, $p < 0.05$ with Bonferroni correction for 19 (session, array) multiple tests).

This result demonstrates that the LFP predictive power not redundant to kinematics features was primarily redundant to information available in the recent 100 ms spiking history in this motor task. In other words, additional single-neuron variability not explained by kinematics seems to be better explained by

fast-timescale features in intrinsic spiking history than by the examined motor cortex LFP features in this reach and grasp task.

Discussion

Neocortical neurons are embedded in large networks possessing highly recurrent connectivity. Recurrent connectivity typically leads to rich spontaneous collective dynamics. The extent to which these spontaneous dynamics contribute to single neuron variability in awake behaving primates, and how these dynamics interact with sensory inputs and behavioral outputs is an important open question in neuroscience. Here we examined this problem in the context of collective dynamics reflected in LFP oscillations at multiple frequencies in three different areas of motor cortex in monkeys performing naturalistic 3-D reach and grasp actions. These LFPs are thought to result, to a large extent, from collective modes of activity driving spatially coherent postsynaptic potentials at multiple spatiotemporal scales (Nunez and Srinivasan, 2006; Buzsáki et al., 2012). LFP features (e.g., amplitude envelope, phase, and analytic signal) in eight different frequency bands predicted single neuron spiking (1 ms time resolution) with significant predictive power for many neurons in all of the three examined motor cortex areas (PMv, PMd, and M1). Neurons for which LFP predictive power was high tended also to show high kinematics predictive power. In fact, this relationship was close to linear (Pearson correlation coefficient ranging from 0.52 to 0.96 across all the studied areas, monkeys and sessions). More importantly, predictive information in the examined LFP features was mostly redundant to the predictive information available in kinematics. In other words, models combining both LFP features and kinematics typically improved only marginally over models using only kinematics in the studied 3-D reach and grasp task. These results should not be dismissed as overfitting artifacts since they were obtained under well controlled L2 regularization aiming to preserve generalization of models with larger number of parameters. Furthermore, in the few cases for which LFP features seemed to add predictive information with respect to kinematics, this information turned out to be redundant to the information available in short term correlations in the intrinsic spiking history. Overall, our findings suggest that multiband LFP oscillations in motor cortex of alert behaving primates, although predictive of single-neuron spiking during movement execution, are primarily related to collective dynamics controlling aspects of motor output (e.g., kinematics) rather than other potential ongoing dynamics not directly related to the task (e.g., arousal levels).

Several previous studies have looked at the relationship between single-neuron spiking and features of LFP oscillations, mostly in sensory cortices and during anesthesia (e.g., Haslinger et al., 2006; Kelly et al., 2010). Recent work by Ecker et al. (2014) has shown that previously reported high correlations between neuronal pairs and strong phase locking to ongoing LFPs in primary visual cortex during stimulation were highly dependent on the anesthesia state, with neuronal ensemble spiking becoming much more asynchronous during awake stimulation tasks. Our analysis goes beyond previous studies by examining motor cortex LFP and spiking in awake behaving non-human primates. Furthermore, to our knowledge, this is the first time that the redundancy between the information available in multiband LFP features and the information available in behavioral output (kinematics) has been systematically assessed in motor cortex. It remains to be seen how much of the residual variability is inherent to stochastic aspects of the biophysics (e.g., noise due to synaptic failure and amplification effects during spike generations; Carandini, 2004), to other motor-related covariates (e.g., torques and muscle activations) not examined in this paper, or to network dynamics not faithfully reflected in LFP features. In the latter, it is possible that the cortical layer from which the electrode tips recorded (likely layer 5 in our data) may impact LFP predictive power. For example, LFPs recorded from layers 2/3 of motor cortex may potentially exhibit different spike prediction performance and different levels of redundancy with respect to kinematics. In addition, we note that typically recorded LFPs might not be as "localized" as previously thought (Kajikawa and Schroeder, 2011). In particular, rhythmic oscillations in electric potentials recorded intracellularly and on broad extracellular fields may share similar frequencies, and yet show very different phase-locking dynamics with respect to neuronal spiking (e.g., Harvey et al., 2009). Thus, the broader LFP spatial average might result in signals that are less predictive of single-neuron spiking and more related to population activity.

The relationship between single-neuron spiking and ongoing LFP oscillations, in particular the locking of neuronal spiking to the phase of oscillations in specific frequency bands, might be highly dependent on the neuron types (e.g., pyramidal vs. fast spiking interneurons; Buzsáki et al., 2012). Recent work by Vigneswaran et al. (2011) has demonstrated that certain types of pyramidal neurons in primary motor and premotor cortices can show features of action potential waveforms and spiking statistics that are indistinguishable from features in inhibitory interneurons. Therefore, an analysis based on such putative classification would remain highly questionable in our motor cortex data.

In our analysis, low frequency (0.3–2 Hz) and higher (>100 Hz) frequency LFP bands tended to contribute the most to prediction of neuronal spiking. The former relate to motor evoked potentials, which are known to be highly correlated with the population spiking (Bansal et al., 2012), and the latter to multiunit activity, whose movement-related modulation also reflects correlated spiking in neuronal populations. Intermediate frequency bands tended to contribute little during movement execution in this type of task. One could raise the possibility that the relationship between LFP features and single-neuron spiking in these intermediate frequency bands could be much more transient than the relationship between spiking and kinematics during movement execution. For example, beta oscillations, even during movement preparation, typically occur in transient, not sustained, events lasting a few or several cycles. Thus, one would like to build models in which spiking phase-locking should be obviously conditioned on the amplitude of the beta filtered LFPs, so that these transients can be properly captured. In this regard, we note that the neural point process models used here should capture such dependence on beta amplitude, since the log-additive form of the models allows

for (nonlinear) multiplicative effects and interactions among different terms (e.g., beta amplitude and phase) in the models. We also note that, although more complex LFP features and models could potentially improve spike prediction, the same could be said about improving the predictive power of motor behavioral covariates by using more complex or a larger set of kinematic features, including for example kinetics (torques) and muscle activation covariates. We hope to be able to examine more complex LFP and motor behavior-related features in the future.

The results reported here on the redundancy between motor-cortex multiband-LFP features and motor behavior are specific to execution of motor tasks in non-human primates who were highly engaged during movement execution. Multiband LFP features also provide reliable biomarkers for broader brain states and their changes. For example, the relationship between single-neuron spiking activity and ongoing LFPs is likely to change substantially depending on anesthesia, drowsiness, resting vs. awake states, attentional and volitional states, as well as stages during motor tasks (e.g., preparation vs. execution). In this broader context, including a larger variety of neural states than examined in this study, we expect multiband LFP features will be an important independent signal to account for neuronal spiking variability not explained by stimuli or behavioral covariates.

Variability in single-neuron spiking activity has often been characterized as of two types: private and shared (e.g., Deweese and Zador, 2004; Churchland and Abbott, 2012; Litwin-Kumar and Doiron, 2012; Goris et al., 2014). Private variability is likely to reflect chaotic nonlinear dynamics in highly recurrent neuronal networks (Litwin-Kumar and Doiron, 2012). Amplification of membrane potential fluctuations by the spiking generation process (Carandini, 2004) in addition to local stochastic factors such as thermal fluctuations and synaptic failure (Faisal et al.,

2008) are also important contributors. On the other hand, shared variability in neuronal ensembles is thought to evolve on slower time scales and reflect representational and computational states in neuronal networks (Churchland and Abbott, 2012; Litwin-Kumar and Doiron, 2012). The examined fluctuations in multiband LFP oscillations seem primarily to be related to this shared variability. Multiband oscillatory LFP activity results in large part from coherent or shared dynamics in neuronal networks. In addition, features in these oscillations that are predictive of single-neuron spiking seemed mostly redundant to parameters in motor behavior. Overall, our finding was that information in the examined multiband LFP features directly relates to these shared representational and computational dynamics across neural populations in motor cortex. Single-neuron activity in motor cortex populations has been shown to be dominated by latent low-dimensional collective dynamics (Truccolo et al., 2010; Churchland et al., 2012). We hope in the future to investigate the relationship between multiband oscillatory LFP activity, in particular slow fluctuations, and latent low-dimensional rhythmic dynamics (Churchland et al., 2012) in motor cortex.

Acknowledgments

References

Arieli, A., Sterkin, A., Grinvald A., and Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273, 1868–1871. doi: 10.1126/science.273.5283.1868

Baker, S. N., Spinks, R., Jackson, A., and Lemon, R. N. (2001). Synchronization in monkey motor cortex during a precision grip task. I. Task-dependent modulation in single-unit synchrony. *J. Neurophysiol.* 85, 869–885.

Bansal, A. K., Truccolo, W., Vargas-Irwin, C. E., and Donoghue, J. P. (2012). Decoding 3D reach and grasp from hybrid signals in motor and premotor cortices: spikes, multiunit activity, and local field potentials. *J. Neurophysiol.* 107, 1337–1355. doi: 10.1152/jn.00781.2011

Bansal, A. K., Vargas-Irwin, C. E., Truccolo, W., and Donoghue, J. P. (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. doi: 10.1152/jn.00532.2010

Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M. A., Logothetis, N. K., et al. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *J. Neurosci.* 28, 5696–5709. doi: 10.1523/JNEUROSCI.0009-08.2008

Brovelli, A., Ding, M., Ledberg, A., Chen, Y., Nakamura, R., and Bressler, S. L. (2004). Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by Granger causality. *Proc. Natl. Acad. Sci. U.S.A.* 101, 9849–9854. doi: 10.1073/pnas.0308538101

Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). The origin of extracellular fields and currents–EEG, ECoG, LFP and spikes. *Nat. Rev. Neurosci.* 13, 407–420. doi: 10.1038/nrn3241

Carandini, M. (2004). Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol.* 2:E264. doi: 10.1371/journal.pbio.0020264

Chen, D., and Fetz, E. E. (2005). Characteristic membrane potential trajectories in primate sensorimotor cortex neurons recorded *in vivo. J. Neurophysiol.* 94, 2713–2725. doi: 10.1152/jn.00024.2005

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., et al. (2012). Neural population dynamics during reaching. *Nature* 487, 51–56. doi: 10.1038/nature11129

Churchland M., and Abbott L. (2012). Two layers of neural variability. *Nat. Neurosci.* 15, 1472–1474. doi: 10.1038/nn.3247

Churchland, M. M., Byron, M. Y., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378. doi: 10.1038/nn.2501

Deweese M. R., and Zador A. M. (2004). Shared and private variability in the auditory cortex. *J. Neurophysiol.* 92, 1840–1855. doi: 10.1152/jn.00197.2004

Ecker, A. S., Berens, P., Cotton, R. J., Subramaniyan, M., Denfield, G. H., Cadwell, C. R., et al. (2014). State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82, 235–248. doi: 10.1016/j.neuron.2014.02.006

Faisal, A. A., Selen, L. P., and Wolpert, D. M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303. doi: 10.1038/nrn2258

Fawcet, T. (2006). An introduction to ROC analysis. *Pattern Recognit. Lett.* 27, 861–874. doi: 10.1016/j.patrec.2005.10.010

Fiser, J., Chiu, C., and Weliky, M. (2004). Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 431, 573–578. doi: 10.1038/nature02907

Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* 33, 1–22.

Goris, R. L. T., Movshon, J. A., and Simoncelli, E. P. (2014). Partitioning neuronal variability. *Nat. Neurosci.* 16, 858–865. doi: 10.1038/nn.3711

Harvey, C. D., Collman, F., Dombeck, D. A., and Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* 461, 941–946. doi: 10.1038/nature08499

Haslinger, R., Ulbert, I., Moore, C. I., Brown, E. N., and Devor, A. (2006). Analysis of LFP phase predicts sensory response of barrel cortex. *J. Neurophysiol.* 96, 1658–1663. doi: 10.1152/jn.01288.2005

Hatsopoulos, N. G., Xu, Q., and Amit, Y. (2007). Encoding of movement fragments in the motor cortex. *J. Neurosci.* 27, 5105–5114. doi: 10.1523/JNEUROSCI.3570-06.2007

Hermes, D., Siero, J. C., Aarnoutse, E. J., Leijten, F. S., Petridou, N., and Ramsey, N. F. (2012). Dissociation between neuronal activity in sensorimotor cortex and hand movement revealed as a function of movement rate. *J. Neurosci.* 32, 9736–9744. doi: 10.1523/JNEUROSCI.0357-12.2012

Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., et al. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* 485, 372–375. doi: 10.1038/nature11076

Jasper, H., and Penfield, W. (1949). Electrocorticograms in man: effect of voluntary movement upon the electrical activity of the precentral gyrus. *Eur. Arch. Psychiatry Clin. Neurosci.* 183, 163–174. doi: 10.1007/bf01062488

Kajikawa, Y., and Schroeder, C. E. (2011). How local is the local field potential? *Neuron* 72, 847–858. doi: 10.1016/j.neuron.2011.09.029

Kayser, C., Montemurro, M. A., Logothetis, N. K., and Panzeri, S. (2009). Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron* 61, 597–608. doi: 10.1016/j.neuron.2009.01.008

Kelly, R. C., Smith, M. A., Kass, R. E., and Lee, T. S. (2010). Local field potentials indicate network state and account for neuronal response variability. *J. Comput. Neurosci.* 29, 567–579. doi: 10.1007/s10827-009-0208-9

Litwin-Kumar, A., and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* 15, 1498–1505. doi: 10.1038/nn.3220

Montemurro, M. A., Rasch, M. J., Murayama, Y., Logothetis, N. K., and Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Curr. Biol.* 18, 375–380. doi: 10.1016/j.cub.2008.02.023

Murthy, V. N., and Fetz, E. E. (1992). Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proc. Natl. Acad. Sci. U.S.A.* 89, 5670–5674. doi: 10.1073/pnas.89.12.5670

Nunez, P. L., and Srinivasan, R. (2006). *Electric Fields of the Brain: the Neurophysics of EEG*. Oxford: Oxford University Press.

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., et al. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454, 995–999. doi: 10.1038/nature07140

Rasch, M. J., Gretton, A., Murayama, Y., Maass, W., and Logothetis, N. K. (2008). Inferring spike trains from local field potentials. *J. Neurophysiol.* 99, 1461–1476. doi: 10.1152/jn.00919.2007

Reimer, J., and Hatsopoulos, N. G. (2010). Periodicity and evoked responses in motor cortex. *J. Neurosci.* 30, 11506–11515. doi: 10.1523/JNEUROSCI.5947-09.2010

Rubino, D., Robbins, K. A., and Hatsopoulos, N. G. (2006). Propagating waves mediate information transfer in the motor cortex. *Nat. Neurosci.* 9, 1549–1557. doi: 10.1038/nn1802

Sanes, J. N., and Donoghue, J. P. (1993). Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proc. Natl. Acad. Sci. U.S.A.* 90, 4470–4474. doi: 10.1073/pnas.90.10.4470

Scheffer-Teixeira, R., Belchior, H., Leão, R., Ribeiro, S., and Tort, A. (2013). On high-frequency field oscillations (>100 Hz) and the spectral leakage of spiking activity. *J. Neurosci.* 33, 1535–1539. doi: 10.1523/JNEUROSCI.4217-12.2013

Shadlen, M. N., and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* 18, 3870–3896.

Truccolo, W., Hochberg, L. R., and Donoghue, J. P. (2010). Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* 13, 1–9. doi: 10.1038/nn.2455

Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., and Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.* 93, 1–17. doi: 10.1152/jn.00697.2004

Truccolo, W. A., Ding, M., Knuth, K. H., Nakamura, R., and Bressler, S. L. (2002). Trial-to-trial variability of cortical evoked responses: implications for the analysis of functional connectivity. *Clin. Neurophysiol.* 113, 1–21. doi: 10.1016/S1388-2457(01)00739-8

Vargas-Irwin, C., and Donoghue, J. P. (2007). Automated spike sorting using density grid contour clustering and subtractive waveform decomposition. *J. Neurosci. Methods* 164, 1–18. doi: 10.1016/j.jneumeth.2007.03.025

Vargas-Irwin, C. E., Shakhnarovich, G., Yadollahpour, P., Mislow, J. M., Black, M. J., and Donoghue, J. P. (2010). Decoding complete reach and grasp actions from local primary motor cortex populations. *J. Neurosci.* 30, 9659–9669. doi: 10.1523/JNEUROSCI.5443-09.2010

Vigneswaran, G., Kraskov, A., and Lemon, R. N. (2011). Large identified pyramidal cells in macaque motor and premotor cortex exhibit "thin spikes": implications for cell type classification. *J. Neurosci.* 31, 14235–14242. doi: 10.1523/JNEUROSCI.3142-11.2011

Waldert, S., Lemon, R. N., and Kraskov, A. (2013). Influence of spiking activity on cortical local field potentials. *J. Physiol.* 591, 5291–5303. doi: 10.1113/jphysiol.2013.258228

Wörgötter, F., Suder, K., Zhao, Y., Kerscher, N., Eysel, U. T., and Funke, K. (1998). State-dependent receptive-field restructuring in the visual cortex. *Nature* 396, 165–168. doi: 10.1038/24157

Zanos, S., Zanos, T. P., Marmarelis, V. Z., Ojemann, G. A., and Fetz, E. E. (2012). Relationships between spike-free local field potentials and spike timing in human temporal cortex. *J. Neurophysiol.* 107, 1808–1821. doi: 10.1152/jn.00663.2011

Zhuang, J., Truccolo, W., Vargas-Irwin, C., and Donoghue, J. P. (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. doi: 10.1109/TBME.2010.2047015