Behavioral/Systems/Cognitive

Encoding of Coordinated Grasp Trajectories in Primary Motor Cortex

Maryam Saleh,[1] **Kazutaka Takahashi,**[2] **Yali Amit,**[3,4] **and Nicholas G. Hatsopoulos**[1,2]

[1]Committee on Computational Neuroscience, [2]Department of Organismal Biology and Anatomy, and Departments of [3]Statistics and [4]Computer Science, University of Chicago, Chicago, Illinois 60637

Few studies have investigated how the cortex encodes the preshaping of the hand as an object is grasped, an ethological movement referred to as prehension. We developed an encoding model of hand kinematics to test whether primary motor cortex (MI) neurons encode temporally extensive combinations of joint motions that characterize a prehensile movement. Two female rhesus macaque monkeys were trained to grasp 4 different objects presented by a robot while their arm was held in place by a thermoplastic brace. We used multielectrode arrays to record MI neurons and an infrared camera motion tracking system to record the 3-D positions of 14 markers placed on the monkeys' wrist and digits. A generalized linear model framework was used to predict the firing rate of each neuron in a 4 ms time interval, based on its own spiking history and the spatiotemporal kinematics of the joint angles of the hand. Our results show that the variability of the firing rate of MI neurons is better described by temporally extensive combinations of finger and wrist joint angle kinematics rather than any individual joint motion or any combination of static kinematic parameters at their optimal lag. Moreover, a higher percentage of neurons encoded joint angular velocities than joint angular positions. These results suggest that neurons encode the covarying trajectories of the hand's joints during a prehensile movement.

Introduction

Whereas a motor act like playing the piano often requires independent control of many of the 20 degrees of freedom of the hand, prehension—preshaping the hand to grasp an object—involves a small set of synergistic movements of the segments of the hand. As such, studies often characterize prehension in a lower dimensional space either by considering grip aperture alone or by reducing the dimensionality using a small set of principal components on the joint angle kinematics (Paulignan et al., 1991; Mason et al., 2002, 2004; Santello et al., 2002; Todorov and Ghahramani, 2004; Hendrix et al., 2009).

Though extensive psychophysics studies suggest that the CNS controls prehension with a limited set of kinematic synergies, few studies in cortex investigate how neurons might encode these synergies. Instead, previous motor cortical encoding studies have shown that the discharge rates of MI neurons, in large time windows (~150 ms), vary with discrete states such as whether a finger joint is flexed or extended, grasp postures, or object categories (Schieber and Hibbard, 1993; Georgopoulos et al., 1999; Poliakov and Schieber, 1999; Umilta et al., 2007; Aggarwal et al., 2009). Recently, one study showed that single MI neurons vary

their discharge rates in time with respect to grip aperture and force and another showed that populations of neurons simultaneously decode the joint kinematics of the arm and hand during reach-to-grasp movements (Hendrix et al., 2009; Vargas-Irwin et al., 2010).

Here, we extend these previous studies by describing how single MI neurons encode the temporal evolution of a set of joint kinematic features of the hand during prehension. In particular, we demonstrate that single MI neurons encode a synergistic set of temporally extensive joint trajectories of the fingers and wrist. We used a generalized linear model framework to predict each neuron's response based on the finger and wrist kinematics as well the spiking history of the neuron (Truccolo et al., 2005, 2008; Hatsopoulos et al., 2007). These spike history terms may capture intrinsic temporal dependencies in the neuron's response (Pillow et al., 2008; Stevenson et al., 2009). We show that both kinematic and spike history terms contributed significantly to the predictive power of the majority of neurons. The model's predictive power significantly increased when we replaced the kinematic parameters at their optimal lag by temporally extensive kinematic trajectories. A larger percentage of neurons encoded joint velocities compared with angular positions and more neurons encoded flexion/extension movements about the wrist and the metacarpophalangeal joint of the index finger compared with other kinematic trajectories. In 51% of neurons, 10 of 12 trajectories were required to reach the maximum predictive power of the kinematics which indicates that neurons tend to encode more than one joint angle parameter. We also tested the contribution of each principal component where each component represented a linear combination of the kinematic trajectories of

the joint angles: in a majority of neurons, one or two principal components accounted for all the predictive power of the kinematics which suggests that neurons encode a low dimensional representation of the joint trajectories that characterize prehension.

Materials and Methods

Behavioral task

Two subjects (female rhesus macaques, monkey O 6.6 kg, monkey A 6.5 kg) were trained to grasp objects with their right hand. The arm was constrained with a thermoplastic brace, at 170° extension in the elbow joint, and 80° elevation (abduction) in the shoulder joint. A grasp trial consisted of 4 different periods: a premovement period, a movement period preceding the grasp, a hold period when the subject held the object, and a release period when the robot retracted the object (Fig. 1A). At the start of a trial, the robot moved the object toward the subject's hand. The trajectory and speed of the robot were designed to approximately simulate actual hand transport. Trajectories were fine tuned so that the object approached the subject's hand in the same orientation that the subject's hand approached the object in a free reaching condition. The subject preshaped her hand to grasp the approaching object. The subject was instructed to hold the object until the robot retracted it. Three objects were presented, requiring a total of 4 hand configurations (Fig. 1B). The first configuration was a key grip, where the monkey held a thin flat object between the thumb and the side of the middle phalanx of the index finger; the second object, the D-ring, was presented twice, in configurations that elicited a whole hand grip (power) at two different wrist orientations. These two different configurations will be treated as separate objects—D-ring horizontal and D-ring vertical—in the rest of this paper. The third object, a sphere, elicited a whole hand grip with a fanning of the fingers. These objects were designed to evoke a variety of coordination patterns among the joint angle kinematic trajectories. Each object was presented to the subject in a block of at least 30 trials.

Neural data acquisition

Each monkey was chronically implanted with a 100 electrode (400 μm interelectrode separation, 1.5 mm electrode length; BlackRock Microsystems) microelectrode array in the arm-hand area of the primary motor cortex in the left hemisphere, contralateral to the hand they used to grasp the objects. The array was placed parallel and as close as possible to the central sulcus, medial from the spur of the arcuate sulcus (Fig. 1C). Surgical details have been described previously (Rousche and Normann, 1992; Maynard et al., 1997, 1999). During each recording session, signals were amplified (gain 5000×), bandpass filtered between 0.3 and 7.5 kHz, and digitized at 30 kHz using the Cerebus Neural Data Acquisition system (BlackRock Microsystems). For each channel, a threshold was set above the noise band: if the signal crossed the threshold, a 1.6 ms window of the signal—as to yield 48 samples given a sampling frequency of 30 kHz—was sampled around the occurrence of the threshold crossing. To classify the waveforms as noise or units on the given channel, the waveforms were spike-sorted offline using Offline Sorter (Plexon) and/or exported to Matlab (MathWorks) and spike-sorted using custom software.
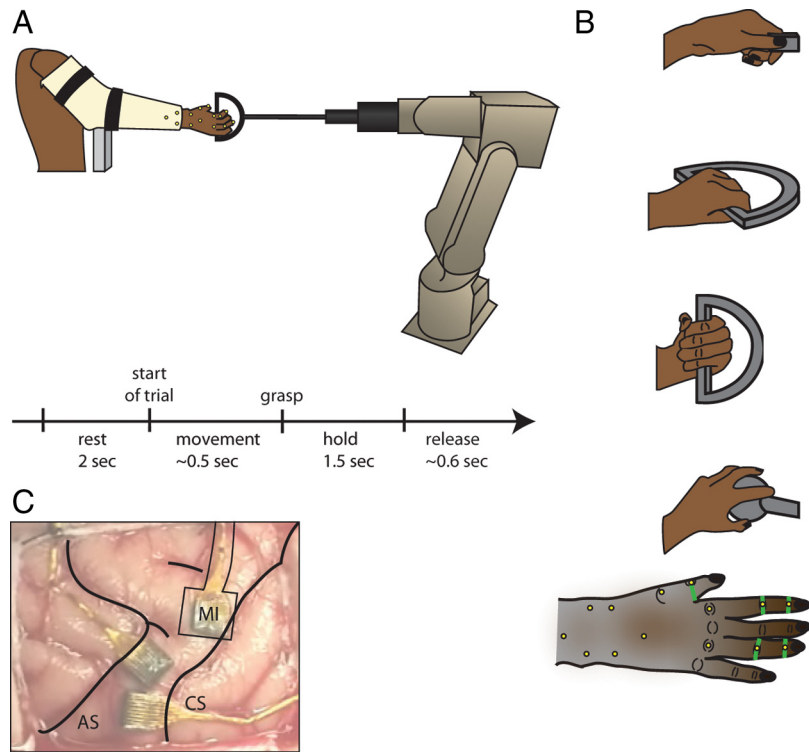

Figure 1. Experimental methods. *A*, Top, The subject's arm was held with a thermoplastic brace while the robot presented and retracted the objects from the subject's grasping range. *A*, Bottom, Task timeline. The trial started as the robot started to move the object toward the subject's hand (start of trial). As the robot moved closer to the hand, the subject preshaped her hand to grasp the object (movement period). The subject was instructed to hold the object until the robot retracted it at which point the hand released the object. *B*, Four grasp postures were tested. The first object elicited a key grip; the second object, the D-ring, was presented in two different configurations that elicited a whole hand grip at two different wrist orientations. The third object, a sphere, elicited a whole hand grip with a fanning of the fingers. These objects were designed to produce a variety of coordination patterns between the joint angle kinematics trajectories. Fourteen retro-reflective markers were placed on the subject's digits and wrist. Six infrared cameras (Vicon) captured the 3-D position of these markers during the task. These data were used to calculate the joint angle kinematics of the hand offline. *C*, While the subjects performed the task, neurons were recorded from the arm-hand area of MI with a 100 electrode microelectrode array. The two other arrays were implanted in MI-orofacial and the ventral premotor cortex and were intended for other studies. AS, Arcuate sulcus; CS, central sulcus.

Motion tracking

The three-dimensional (3-D) positions of monkeys' wrists, hands, and fingers were recorded using a video-based motion analysis system (Vicon Motion Tracking System, Workstation 460). We used six M2 cameras at 1.2 megapixel resolution. Fourteen spherical retro-reflective markers (3 mm diameter) were glued to the monkeys' fingers, hand and forearm (Fig. 1A). Their 3-D positions were used to calculate the angles of the following joints: wrist flexion/extension (WR), the metacarpophalangeal joint of the index finger (IMP), the proximal interphalangeal joint of the index finger (IPIP) and ring finger (RPIP), the angle between the proximal interphalangeal joints of the index and ring fingers with the metacarpophalangeal joint of the middle finger as the origin (FAN), and, last, the angle between the metacarpals of the thumb and index finger with the vertex at the base of the radius (TI). The angle TI is projected on the plane that is formed by the metacarpal bone of the index finger and the metacarpal bone of the thumb. A custom software package (TheGame, version 2) was used to remotely control the trial capture and align the kinematics data with the neural data via a synchronizing pulse that was sent to both systems every 150 ms. The kinematics data were sampled at 250 Hz and bidirectionally filtered with a fourth order Butterworth low-pass filter with 6 Hz cutoff. All data filtering and calculations were done in MATLAB. The start of movement and the time of grasp were determined by analyzing the kinematics—the start of movement corresponded to the time at which the first marker's speed profile exceeded 10 cm/s. To determine the grasp time, we calculated the time at which the last marker's speed profile dropped below 10 cm/s.
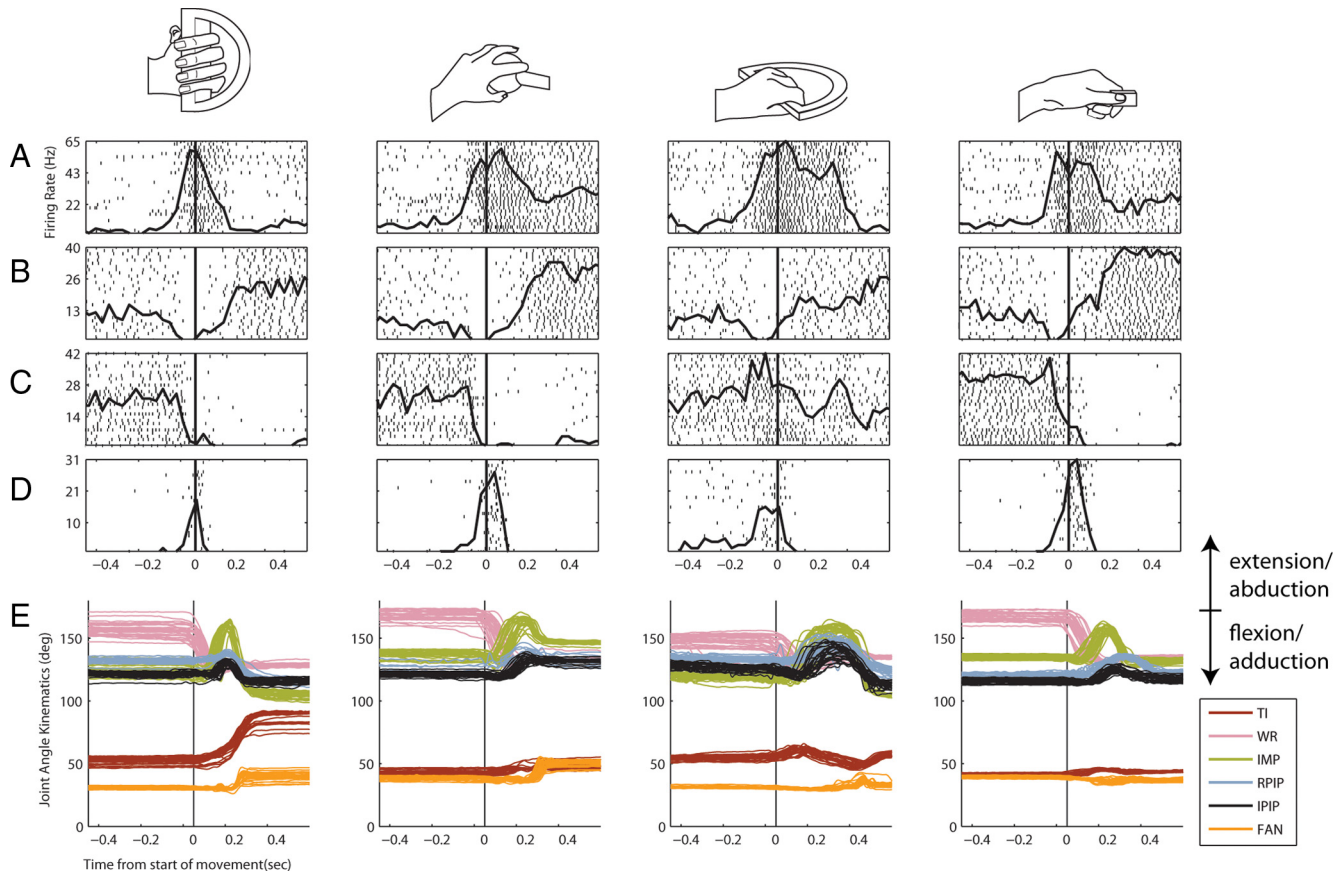
Figure 2. Spike rates, spike trains, and kinematics traces during the grasping task, for all objects. ***A–D***, Spike trains across trials for 4 neurons, aligned at the start of movement. The average firing rates are superimposed as solid black lines. ***E***, Joint angle trajectories for all objects in one data set. The 6 joint angles that were collected were the angle between the metacarpals of the thumb and index finger with the origin at the base of the radius (TI, red), the wrist flexion/extension (WR, pink), the metacarpophalangeal joint of the index finger (IMP, green), the proximal interphalangeal joint of the ring finger (RPIP, blue) and index finger (IPIP, black), and the angle between the proximal interphalangeal joints of the index and ring fingers with the metacarpophalangeal joint of the middle finger as the origin (FAN, orange).

Computational model

We developed an encoding model to predict the spiking response of each neuron based on a set of external and internal covariates. The external covariates consisted of the kinematic trajectories of a group of finger and wrist joints. The internal covariates included the history of spiking responses of the neuron on multiple time scales. These covariates may capture effects such as intrinsic temporal correlations due to refractoriness, bursting, rhythmicity, and short-term plasticity. We used a generalized linear modeling framework where the conditional intensity function λ_t, the neuron's instantaneous spike count in a 4 ms time bin, is estimated as a function of a set of covariates—the joint kinematics and the neuron's firing history. A generalized linear model has the advantage over standard multiple linear regression in that it is designed for discrete, integer responses (i.e., numbers of spikes emitted in a time bin) that are not normally distributed as assumed by standard regression. We linked the rate function to the covariates through a log function, to ensure non-negative estimated rates. That is, the logarithm of the firing rate of the neuron is a linear function of the covariates. More details about the application of the generalized linear model to neural data can be found previously (Paninski et al., 2004a; Truccolo et al., 2005).

Spike history. Setting aside the effects of external covariates, the current spiking response of a neuron depends on the neuron's spiking response in the past. Such spike-history effects may include short time-scale phenomena such as refractoriness and bursting and longer time-scale oscillations, facilitation and depression phenomena occurring on the order of tens to hundreds of milliseconds (Hille, 2001; Keat et al., 2001; Chen and Fetz, 2005). To model these spiking history

components, we filtered the spike train with a basis set of raised cosines of the form:

$$b_j(t) = \frac{1}{2} \cos(a \log [t + c] - \phi_j) + \frac{1}{2} \quad (1)$$

for t, such that $a \log(t + c) \in [\phi_j - \pi, \phi_{j\,+}\,\pi]$, and 0 previously (Pillow et al., 2008). The ϕ_j are the x-axis positions of the peaks of the cosine curves. Their positions were selected as equidistant on the regular x-axis scale—that is, before applying the log transform, the basis centers were equidistant from each other. After the log transform, the basis centers were at 8 ms, 12 ms, 20 ms, 32 ms, 60 ms, 108 ms and 208 ms respectively. The resulting basis set b_j, where $j \in [1,7]$ accounts for both fine temporal correlations near the spike and broader correlations at longer time lags from the spike. Each basis vector was convolved with the spikes $\rho(t)$ to give rise to $J = 7$ spike history vectors H_j, representing the effect of spike history up to 200 ms in the past:

$$H_{t,j} = \sum_{\tau=t-T}^{t-1} b_j(\tau)\, \rho(t - \tau) \quad (2)$$

where T is the duration of the basis vector.

The following spike history model is used as a control model excluding all external covariates:

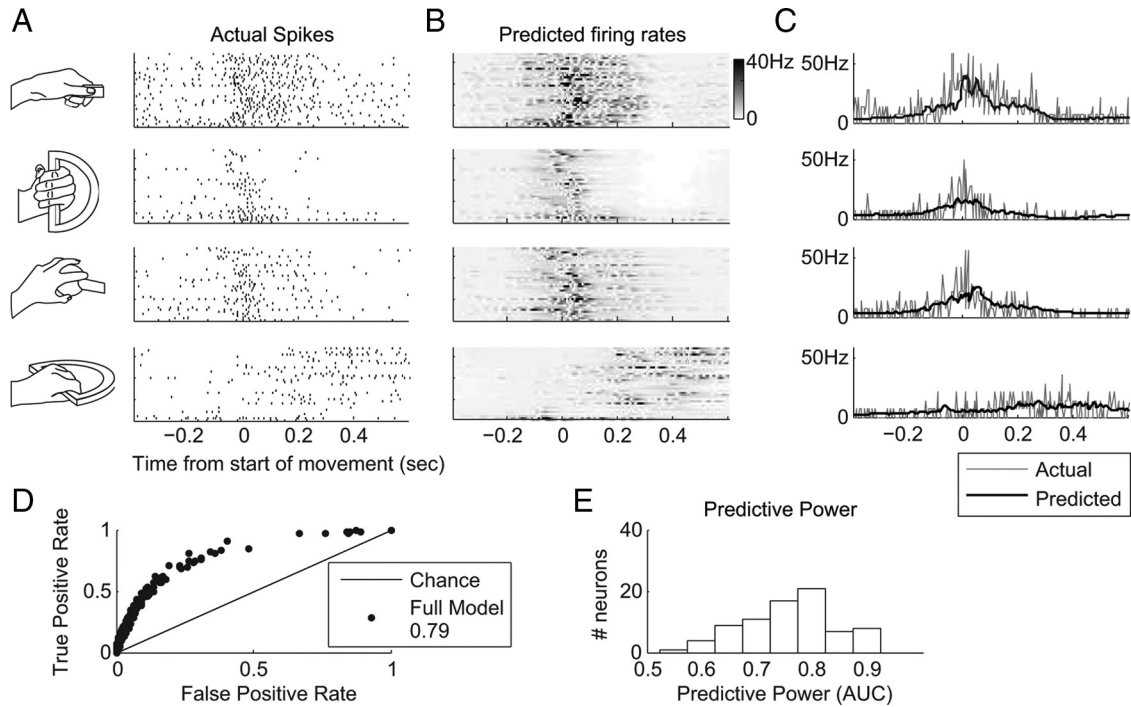$$\log \lambda_t(H_t) = \beta_0 + \sum_{j=1}^{J} \beta_j^H H_j \quad (3)$$

Figure 3. ROC analysis: prediction results. *A*, Spike trains for a given neuron: 30 rows of trials aligned at the start of movement for the four object conditions. *B*, Predicted firing rates across the same 30 trials as in *A* using the full model. Grayscale color bar indicates the range of firing rates, where a white color would indicate a firing rate of 0 and a black color would indicate a firing rate of 40 ms bins. *C*, Comparison of actual (thin line) and predicted firing rates (thick line) calculated in 4 ms bins. *D*, Example of ROC curve for the given neuron. The predictive power of the model for the neuron is the area under the ROC curve, which was calculated to be 0.79. *E*, Distribution of predictive power (AUCs) for all neurons. AUCs can range between 0.5 and 1.

The β_j^H terms represent the weights associated with the history terms and β_0 can be taken to be the baseline firing rate of the neuron.

Kinematic variables. To model velocity and position together with spike history, we used the following model:

$$\log \lambda_t(H, \dot{\theta}, \theta) = \beta_0 + \sum_{j=1}^{J} \beta_j^H H_j + \sum_{f=1}^{F} \sum_{r=t-164}^{t+200} \beta_{r,f}^{\dot{\theta}} \dot{\theta}_{r,f}$$

$$+ \sum_{f=1}^{F} \sum_{r=t-164}^{t+200} \beta_{r,f}^{\theta} \theta_{r,f} \quad (4)$$

where r represents the time lags between the kinematics and the spiking, ranging from -164 ms to 200 ms after the spiking activity, with a step size of 52 ms, f denotes a particular kinematic feature, and F is the total number of angular or angular velocity features ($F = 6$; θ is the joint angle and $\dot{\theta}$ is the joint angle velocity). Each term associated with a coefficient β is a covariate. The total number of kinematic terms is, therefore, 96 = (6 angles + 6 angular velocities) $\times$ 8 time lags. In the remainder of the text, we use "trajectories" to denote joint angles or angular velocities of the fingers and wrist at multiple time lags relative to the spikes. The spike history terms and kinematic terms on the right hand side of the equation are also referred to as covariates of the conditional intensity function.

Assessing the significance of the kinematic terms at each lag. For each neuron, we assessed the significance of the kinematic terms at each lag by using a goodness-of-fit test: the χ^2 test on the deviances between two models. The first model was the full model in Equation 4. The second model was the same as the full model but with the tested term removed. As such, the second model's terms were a subset of the terms in the full model. When testing two models within the generalized linear model framework, one of which has a subset of the terms of the other model—a model that is "nested" in the first, the difference in the deviances between the two models follows a χ^2 distribution (McCullagh and Nelder, 1989). An increase in deviance from the full model to the nested model would indicate an increase in fitting error due to the exclusion of the tested

term. We used the χ^2 test to assess whether the change in deviance was significant at a *p*-value smaller than 0.05.

Principal components. We took advantage of the fact that the kinematic terms are correlated during grasping and computed the principal components of all the 96 kinematic terms to reduce the number of parameters in the model. The kinematic data were sampled in 4 ms bins, where one observation corresponded to the trajectories of all the kinematic parameters (96 terms). The next observation corresponded to the same parameters sampled 4 ms after the samples from the previous observation such that total number of observations corresponded to the number of trials multiplied by the number of 4 ms samples per trial (Hatsopoulos et al., 2007). We computed the principal components of the 96 kinematic variables for these observations such that the last two terms from the right hand side of Equation 4 were replaced by the number of principal components that accounted for 90% of the variance. The number of principal components ranged between 10 and 15 principal components (P in Eq. 5) depending on the dataset. The covariates PC in Equation 5 refer to the first P principal components; that is, the projections of the kinematic data onto the first P eigenvectors.

$$\log \lambda_t(H, PC) = \beta_0 + \sum_{j=1}^{J} \beta_j^H H_j + \sum_{p=1}^{P} \beta_p PC_p \quad (5)$$

We refer to the "full model" as a model that included all the terms in the right hand side of Equation 5.

Prediction accuracy. To test for the prediction accuracy of a model, we computed the area under the receiver operating characteristic (ROC) curve (AUC) for all test data, using 10 folds of cross-validation (i.e., 10 distinct sets of test data that were not used to build the model). The receiver operating characteristic curve is a plot of the true-positive versus false-positive rates of prediction. To calculate true-positive and false-positive rates at different thresholds, we started by aligning the real spike counts with the predicted spike counts. We then looped through a set of thresholds that matched the range of conditional intensity estimates (the predicted data). The conditional intensity estimate represents the firing
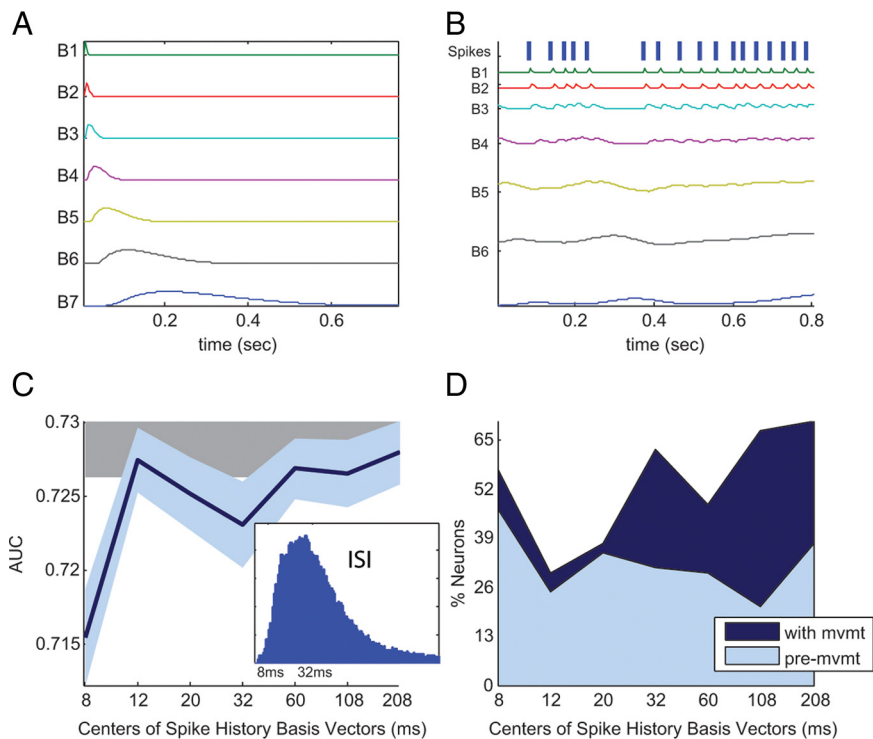
Figure 4. Spike history. **A**, Spike history basis vectors B1 to B7. **B**, Actual spike raster and the same spikes convolved with the spike basis vectors B1 to B7 to yield the spike history terms in the encoding model. Time 0 does not reflect the start of data collection so modulation before the occurrence of the first spike corresponds to history effects from spikes that occurred before time 0. **C**, Prediction results for a given neuron. The dark gray bar represents the predictive power [area under ROC curve (AUC)] of the full model. The mean predictive power of the full model with one history term removed at a time is shown in dark blue. The error bars are SEs across folds and are depicted in the lighter shade of blue. For this cell, removing the first, third or fourth spike history terms significantly decreased the predictive power of the neuron at a significance level of $p < 0.05$. Spike history terms are referred by the centers of their respective basis vectors B1 to B7. Inset, Interspike interval distribution of a given cell. **D**, Percentage neurons (total of 78) that encoded each spike history term across all experiments during movement (dark blue) and during premovement (light blue) periods.

rate estimation in a 4 ms bin and typically ranged from 0 to ~0.2. A value of 0.2 would correspond to a firing rate of 50 Hz. We looped through 50 equidistant thresholds between 0 and the maximum value of the conditional intensity estimate. For each threshold, we calculate the total number of positives, namely the number of timestamps for which the estimated conditional intensity was above the threshold. Of these, the hits or true positives correspond to actual spikes, and the rest are false positives. Each point in Figure 3D (see below) corresponds to one of these calculations; that is, one false-positive and true-positive reading at a given threshold for one cross-validation fold. For each fold, we computed the AUC. Numerically, this means we connected the dots associated with one fold in Figure 3D and computed the piecewise trapezoidal area between each set of two consecutive dots. These areas were summed to yield the total area under the ROC curve for each fold. The AUC provides an assessment of the predictive power of the encoding model and ranges between 0.5 and 1.0, where 0.5 represents chance level prediction and 1.0 is a perfect prediction (see Fig. 3D). Specifically, when two different samples are randomly drawn from the data (one sample contains a spike and the other does not), the AUC represents the probability that the model will yield a higher probability for the sample with a spike. More details about ROC curve analysis can be found in other articles (Fawcett, 2006; Hatsopoulos et al., 2007; Truccolo et al., 2008, 2010). We used a Wilcoxon signed rank test to compare the AUCs of neurons for different tested models. If the test involved multiple comparisons, we applied a Bonferroni correction to the p-values.

Results

We collected two datasets that were 10 days apart per subject, where each dataset included at least 30 grasps to each of the 4

objects. In the first monkey, we found 23 of 41 and 22 of 39 task-related MI neurons for the first and second experiments respectively; in the second monkey, we found 20 of 35 and 13 of 17 task-related MI neurons. A neuron was considered to be task-related if its mean firing rate during the movement period was significantly different from its mean firing rate during the premovement period (paired t test across trials, $p < 0.05$). In the remainder of the text, when referring to neurons, we mean cell samples that were task-related. We have previously shown that for a given subject implanted with a Utah array, 43% of units are stable across 10 d on average (Dickey et al., 2009). Taking this estimation into consideration, we expect that in the same animal, ~9 task-related cell samples overlapped from one experiment to the next. Firing rates varied across a trial, across objects and across neurons (Fig. 2A–D). The starting postures, the end postures, and the kinematic trajectories varied across object task conditions (Fig. 2E). IMP (green), IPIP (blue) and RPIP (black) show patterns that reflect an extension and flexion of the proximal and middle phalanges of the fingers. Flexion of the fingers is least pronounced for the sphere object indicating that the end posture had the hand in a semi-open configuration. TI (red) is the angle between the metacarpal bones of the thumb and index finger with the origin at the base of the radius and represents the abduction/adduction of the thumb. The thumb was abducted most prominently when the monkey grasped the vertical D-ring. This subject did not tend to oppose her thumb to the rest of the fingers when grasping the sphere, horizontal D-ring, or key-grip object. Instead, when grasping these objects, she tended to keep her thumb adjacent to the metacarpal bone of the index finger. FAN (orange) is an analog for the abduction/adduction of the proximal phalanges. The monkey fanned her fingers to grasp both the sphere and vertical D-ring but not the other two objects. WR (pink) represents wrist flexion and extension. The robot brought the objects to the monkey's hand in a manner that prevented the need for an extension of the wrist before the grasp. This was a necessary precaution to prevent any injuries of their hands in the event that they would not grasp the object.

Prediction results

The encoding model accurately predicted the precise timing and firing rates of the cells at a 4 ms resolution. The spiking responses as well as the predicted responses from the model for an example neuron are shown for each of the four objects (Fig. 3A–C). The predictive power of the model for this neuron—the area under the ROC curve (AUC)—yielded a value of 0.79. On a population level, the median predictive power of the model for all neurons was 0.7608 which was significantly above chance (Fig. 3E, Wilcoxon signed rank test—rank$_{\text{Wil}}$—to as-

sess whether the AUCs of all neurons come from a distribution whose median is 0.5, $p < 0.001$).

Spike history

We used 7 basis functions to model the effect of a neuron's own history at different time scales (Fig. 4A,B) (see Materials and Methods). We tested whether particular spike history terms were encoded by each neuron. To test the significance of each history term for a neuron, we compared the full model against one with the test history term removed. A significant decrease in AUC indicated that the given history term was significant for that neuron. Figure 4C displays the results of this analysis for a given neuron: for this neuron, removing the first and third spike history terms (Fig. 4A,B; B1 in green and B3 in cyan) yielded a significant decrease in predictive power (paired rank$_{Wil}$, $p < 0.05$). We found that between 29% and 68% of neurons encoded each of the seven history terms (Fig. 4D, dark blue). We also examined whether history effects differed when no movement was occurring. We recomputed the encoding model using data that were collected during the premovement period. The proportion of neurons with significant history effects on a short time-scale (i.e., first through third history terms) did not differ substantially during the premovement period. In contrast, there was at least a 45% drop in the number of neurons that encoded the fourth, sixth and seventh spike history terms (Fig. 4D, light blue) though each history term was encoded by at least 21% of the neurons.

Kinematics

We tested whether neurons encode temporally extensive, kinematic trajectories compared with kinematic parameters measured at a single time lag with respect to the spike times. To do this, we first built an encoding model using single time lags to determine the optimal lag time for each neuron assessed with the goodness-of-fit test (see Materials and Methods). Figure 5, A and B, displays the distribution of optimal lags for all neurons, calculated for the joint angle velocities and joint angle positions respectively. A plurality of neurons yielded the lowest p-value (see Materials and Methods, χ^2 test on deviance) at a positive lag of 148 ms or 200 ms. A positive lag means the kinematics lag the neural activity whereas a negative lag indicates the kinematics lead the neural activity. We then built a model by cumulatively adding neighboring lags, starting at a lag of 200 ms, up to a negative lag of −164 ms yielding a largest trajectory length of 364 ms. Figure 5C displays the mean difference in predictive power between the tested trajectory length and the trajectory length at which the maximum predictive power was achieved. The trajectory lengths at which the lowest absolute values on this graph are shown represent the lengths at which the highest predictive power is achieved. The error asymptotically converged to 0 at a path length between 312 ms and 364 ms. Since the errors at trajectory lengths of 312 ms and 364 ms were significantly different (paired rank$_{Wil}$, $p = 0.0092$), we concluded that a trajectory
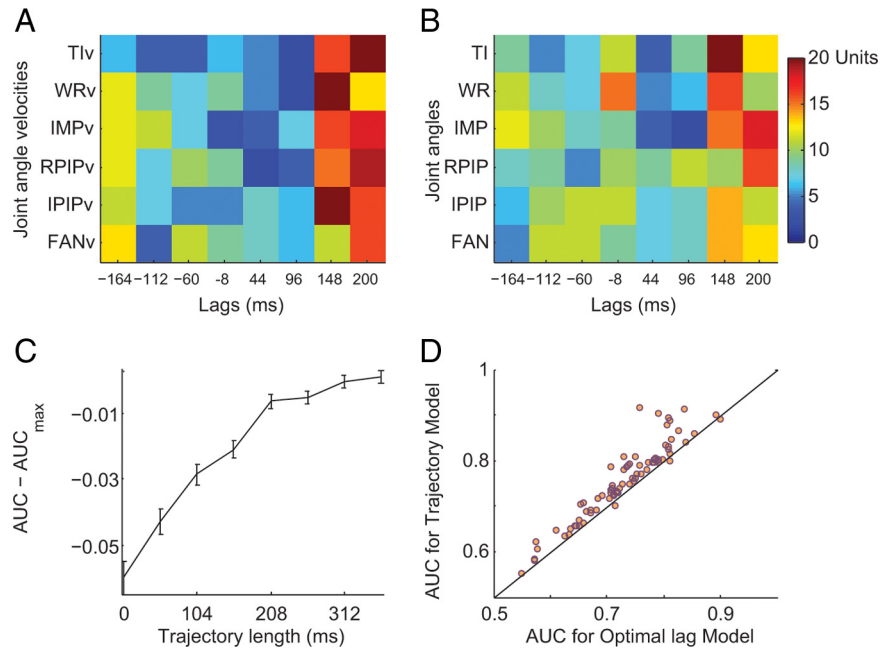
Figure 5. Kinematics: lag analysis. ***A***, Distribution of optimal lags (see Materials and Methods) for joint angle velocities across neurons. A plurality of neurons showed an optimal lag at 148 and 200 ms. ***B***, Same as *A*, but for joint angles. More neurons also showed an optimal lag at 148 and 200 ms. ***C***, Increase in predictive power as a function of trajectory length. The predictive power is given as AUC − AUC$_{max}$, which is the difference in predictive power between the model at trajectory length *x* and the maximum predictive power across all trajectory lengths. We report the mean and SE of these values across all the neurons. The maximum predictive power is generally achieved at a trajectory length of 364 ms where the difference converges to 0. Trajectories were tested starting at a lag time of 200 ms, cumulatively extended to a negative lag time of −150 ms relative to the timing of the spikes. ***D***, AUC values across neurons between a model with a trajectory length of 364 ms (*y*-axis) versus values with all the features only at their respective optimal lags (*x*-axis).

length of 364 ms was optimal. Next, we compared the predictive power of an encoding model with the 12 kinematic terms sampled at a trajectory length of 364 ms to one with the same terms sampled only at their optimal lag (Fig. 5D). In 90% of the neurons, the AUC was significantly higher for the full model compared with one with spike history terms and the kinematics at just their optimal lags (paired rank$_{Wil}$, $p < 0.05$). The difference in AUC between these two models represents the added predictive power when incorporating the 364 ms trajectory. At a population level, the AUC was higher for the full model compared with a model with the kinematics covariates only at their optimal lags (paired rank$_{Wil}$, $p < 0.001$).

We also tested whether particular kinematic trajectories were encoded by each neuron. Each kinematic trajectory will be referred to as a feature. To test the significance of each feature for a neuron, we compared the full model against one with the test feature removed. A significant decrease in predictive power indicated that the given feature was significant for that neuron. As an example for a single neuron, Figure 6A displays the mean AUC for the full model in dark gray and the mean AUC for a model with one of the 12 features removed at a time in blue. The feature that was removed from the model is indicated on the *x*-axis. For this given neuron, removing either of WRv, IMPv, IPIPv, TI, TIv, or WR yielded a significant decrease in predictive power (paired rank$_{Wil}$, $p_{WRv} = 0.002$, $p_{IMPv} = 0.002$, $p_{IPIPv} = 0.003$, $p_{TI} = 0.037$, $p_{TIv} = 0.041$, and $p_{WR} = 0.049$). All features were represented though IMPv and WRv were more frequently represented than other features and joint angle velocities as a whole were also 3 times more frequently represented than joint angle positions (Fig. 6B). Across the entire population of neurons, a median of 3
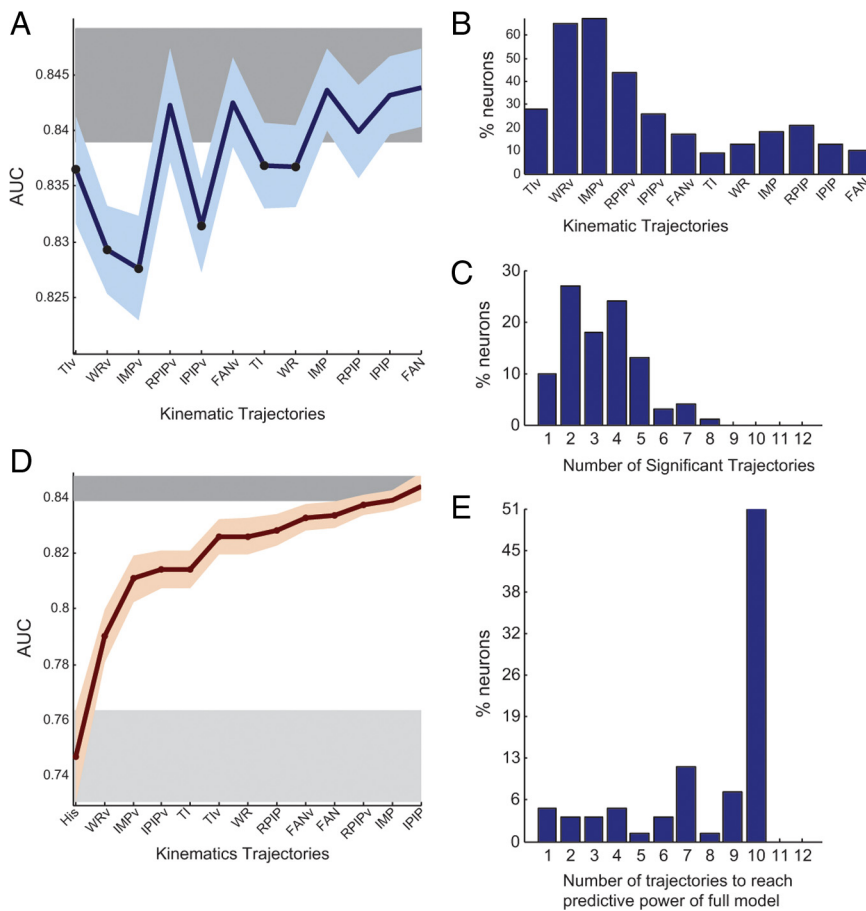
Figure 6. Single versus multiple kinematic trajectories. **A**, Predictive power (AUC in blue) with the removal of one kinematic trajectory at a time for a given cell. AUC values for full model are shown in dark gray where the bar width represents the mean ± SE. AUC values for the full model with just one kinematics trajectory extracted at a time are shown in blue. The error bars are SEs across folds and are depicted in the lighter shade of blue. For this cell, the trajectories of WRv, IMPv, IPIPv, TI, TIv, and WR significantly decreased the predictive power of the cell when they were removed from the encoding model (Wilcoxon signed rank test at $p <$ 0.05 in all cases). **B**, Percentage of neurons exhibiting a significant drop in predictive power when removing each kinematic trajectory individually. More cells encoded the WRv and IMPv trajectories than any other kinematic trajectory though at least one cell encoded any given kinematic trajectory. **C**, Number of significant trajectories across cells. **D**, Starting from the predictive power of the history model, kinematic trajectories were added in their order of significance for the given cell. The predictive power of the full model was achieved with 10 trajectories (signed rank test, $p < 0.05$). **E**, Population results. Fifty-one percent of cells required 10 features to reach the full model's predictive power when the trajectories were added in their order of significance (most significant first, signed rank test. $p < 0.05$ in all cases).

features were significant per neuron (paired rank$_{Wil}$, $p < 0.05$ in all cases; Fig. 6C).

We then tested how many features were required to attain the full model's predictive power for each of the neurons. To do so, we cumulatively added features to the model until its predictive power was not significantly different from the predictive power of the full model. Figure 6D displays the increase in predictive power with the addition of each feature for a given neuron, when the features were added in their order of significance For the example cell, the full model's predictive power was achieved when 10 features were added, even though only 5 features were significant (paired rank$_{Wil}$, $p = 0.047$). Across cells, a median of 10 features were required to attain the full model's predictive power (Fig. 6E). These results suggest that a majority of features are necessary to account for the response variability of the neurons.

Given that our model requires combinations of multiple kinematic features to most accurately predict neuronal response, we sought to test whether neurons encode synergistic, coordina-

tion patterns of hand motion that occur during grasp. The principal components (PC) of all the kinematics features represent linear combinations of the joint angle kinematics across a 364 ms time window and can be interpreted as coordination patterns among the 12 kinematic feature trajectories. The first two eigenvectors of the principal components were highly consistent across datasets and across monkeys (Fig. 7A). Looking at the joint angle position terms, PC1 models a simultaneous flexion of the wrist and extension of the fingers, followed by flexion of the fingers. The IPIP feature for the second experiment in the first subject yielded a slightly larger magnitude in the movement trajectory. This is likely due to a different strategy used on that day. The IMP feature also shows magnitude differences in the movement trajectories across subjects and this is also likely due to differences in the strategies used by the monkeys. TI and FAN do not show as much modulation as the other features and are not as consistent across experiments; however, for three of the four experiments, PC1 shows an abduction of the fingers followed by an adduction. PC2 shows an extension followed by flexion of the wrist with a simultaneous flexion and extension of the fingers. In 3 of 4 experiments, it also shows a later abduction followed by a slight adduction of the fingers. We applied the same procedure to the PCs as we did to the features: we systematically removed one PC at a time and determined whether there was a significant decrease in the model's predictive power (Fig. 7B). In the given example cell, a significant decrease in predictive power could only be achieved by removing the second principal component (Fig. 7B, paired rank$_{Wil}$, $p < 0.001$). On a population level, 36% of neurons encoded the first and second principal components (Fig. 7C). To assess how many components were required to achieve the full predictive power of the cell, we cumulatively added the components to the encoding model, in their order of significance. When the components were added in this order, 81% of cells required between 1 and 3 components to attain the full model's predictive power of the model (Fig. 7E).

Spike history versus kinematics

The actual spike trains and the predicted firing rate using the full model for an example neuron are shown in Figure 8B. To examine the relative contribution of the kinematics and history terms of the model, we built and tested an encoding model that included just the kinematics terms (Fig. 8C; data from all 4 objects are shown in sequential blocks) and one that included just the spike history terms (Fig. 8D; data from all 4 objects are shown in sequential blocks). We observed that the kinematics model captured the coarse temporal structure of the spike train while the spike history model captured the fine and coarse temporal

structure of the spike train but at a temporal lag from the actual spiking activity (Fig. 8*E*). In fact, the maximum correlation between the actual spikes and the predicted firing rates of the kinematic model occurred at an offset of −4 ms (median of offsets across neurons). The distribution of offsets between the actual spikes and the predicted firing rates of the spike history model was bimodal, with a peak at 0 ms and one at −28 ms. Many more neurons yielded an offset around −28 ms such that the overall median of the distribution was at −28 ms. These results indicate that the kinematics yield more precise onset timing information than the spike history terms. In terms of predictive power, both the kinematics and spike history models performed significantly better than chance on a population level (rank$_{Wil}$, p_{kins}<0.001; p_{hist}<0.001). The kinematics and the history terms contributed to varying degrees to the predictive power of cells: for 38% of neurons, the predictive power of the kinematic model was significantly higher than the predictive power of the history model and for another 40% of neurons, the predictive power of the spike history model was significantly higher than the predictive power of the kinematic model (paired rank$_{Wil}$ s <0.05). However, as a population, there was no significant difference between the kinematics and history models (Fig. 8*F*, rank$_{Wil}$, $p = 0.0319$). In contrast, for 84% of neurons, the full model, which included both history and kinematic terms, performed significantly better than the kinematics model (Fig. 8*G*; paired rank$_{Wil}$, $p < 0.05$ in all cases). And the full model performed better than the history models alone in 96% of the neurons (Fig. 8*H*; paired rank$_{Wil}$, $p < 0.05$ in all cases).

Generalization
A stronger test for the validity of an encoding model is to test its predictive power in a different context. We built the encoding model with data from 3 of the 4 object conditions and tested it on the data from the fourth object condition. We call these the "restricted" models. We repeated this procedure 4 times, to test whether the kinematics and spike history from any 3 object conditions are sufficient to predict the firing rate of a neuron recorded while a fourth object was grasped. The actual spike trains and predicted firing rates for a given neuron in each test condition are displayed in Figure 9, *A* and *B*. To assess how well the model generalizes in each test condition, we compare the difference in predictive power of the cells between the 4-object model and each of the restricted models (Fig. 9*C,D*). Though the mean ranks of the AUCs between the 4-object model and any of the restricted models were significantly different (Kruskal Wallis ANOVA test followed by paired rank$_{Wil}$, $p < 0.05$ in all cases), the predictive power of the neurons were well above chance (rank$_{Wil}$
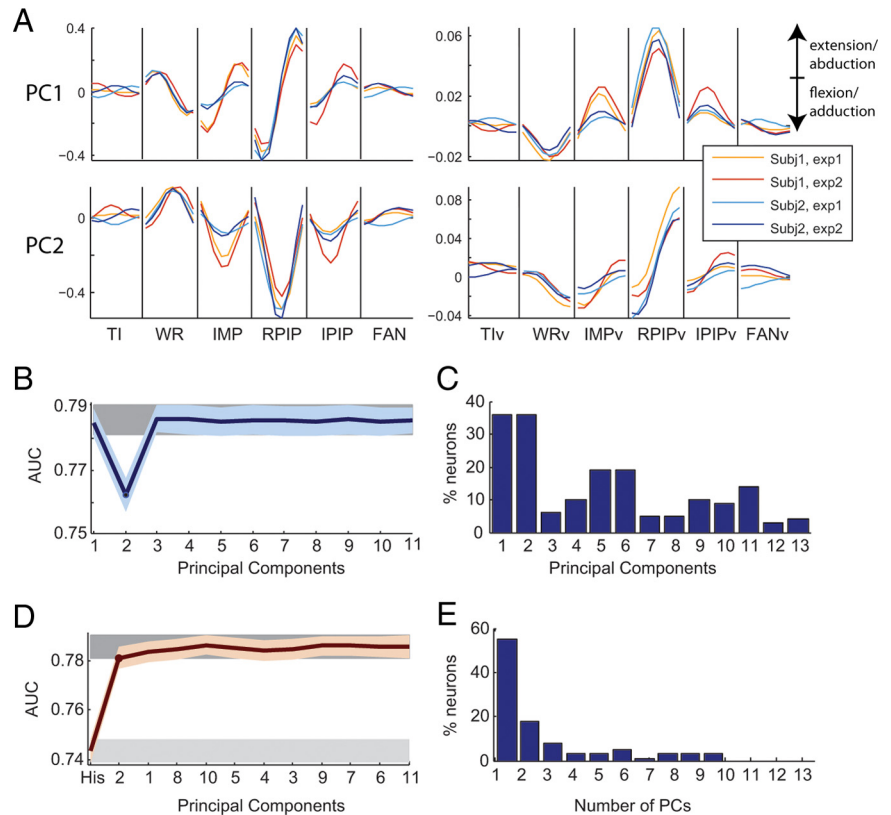

Figure 7. Principal components as kinematic synergies. *A*, Eigenvectors for the first (top) and second (bottom) principal components calculated for the 6 trajectories of joint angles (left) and joint angle velocities (right). The first point of each feature's trajectory corresponds to the same time lag. *B*, Predictive power (AUC in blue) with the removal of one principal component for a given cell. For this cell, the second component significantly decreased the predictive power of the cell when it was removed from the encoding model. We used a Wilcoxon signed rank test at $\alpha = 0.05$ to test for a significant drop between the predictive power of the full model and the predictive power of the same model with one extracted component. *C*, Percentage of cells exhibiting a significant drop in predictive power when removing each principal component individually. Thirty-six percent of neurons encoded the first and second principal components though all neurons encoded some components. *D*, Single cell example. Starting from the predictive power of the history model, principal components were added in their order of significance for the given cell. The predictive power of the full model was achieved by adding just the second component. *E*, Population results. Eighty-one percent of cells required 3 or fewer principal components to reach the full model's predictive power.

against a median of 0.5, $p < 0.001$ in all cases) indicating the restricted models did generalize to a large extent.

Discussion
We found that MI neurons in the hand/arm area encode joint trajectories that describe prehensile movements. Though a large percentage of cells encoded the angular velocities of the flexion/extension of the wrist and proximal phalanx of the index finger, at least 10 of the features were required to account for all the predictive power of the kinematics in at least half the tested population. By considering the principal components on the movement trajectories as features, we were able to account for all the predictive power of the majority of neurons with just up to three features. This indicates that neurons preferentially encode coordinated movements that describe a grasping movement, rather than a movement about any one joint.

Task-related neurons
Intracortical microstimulation studies suggest that MI has a horseshoe-like somatotopic organization, where the center is located in the anterior bank of the central sulcus and evokes movements of the hand and fingers while the periphery evokes movements of the arm (Park et al., 2001). We placed the array in
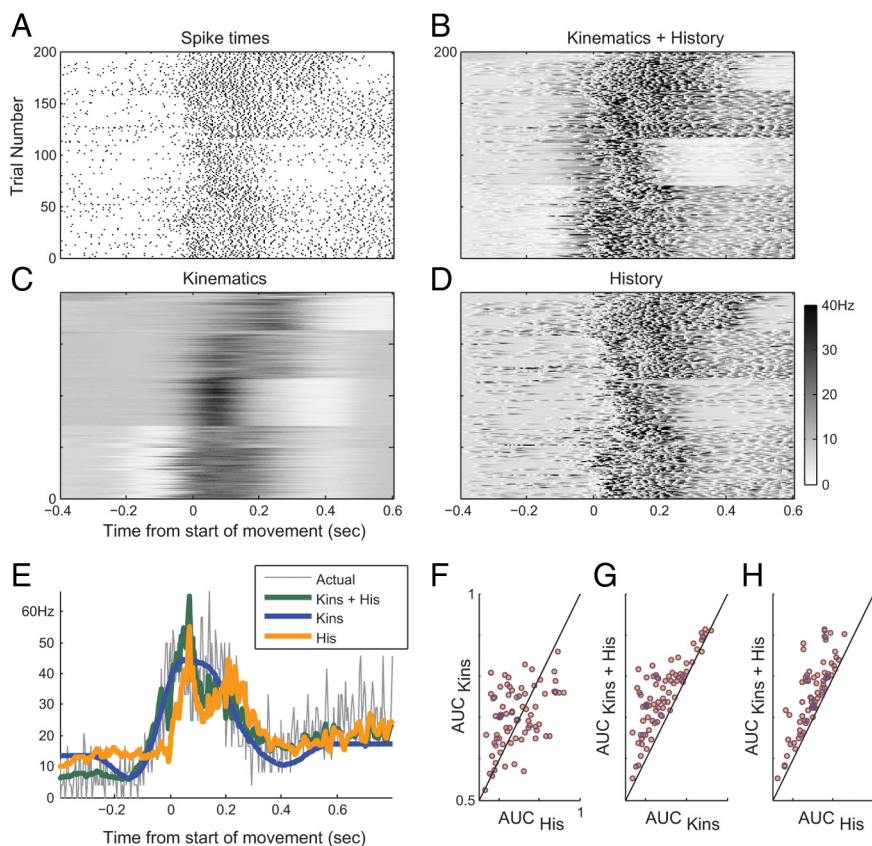
Figure 8. Kinematics vs spike history. **A**, Spike trains aligned on the start of movement across all trials and all four object conditions for a given cell. **B**, Predicted firing rates using the full model where the grayscale ranges between 0 and 40 Hz. **C**, Predicted firing rates using an encoding model that only included the kinematic terms. **D**, Predicted firing rates using an encoding model that only included the spike history terms. **E**, Comparison of the trial-averaged firing rates during grasps to the key grip object for the actual data (gray), the full model (green), the kinematic model (blue) and the history model (yellow). **F**, Comparison of the predictive power for all cells (AUCs) for the kinematic versus history models. **G**, Comparison of the predictive power for all cells (AUCs) for the full model versus the kinematic model. **H**, Comparison of the predictive power for all cells for the full versus history models. Each red circle represents a single motor cortical cell.

Spike history terms

Neurons in MI exhibit phasic and tonic firing rates that reflect fast and slow changes in their biophysical properties (Cheney et al., 1985; Kalaska et al., 1989). The spike history terms were used to model temporal correlations such as these. When we compared the predictive power of models that were built with data from the premovement period only, we observed that the number of cells that encoded the long time-scale effects dropped by at least 45%. Two possible interpretations are that these low-frequency spike history terms capture some aspect of the kinematics that is not modeled by the kinematic parameters or that they capture slow changes in the biophysical properties of cells that only occur during movement.

Time lag analysis

Assuming that neurons encode kinematics at a single time lag, the most significant lag time across neurons was found to be at 148 ms or 200 ms for all kinematic features, consistent with the lag times observed in reaching studies (Moran and Schwartz, 1999a; Paninski et al., 2004b). This study shows that that the predictive power of neurons was maximal for trajectory lengths of 364 ms, compared with shorter lengths including the kinematics at just their significant lag. Although the predictive power of the model increased significantly for a temporally extensive trajectory, the improvement was somewhat modest. Nevertheless, there was a systematic improvement in predictive power as the trajectory length was increased suggesting that grasp trajectory information is contained in the responses of MI neurons (Fig. 5C). Previous studies of reaching show that neurons shift their directional tuning in time during a movement and, as such, suggest that neurons encode trajectories instead of static kinematic parameters at their optimal lag (Hocherman and Wise, 1991; Hatsopoulos et al., 2007).

the intermediate zone, an area between the center and peripheral zones that evokes hand and arm movements (Fig. 1C). It is possible that the 39% of neurons that did not modulate with the grasp task were reach-related. In fact, we recorded a reach-to-grasp task in the same session and found that more neurons modulated with the task (results not reported), presumably due to arm movements. This is consistent with studies that show that neurons in this area modulate their firing rate with arm movements (Georgopoulos et al., 1986; Moran and Schwartz, 1999b; Stark et al., 2007; Truccolo et al., 2008).

Grasping kinematics and relationship to unconstrained reach-to-grasps

Our experimental task aimed to isolate neural responses that are tied to grasping movements. Though the arm was constrained in the brace, the kinematic trajectories of the hand closely match those reported in studies that examined unconstrained reach-to-grasp movements (Fig. 2E) (Mason et al., 2001; Santello et al., 2002). As the kinematic trajectories in Figure 2E depict, a prehensile movement begins with a simultaneous extension and abduction of the fingers including the thumb. The degree to which the hand opens is related to the size of the object (Paulignan et al., 1991). As the hand approaches the object, all the fingers flex and adduct to grasp the object.

Single versus multiple features

Previous cortical encoding studies suggested that single neurons encode more than one individuated finger movement (Schieber, 1990; Schieber and Hibbard, 1993). If we define a finger movement as an angular position and velocity trajectory of a joint, we found similar results: the predictive power of the kinematics was only fully described by 10 kinematic trajectories in a majority of neurons. Despite this multifeature encoding, neurons encoded joint angular velocities more frequently than joint angular positions and preferentially encoded the velocity trajectories of the wrist and the metacarpophalangeal joint of the index finger compared with other trajectories.

Psychophysics studies suggest that the CNS controls prehension by means of two synergies, represented by the first two principal components on the joint angle kinematics of the fingers and wrist (Mason et al., 2001; Santello et al., 2002). Santello et al. (2002) describe the first principal component as an extension of
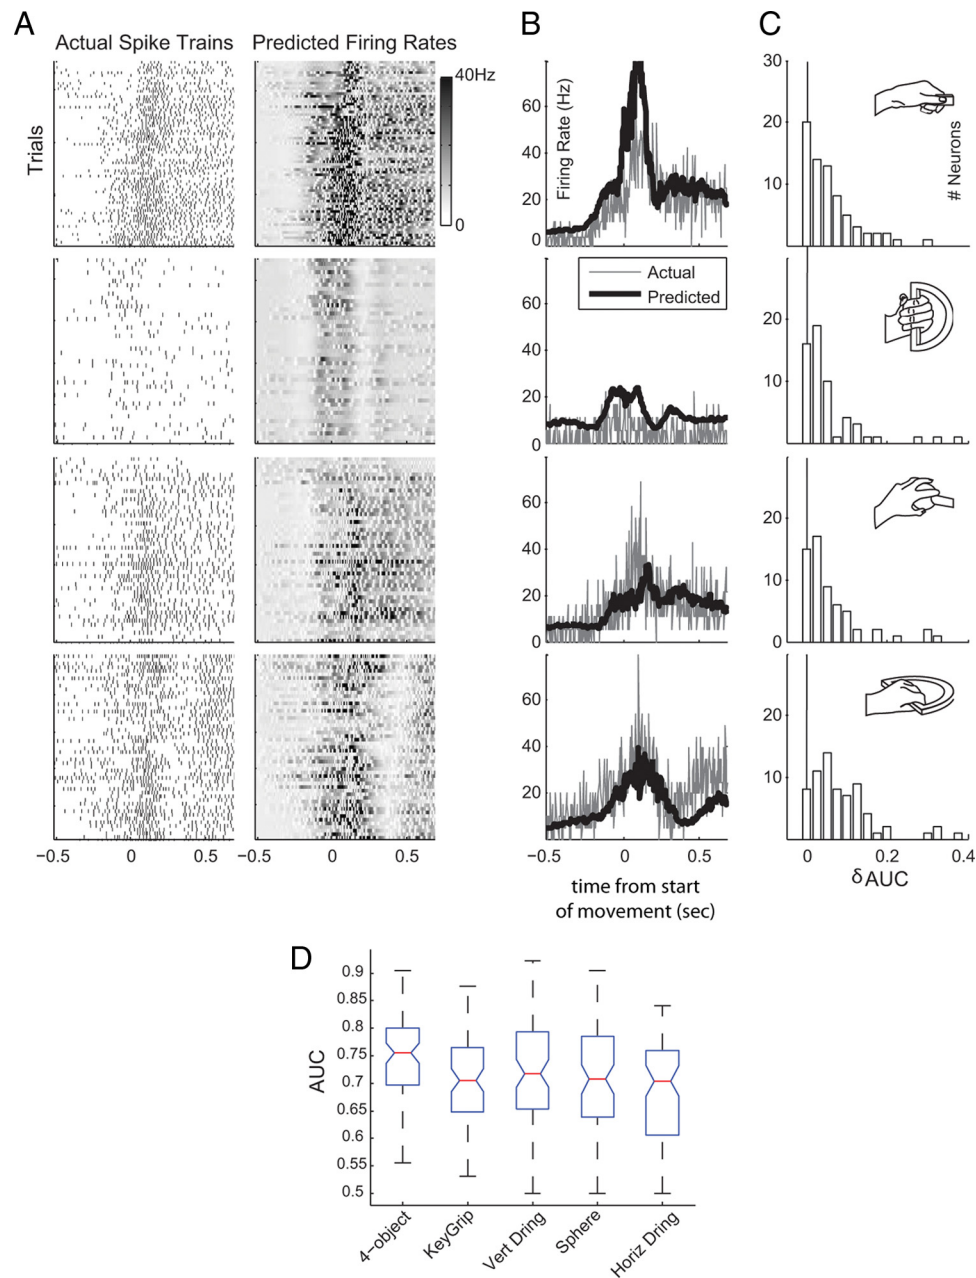
Figure 9. Generalization tests. The restricted model was trained on data from three objects and used to predict data collected from the fourth object where the fourth object is displayed on the upper right of column C. ***A***, Actual spike times and predicted firing rates for each restricted model. ***B***, Comparison of mean firing rates of actual and predicted data. ***C***, Distribution of differences in predictive power (AUC) across neurons between the 4-object model and the restricted model. A difference of 0 would indicate that there was no difference between the predictive powers obtained from the 4-object model and the given restricted model. ***D***, Box plot of predictive power (AUC) of 4-object model and restricted models. The medians are illustrated in red, the lower and upper quartiles are given by the bottom and top edges of the blue boxes and the full range of predictive powers are given by the dotted lines in black.

the metacarpophalangeal and proximal interphalangeal joints to a maximum excursion, followed by flexion of the same joints. At the same time, the digits, are gradually abducted and later adducted toward the end of the reach while the thumb is only abducted. In this paper, we demonstrate that the first two principal components (PCs) were more frequently encoded by the neurons (Fig. 7C) and that a small set of principal components (between 1 and 3) account for the response variability of most MI neurons. Together, these results suggest that each neuron encodes a particular spatiotemporal trajectory that involved the coordinated kinematics of all the joint angles of the hand.

Spike history versus kinematics

This study extends previous encoding models of the hand by considering the continuous hand kinematics as well as the spiking history to predict the firing activity of a neuron on a fine time scale. We found that the spike history terms account for the fast and slow changes in firing rate whereas the kinematics only account for the slower changes. This is understandable given the relatively slow fluctuations of the kinematics. Since the overall modulation onset of an MI neuron is intimately linked to movement onset of the hand, it follows that the kinematics tended to more accurately model the onset of firing rate modulation. The

spike history predicted the fine and coarse temporal structure of the firing activity but predicted modulation onset with a time lag of 28 ms. Spike history cannot predict modulation onset accurately because it carries only past spiking information. The degree to which the fine temporal structure of the neuron could be modeled is limited by the cutoff frequencies of the spike history bases, which are a form of low-pass filter. The first spike history term could capture spike frequencies as high as 60 Hz. Any faster changes in spike frequency could not be captured by the spike history terms since the other bases capture modulations at lower frequencies. For the majority of neurons, a full model yielded significantly higher predictive power than a model with either spike history or kinematics terms alone. This result indicates that spike history and kinematic terms hold non-redundant information about the firing rate of a neuron and can be used together to predict the firing rate of neurons at a fine time scale.

Generalization

We tested whether building an encoding model from the data from three objects could generalize to predict the firing rate collected when the monkey grasped a fourth object. The predictions from these restricted models were well above chance. Compared with a model that was built on data from all four objects, all four restricted models yielded significantly lower prediction accuracy (Fig. 9*C*,*D*). This slight decrease in performance is likely related to the fact that each object elicited some unique aspect of movement. This study was limited in the number of objects that were used in the analysis to get repeated trials of the same object condition. However, although we used a small number of objects, at least the first principal component closely matched the one reported in a human psychophysics study where 20 objects were tested.

Conclusion

MI, as a major output pathway to the spinal cord, is thought to elaborate the motor plan into a more detailed movement. Few encoding studies have characterized what parameters of the grasping movement are particularly relevant for MI neurons: one study recently showed that neurons in MI respond to grasp dimension and grasp force (Hendrix et al., 2009). Our model focuses on an encoding model of single neurons and accounts for the temporal progression of the grasp movement during preshaping before grasp, an important aspect of prehension (Mason et al., 2001; Santello et al., 2002).

We found that neurons preferentially encode different subsets of principal components, which model the coordinated movements that elicit a prehensile movement. This result suggests that the primary motor cortex encodes parameters that reflect ethologically relevant movements, rather than abstract, Newtonian parameters such as joint angular velocity. Though it is generally assumed that simple variables will predict responses to more complex ones, studies in different model systems have shown that using a variable that reflects an ethologically relevant behavior or natural sensory input, greatly sharpens the tuning properties of neurons (Margoliash, 1983; Margoliash and Fortune, 1992; David et al., 2004). With the advent of telemetric neural recording devices and marker-less motion tracking systems, it will be possible to further investigate primate behavior in more ethologically relevant contexts (Obeid et al., 2004).

References

Aggarwal V, Tenore F, Acharya S, Schieber MH, Thakor NV (2009) Cortical decoding of individual finger and wrist kinematics for an upper-limb neuroprosthesis. Conf Proc IEEE Eng Med Biol Soc 1:4535–4538.

Chen D, Fetz EE (2005) Characteristic membrane potential trajectories in primate sensorimotor cortex neurons recorded in vivo. J Neurophysiol 94:2713–2725.

Cheney PD, Fetz EE, Palmer SS (1985) Patterns of facilitation and suppression of antagonist forelimb muscles from motor cortex sites in the awake monkey. J Neurophysiol 53:805–820.

David SV, Vinje WE, Gallant JL (2004) Natural stimulus statistics alter the receptive field structure of v1 neurons. J Neurosci 24:6991–7006.

Dickey AS, Suminski A, Amit Y, Hatsopoulos NG (2009) Single-unit stability using chronically implanted multielectrode arrays. J Neurophysiol 102:1331–1339.

Fawcett T (2006) An introduction to ROC analysis. Pattern Recognit Lett 27:861–874.

Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. Science 233:1416–1419.

Georgopoulos AP, Pellizzer G, Poliakov AV, Schieber MH (1999) Neural coding of finger and wrist movements. J Comput Neurosci 6:279–288.

Hatsopoulos NG, Xu Q, Amit Y (2007) Encoding of movement fragments in the motor cortex. J Neurosci 27:5105–5114.

Hendrix CM, Mason CR, Ebner TJ (2009) Signaling of grasp dimension and grasp force in dorsal premotor cortex and primary motor cortex neurons during reach to grasp in the monkey. J Neurophysiol 102:132–145.

Hille B (2001) Ion channels of excitable membranes. Sunderland, MA: Sinauer Associates.

Hocherman S, Wise SP (1991) Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys. Exp Brain Res 83:285–302.

Kalaska JF, Cohen DA, Hyde ML, Prud'homme M (1989) A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. J Neurosci 9:2080–2102.

Keat J, Reinagel P, Reid RC, Meister M (2001) Predicting every spike: a model for the responses of visual neurons. Neuron 30:803–817.

Margoliash D (1983) Acoustic parameters underlying the responses of song-specific neurons in the white-crowned sparrow. J Neurosci 3:1039–1057.

Margoliash D, Fortune ES (1992) Temporal and harmonic combination-sensitive neurons in the zebra finch's HVc. J Neurosci 12:4309–4326.

Mason CR, Gomez JE, Ebner TJ (2001) Hand synergies during reach-to-grasp. J Neurophysiol 86:2896–2910.

Mason CR, Gomez JE, Ebner TJ (2002) Primary motor cortex neuronal discharge during reach-to-grasp: controlling the hand as a unit. Arch Ital Biol 140:229–236.

Mason CR, Theverapperuma LS, Hendrix CM, Ebner TJ (2004) Monkey hand postural synergies during reach-to-grasp in the absence of vision of the hand and object. J Neurophysiol 91:2826–2837.

Maynard EM, Nordhausen CT, Normann RA (1997) The Utah intracortical electrode array: a recording structure for potential brain-computer interfaces. Electroencephalogr Clin Neurophysiol 102:228–239.

Maynard EM, Hatsopoulos NG, Ojakangas CL, Acuna BD, Sanes JN, Normann RA, Donoghue JP (1999) Neuronal interactions improve cortical population coding of movement direction. J Neurosci 19:8083–8093.

McCullagh P, Nelder JA (1989) Generalized linear models. Boca Raton, FL: Chapman and Hall/CRC.

Moran DW, Schwartz AB (1999a) Motor cortical activity during drawing movements: population representation during spiral tracing. J Neurophysiol 82:2693–2704.

Moran DW, Schwartz AB (1999b) Motor cortical representation of speed and direction during reaching. J Neurophysiol 82:2676–2692.

Obeid I, Nicolelis MA, Wolf PD (2004) A multichannel telemetry system for single unit neural recordings. J Neurosci Methods 133:33–38.

Paninski L, Pillow JW, Simoncelli EP (2004a) Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. Neural Comput 16:2533–2561.

Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP (2004b) Spatiotemporal tuning of motor cortical neurons for hand position and velocity. J Neurophysiol 91:515–532.

Park MC, Belhaj-Saïf A, Gordon M, Cheney PD (2001) Consistent features in the forelimb representation of primary motor cortex in rhesus macaques. J Neurosci 21:2784–2792.

Paulignan Y, Jeannerod M, MacKenzie C, Marteniuk R (1991) Selective per-

turbation of visual input during prehension movements. 2. The effects of changing object size. Exp Brain Res 87:407–420.

Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature 454:995–999.

Poliakov AV, Schieber MH (1999) Limited functional grouping of neurons in the motor cortex hand area during individuated finger movements: a cluster analysis. J Neurophysiol 82:3488–3505.

Rousche PJ, Normann RA (1992) A method for pneumatically inserting an array of penetrating electrodes into cortical tissue. Ann Biomed Eng 20:413–422.

Santello M, Flanders M, Soechting JF (2002) Patterns of hand motion during grasping and the influence of sensory guidance. J Neurosci 22:1426–1435.

Schieber MH (1990) How might the motor cortex individuate movements? Trends Neurosci 13:440–445.

Schieber MH, Hibbard LS (1993) How somatotopic is the motor cortex hand area? Science 261:489–492.

Stark E, Drori R, Asher I, Ben-Shaul Y, Abeles M (2007) Distinct movement parameters are represented by different neurons in the motor cortex. Eur J Neurosci 26:1055–1066.

Stevenson IH, Rebesco JM, Hatsopoulos NG, Haga Z, Miller LE, Körding KP (2009) Bayesian inference of functional connectivity and network structure from spikes. IEEE Trans Neural Syst Rehabil Eng 17:203–213.

Todorov E, Ghahramani Z (2004) Analysis of the synergies underlying complex hand manipulation. Conf Proc IEEE Eng Med Biol Soc 6:4637–4640.

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

Truccolo W, Friehs GM, Donoghue JP, Hochberg LR (2008) Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. J Neurosci 28:1163–1178.

Truccolo W, Hochberg LR, Donoghue JP (2010) Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. Nat Neurosci 13:105–111.

Umilta MA, Brochier T, Spinks RL, Lemon RN (2007) Simultaneous recording of macaque premotor and primary motor cortex neuronal populations reveals different functional contributions to visuomotor grasp. J Neurophysiol 98:488–501.

Vargas-Irwin CE, Shakhnarovich G, Yadollahpour P, Mislow JM, Black MJ, Donoghue JP (2010) Decoding complete reach and grasp actions from local primary motor cortex populations. J Neurosci 30:9659–9669.