Table S1

		Exemplar		Decision	
		mean	SEM	mean	SEM
Observed data PLV	Category Performance	0.4663	0.0038	0.4634	0.0033
	Category Learning	0.4656	0.0033	0.4697	0.0033
	SR Learning	0.4571	0.0031	0.4506	0.0029
Surrogate data PLV	Category Performance	0.3925	0.0007	0.3937	0.0007
	Category Learning	0.3913	0.0005	0.3977	0.0006
	SR Learning	0.3885	0.0005	0.3964	0.0007

Table S1 – Synchrony in observed and surrogate data. Mean PLV and the corresponding SEM of all PFC-STR pairs of electrodes as a function of trial epoch and experimental stage, in observed, uncorrected data (top 3 rows) and in surrogate data (bottom 3 rows). Surrogate data were generated by randomly shuffling the trials 200 times prior to computation of PLV. The surrogate-data PLV was an estimate of the bias in the PLV measure, and, to generate the main paper's results, the surrogate-data PLV was subtracted from the observed PLV. This correction was implemented at each trial.

Supplemental Experimental Procedures

Animals

Data were collected from 2 adult female rhesus monkeys (*Macaca mulatta*), 5-9 kg. The animals were taken care of in accordance with the National Institutes of Health guidelines and the policies of the Massachusetts Institute of Technology Committee for Animal Care. We had spent approx. 1.5 years training the animals on this task and exploring various different experimental designs before we started collecting neurophysiological data using the current design (see below). Both animals were trained on the category learning task until they reached similar levels of proficiency.

Task Design

The task design is shown in Fig. 1 and has been described in detail in our previous report from this dataset (Antzoulatos and Miller, 2011). Experimental control was implemented via Cortex (NIMH, Laboratory of Neuropsychology), infrared eye-tracking via Eyelink 1000 (SR Research Ltd, Mississauga, Canada) and neurophysiological recordings via the MAP system (Plexon Inc, Dallas TX). Visual stimuli were presented at full contrast, on a CRT monitor (at a distance of approx. 50 cm from the animal), refreshing at 100 Hz. Trials began when the animal maintained fixation on a central target for 0.7 s, following which, a randomly chosen exemplar from either category was presented for 0.6 s. One second after the offset of the exemplar, the fixation target was extinguished and 2 saccade targets appeared at 5^{o} to the left and right of the center of fixation. The animal had to make a single and direct saccade to the correct target within 1s, and maintain fixation on it for 200 ms for reward (drops of juice). In the case of incorrect response (error trials), there was a 5-s timeout, during which the exemplar was presented again, at the location of the correct target.

Category exemplars were static constellations of 7 randomly located dots (0.4^{o} in diameter; Fig. 1), and subtended a 6x6-degree spatial window, centered on the fixation target. For each daily recording session, we started by constructing a new pair of category prototypes. We took certain precautions (Antzoulatos and Miller, 2011; Vogels et al., 2002) to ensure an intermediate level of task difficulty and to increase the likelihood that the animals would learn the new categories in a single, daily, recording session (no categories were tested in more than 1 sessions): Within each prototype, all dots had to be more than 0.8^{o} away from each other (i.e., a distance equal to 2 dots), no more than 3 pairs of dots between prototypes were allowed to be closer than 1^{o}, and the average Euclidean distance between prototypes had to be between 1.6^{o} - 2.2^{o}. The next step was to create the category exemplars. To do so, each dot of the corresponding prototype was shifted to a random direction and distance from its original location. Distance from original location was tiered at 5 levels, each tier being one extra dot-diameter (0.4^{o}) away from the original location. The probability of each dot to land in any one of the 5 tiers was determined based on the originally published levels of distortion 2 and 3 (Antzoulatos and Miller, 2011; Posner et al., 1967). This led the majority of dots (63%) to shift by only 1 dot away (i.e. to tier 1). However, because the direction of shift could vary randomly by 360^{o} and independently for each of the 7 dots comprising the exemplar, the overall difference of an exemplar from the original prototype was substantial. No dot was allowed to stay in the same location as in the prototype, or at a distance closer than 0.8^{o} from another dot, and no more than 2 pairs of dots across exemplars could be closer than 0.5^{o} (approx. 1 dot-diameter).

On each trial, a randomly chosen exemplar was tested, with trials from both categories randomly interleaved throughout the recording session. The animals advanced through a minimum of 8 blocks, each of which included twice as many exemplars as the block before it (block 1 included 1 exemplar per category). Each block was complete when the animals' performance reached 80% correct at the last 20 trials (all analyses were performed on the minimum of 16 correct trials per block). No exemplar was allowed to be tested in more than 2 consecutive blocks. The behavioral performance criterion we employed led some blocks to be terminated before all exemplars had been tested. It was necessary to avoid prolonging a block too much, so the animal would not form habitual SR associations that would impede abstraction of the essence of each category(Antzoulatos and Miller, 2011). Distinction of learning in 3 stages relied on the same criteria we previously employed (Antzoulatos and Miller, 2011): The first 2 blocks in each experiment were assigned to learning stage 1 (SR Learning) because they only tested 1-2 exemplars per category, whose frequent repetitions allowed for learning of individual SR associations. All blocks after the 2nd, and until a category performance criterion was met (see below), were classified as learning stage 2 (Category Learning). At this stage, the likelihood of any single exemplar to be repeated was gradually diminishing, as more and more exemplars were introduced in each block. We considered category learning to be complete when the animals could classify correctly the majority (75%) of each category's exemplars on their first trial. The first 2 blocks meeting this criterion were classified as stage 3 (Category Performance), and asymptotic performance at this stage relied almost exclusively on a single trial per exemplar. The neurophysiological results were averaged across all blocks of each stage.

Neural Recordings and Data Analysis

Guided by the animals' structural MRI images, simultaneous multi-electrode recordings were made from the right lateral prefrontal cortex (PFC; dorsal and ventral regions) and the head and body of the right caudate nucleus of the striatum (STR; see Antzoulatos and Miller, 2011 for precise anatomical locations). Two custom-made multi-electrode arrays (8-16 tungsten electrodes, FHC) were lowered at different sites in the animals' brain every day of recording. Because the electrodes could be guided either individually or in pairs, apart from varying their anteroposterior and mediolateral coordinates, we could also vary their exact depth so as to maximize the yield of neural signals. Electrode recordings were first fed to a unity-gain headstage and were referenced to ground. Local field potentials (LFPs) were separated from spiking signals online at the preamplifier, using a 0.7Hz - 300Hz bandpass filter, were amplified x1000, and then sampled at 1 KHz rate. To ensure that only signals from active regions of the brain were collected, LFPs were recorded only from sites that also displayed spiking activity (a total of 84 PFC sites and 65 STR sites). Line frequency noise (60 Hz) was removed from the signal offline, by applying a 10th-order Butterworth bandstop filter (59Hz - 61Hz) at both forward and backward time directions. To remove any stimulus-evoked potentials added on top of the ongoing oscillations, prior to wavelet transform, all time-aligned LFP signals were mean centered, that is, the cross-trial average time-locked LFP signal (separate by category and outcome) was subtracted from each trial's LFP. All data analyses were performed on MATLAB (Mathworks, Natick MA), and statistical tests were corrected for multiple comparisons.

Spectral decomposition of LFPs relied on a wavelet transform by convolution of the LFP signal with a Morlet wavelet (Torrence and Compo, 1998), at 5 octaves from 2-64 Hz, at a frequency resolution of 0.1 octave. Wavelet analyses were conducted based on the MATLAB-based software Wavelet, offered by C. Torrence and G. Compo at the URL: http://atoc.colorado.edu/research/wavelets/.

After wavelet transform, the frequency-specific phase of each wave (inverse tangent of real and imaginary components) was extracted using the MATLAB function *angle* and the amplitude (from Pythagorean equation with the real and imaginary components as x and y) using the function *abs*. Spectral power was computed as the squared amplitude, and normalized to frequency^{-1} to correct for the power-law decay and enhance visibility of the higher frequencies (Siegel et al., 2009). Synchrony between pairs of LFP signals was evaluated as a Phase-Locking Value (PLV):

$$PLV = \left| \frac{1}{n} \sum_{1}^{n} e^{i(\varphi_1 - \varphi_2)} \right|$$

which is the length of the vector average of a sample (n) of phase-differences (φ_1 - φ_2). The phase (φ) is a function of frequency and time. Being a circular mean, PLV = 1-circular variance. As such, it varies between 0 (maximum variance), when all phase-differences are uniformly distributed over 360^{o}, and 1 (minimum variance), when all phase-differences are a single value. PLV quantifies, therefore, how consistent the phase-difference is between 2 waves over a set of observations. Averaging of phase-differences can be computed over a time segment, over a sample of trials, or across several pairs of electrodes. Each of these methods has been used in the past, depending on the specific question of interest and other factors (e.g., Lachaux et al., 2000; Wang et al., 2006). Our analyses utilized the first approach: For each of the frequencies of interest, the momentary (at 1-ms resolution) phase-difference was averaged over a 500-ms long time window. This analysis has the advantages of, first, utilizing a large enough sample of phase-differences, and second, preserving the trial resolution, which was useful for further computation of category selectivity (see below). The autoregressive nature of the LFP can bias the PLV: Imagine 2 perfect sinusoids simultaneously recorded over a period of time. Regardless of the size of their phase-difference, this difference will be constant and the sinusoids will appear perfectly phase-locked (PLV=1). In reality, because of multiple momentary shifts in phase/amplitude, LFPs are far from perfect sinusoids and their PLV never reaches its maximum. Still, in order to correct PLV for this bias, the trials were randomly shuffled 200 times and the average randomization PLV (i.e., the PLV expected by chance) was computed. The observed PLV was always greater than the randomization PLV, indicating that the bias was superimposed on the true synchrony between a pair of LFP signals. All PLVs we used in our analyses were bias-corrected by subtracting this bias from the observed PLV.

Synchrony was evaluated from simultaneously recorded pairs of electrodes in PFC and STR (PFC-STR, n=426), within PFC (PFC-PFC, n=240) or within STR (STR-STR, n=141). Other than PLV, the results were similar when we evaluated synchrony across trials using the standard coherence measure (e.g., Buschman et al., 2012), or pairwise phase consistency (Vinck et al., 2010). Synchrony between proximal sites (i.e., within PFC or STR) may also be vulnerable to electrotonic volume conduction. By also evaluating synchrony separately for proximal vs. distal pairs of electrodes, we confirmed that the lack of learning-induced changes in PLV within PFC and STR (Fig.3) was not due to masking by volume conduction: Neither proximal, not distal electrode pairs in PFC or STR displayed a change in synchrony across learning stages. Finally, because of the multidimensionality of a dataset like the one in this study (with dimensions like time, trials, learning stages, spectral frequency, area of recording, etc.), to increase the statistical power and computational efficiency of the analyses, it was necessary to only focus at the most critical dimensions. However, we ensured that the dimensions we chose to collapse (e.g., the trial dimension by averaging the first 16 correct trials and the time dimension by looking at 2 critical trial epochs) were inconsequential for the

results of our analyses. The results were, therefore, similar when we examined different trial epochs, or when we evaluated the evolution of synchrony across trials within a learning stage.

Similar to the LFP-LFP PLV (above) is the computation of the spike-LFP PLV:

$$PLV = \left| \frac{1}{n} \sum_{1}^{n} e^{i\varphi} \right|$$

except that n is now a sample of spike timestamps, and φ is the instantaneous frequency-specific phase at the time of each spike. This form of PLV evaluates how consistently spikes are fired at specific phase-bins of each frequency, and has two caveats: First, it is sensitive to the number of spikes, and second, it is sensitive to the temporal profile of spiking activity (see below). To partly address the former, these analyses are typically performed on multiunit activity (MUA) instead of single-unit activity. We also did the same: we pooled all previously sorted spikes into a single MUA signal per electrode (Antzoulatos and Miller, 2011). Not only did this increase the size of spike samples, but also ensured that each electrode provided a single LFP and a single MUA signal. Another approach that is sometimes followed to address this sensitivity on spike numbers is stratification: It involves keeping the minimum number of spikes from all trials, and discarding the (randomly chosen) excess spikes. This approach was not straightforward to implement in our dataset, because the number of spikes was considerably fluctuating in the course of learning. Additionally, even stratification does not fully address the second caveat of spike-LFP PLV, which is the dependence on the temporal profile of spikes: Consider a burst of spikes occurring over a brief (say, 50-ms) time window after a visual display. Inevitably, due to their temporal clustering, these spikes will appear phase-locked at a low frequency (because at low frequencies each phase-bin lasts longer than at high frequencies). Even after stratification the spikes will most likely remain in that brief time window, thus leading to spurious PLV. Therefore, in order to correct for these biases of the MUA-LFP PLV, we randomly shuffled the trials 200 times, and subtracted the average randomization PLV, as we did for LFP-LFP PLV (above). Although this analysis did not yield evidence of general spike-LFP synchrony, evaluation of category selectivity (as described below) revealed that spike-LFP synchrony was actually category-specific (see Figs 5 and S2).

As in our previous report of category-selective spiking activity (Antzoulatos and Miller, 2011), for category selectivity in LFP-LFP and MUA-LFP synchrony (PLV) we used the discrimination index d'. This index is calculated as the absolute difference between mean synchrony in the 2 sets of trials (category A vs. B), normalized to their pooled standard deviation s_p, i.e.,

$$d' = \frac{|\langle PLV_A \rangle - \langle PLV_B \rangle|}{s_p}$$

where

$$s_p = \sqrt{\frac{s_A^2 * (n_A - 1) + s_B^2 * (n_B - 1)}{(n_A + n_B - 2)}}$$

s^2 being the variance and n the number of each set of trials. To correct for biases generated by the variable number of trials in each category, any given set of trials was randomly shuffled between the 2 groups 200 times, thus creating a surrogate dataset. The observed d' was subsequently transformed into a z score on the basis of the surrogate data mean and variance. Similarly, we

expressed the difference between error and correct trials as a z-transformed d', except that we computed the signed (i.e., not absolute) difference of error-correct trial mean PLV (Fig.4).

Finally, our Granger causality analyses (Fig.6) relied on a nonparametric spectral matrix factorization algorithm (Dhamala et al., 2008a, 2008b), which can be found in the MATLAB-based software Fieldtrip (Oostenveld et al., 2011). Preliminary analyses with parametric Granger causality tests indicated that the optimal multivariate autoregressive (MVAR) model order (estimated using the Bayesian and Akaike criteria) was highly variable. However, because our finding that STR exerts stronger net influence on PFC was somewhat surprising (given the direct inputs from PFC to STR), we did replicate this finding using the parametric Granger test (with MVAR orders from 5 to 100) in both the time and frequency domains. These analyses were performed on the MATLAB-based software GCCA (Seth, 2010).

SUPPLEMENTAL REFERENCES:

Dhamala, M., Rangarajan, G., and Ding, M. (2008b). Estimating Granger Causality from Fourier and Wavelet Transforms of Time Series Data. Phys. Rev. Lett. *100*.

Lachaux, J.-P., Rodriguez, E., Le Van Quyen, M., Lutz, A., Martinerie, J., and Varela, F.J. (2000). Studying Single Trials of Phase Synchronous Activity in the Brain. Int. J. Bifurc. Chaos *10*, 2429–2439.

Oostenveld, R., Fries, P., Maris, E., and Schoffelen, J.-M. (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. Comput. Intell. Neurosci. *2011*, 1–9.

Posner, M.I., Goldsmith, R., and Welton, K.E., Jr (1967). Perceived distance and the classification of distorted patterns. J. Exp. Psychol. *73*, 28–38.

Seth, A.K. (2010). A MATLAB toolbox for Granger causal connectivity analysis. J. Neurosci. Methods *186*, 262–273.

Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., and Pennartz, C.M.A. (2010). The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization. NeuroImage *51*, 112–122.

Wang, Y., Hong, B., Gao, X., and Gao, S. (2006). Phase Synchrony Measurement in Motor Cortex for Classifying Single-trial EEG during Motor Imagery. (IEEE), pp. 75–78.