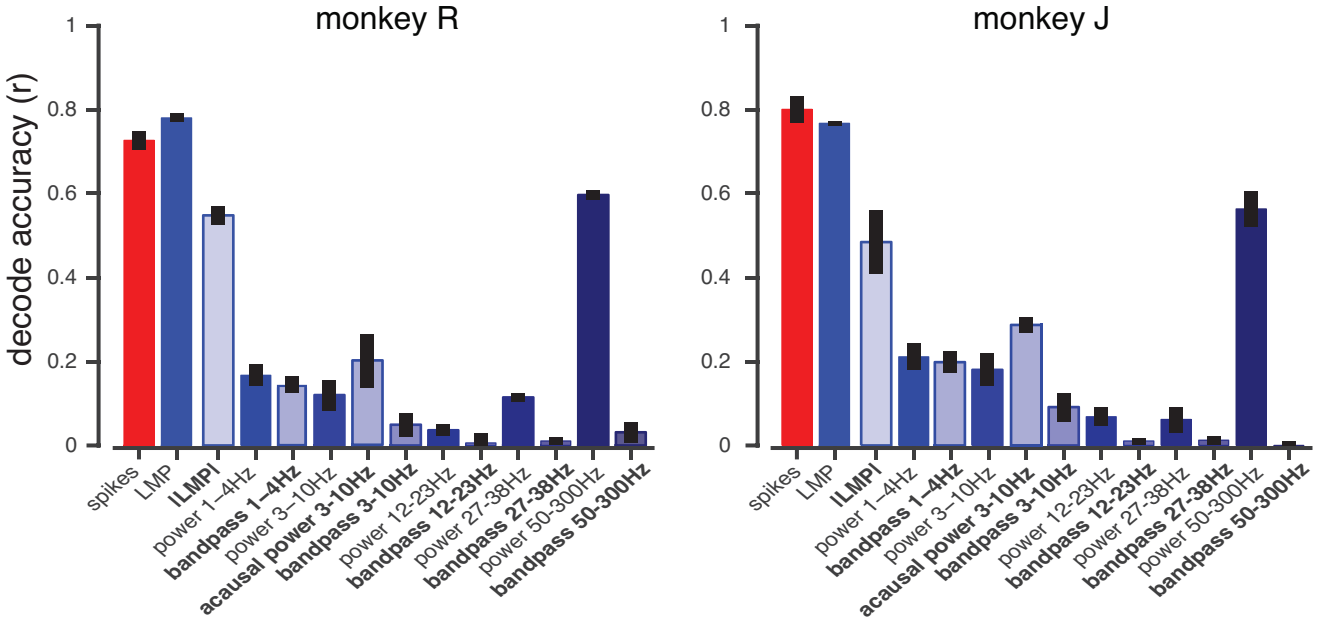


**Supplementary Figure 1**. Offline decode analysis of additional candidate LFP features. Solid bars are the same features as in Figure 2(a), plotted here for comparison; outlined bars with bolded labels show offline decode accuracy using additional features. The offline decode was performed using the same data and methods as in Figure 2. None of these additional features performed as well as the LMP feature. "|LMP|" is the absolute value of the LMP signal; note that this is not the same as the half-wave rectified LMP, which had similar offline decode accuracy to the LMP results shown here. The "acausal power 3-10Hz" feature was computed by filtering the signal both forwards and backwards in time to yield a zero phase delay filter. While this improved offline decode accuracy, we note that this acausal filter would be impractical for closed-loop decoding because of the required delay introduced into the control loop. We also evaluated the other frequency bands' power after filtering in this acausal manner (including the LMP), and found that these features' offline decode accuracy (not shown) were very similar to their causally filtered counterparts. The "bandpass" features are the time-domain signal amplitude (sampled every 50 ms) after filtering the LFP using the same filter as the corresponding band power feature.

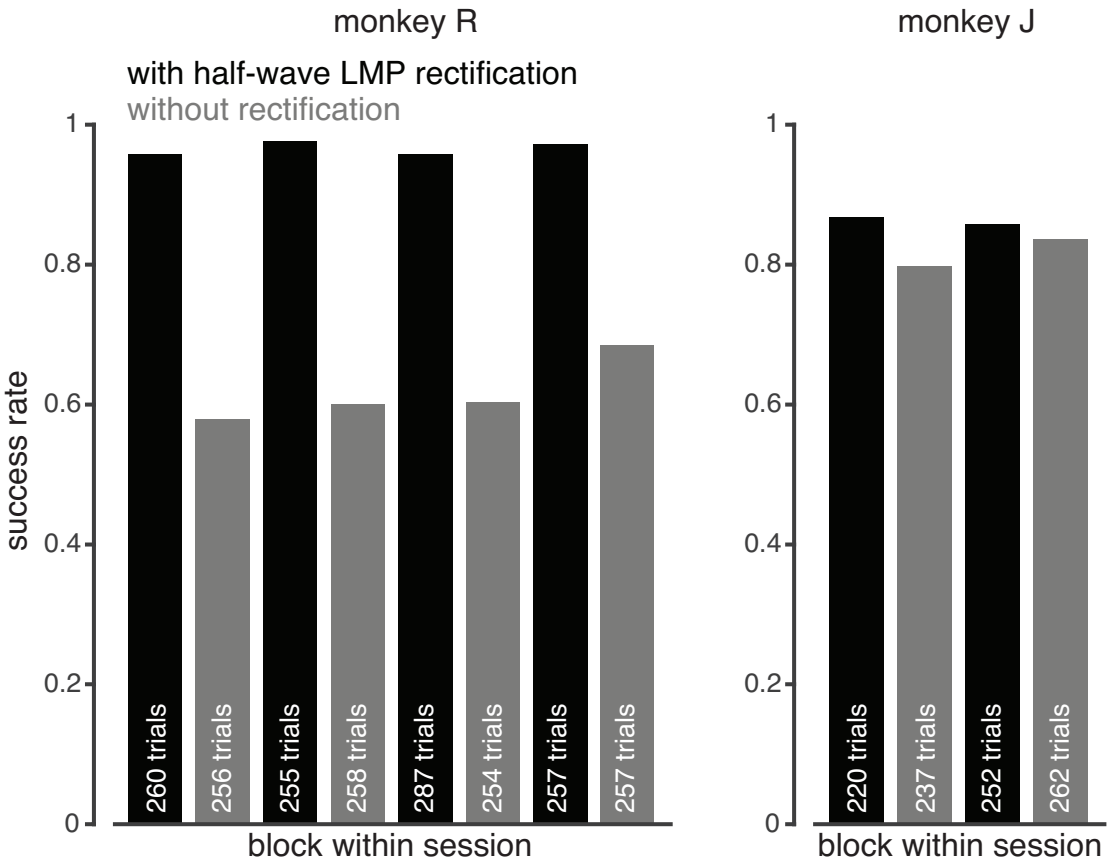


**Supplementary Figure 2**. Closed-loop performance of the LMP decoder with and without half-wave rectification. We evaluated performance on the Continuous Random Target Task using LMP decoders either with the half-wave rectification step described in Methods (black) or without this step (grey). When used, half-wave rectification was applied both to the training data and during closed-loop use. The two decoders were compared in alternating blocks within the same experimental session. Success rates for each block are shown in chronological order within an experiment session. The number of trials in each block is labeled. Monkey R's overall success rate was significantly higher with half-wave rectification ($p < 0.001$, binomial test), but monkey J's was not ($p > 0.05$).

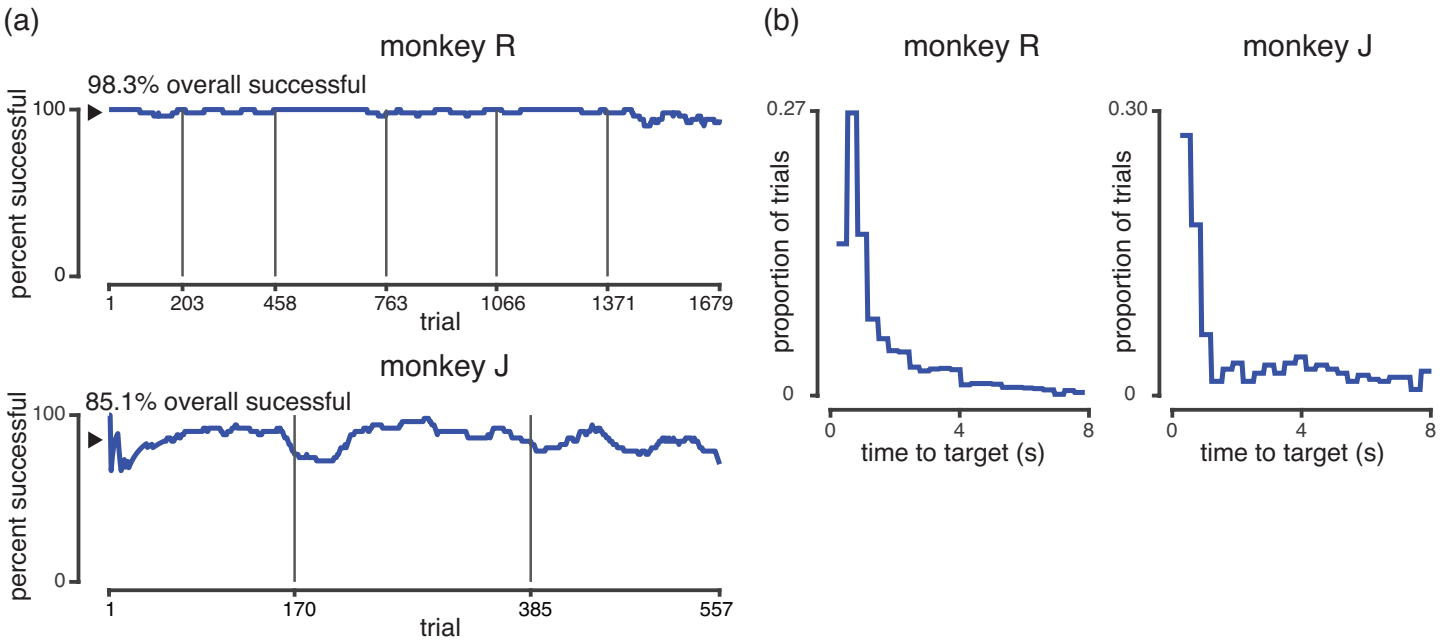


**Supplementary Figure 3**. Performance of the LMP decoder on a Radial 8 Task. Targets were pseudorandomly positioned at one of 8 locations evenly spaced along a 13 cm diameter circle. The monkey initiated the trial by first acquiring a central target, after which the target jumped to a peripheral location. To successfully acquire the target, the cursor had to be moved within an acquisition region of width 3.4 cm and held there for a contiguous 300 ms. (a) Running mean of success rate for the previous 50 trials across 6 days for monkey R and 3 days for monkey J. Vertical grey lines mark transitions between experiment days. Overall percent successful across all experimental sessions' trials is shown with a black arrow. (b) Histograms showing the distribution of times to target, also known as reach times, for all trials of this task. The 300 ms hold time is not included in this calculation of time to target.

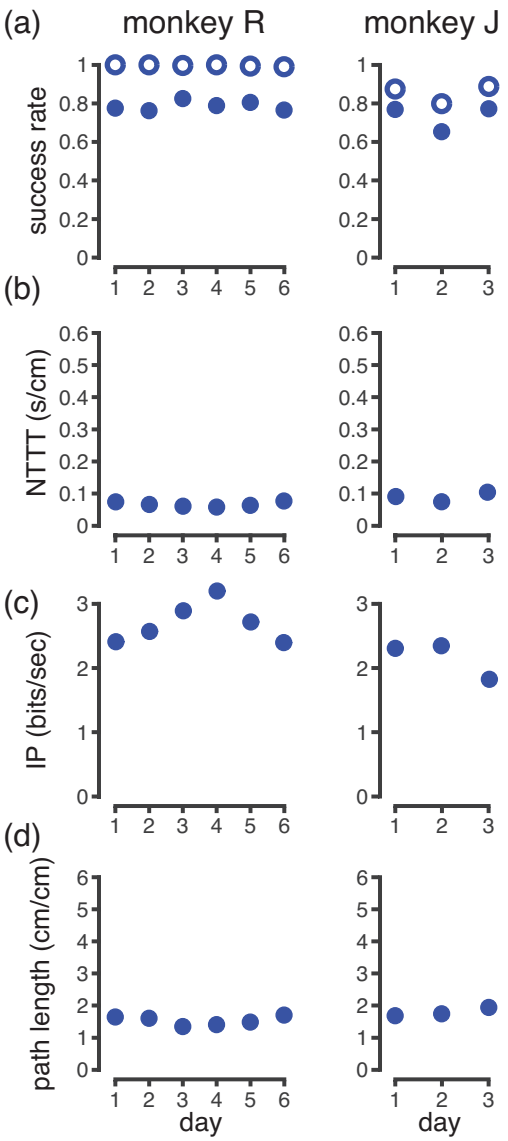


**Supplementary Figure 4**. Performance on a Continuous Random Target Task with short hold time. Targets randomly appeared anywhere in a 20 x 20 cm region. The monkey had up to 10 seconds to move the cursor into a 4 x 4 cm target area and hold it there for 100 ms. To be consistent with the metrics used in (Flint et al. 2013), we restricted our definition of a successful trial to those where the cursor stayed within the target for 100 ms the first time they entered the target region; that is, trials with 'dial-in' (leaving the target within 100 ms and then reacquiring it) were counted as failures. Note that this is a posthoc restriction of what constitutes a successful trial, and that our monkeys were trained to expect that dial-in is allowed and performed the task under this reward contingency. (a) Mean success rate for each day when this task was run is shown in filled circles. Open circles shown the success rate had we counted trials where the target was acquired after dial-in as successful. (b) Normalized time to target. (c) IP = Index of Performance, also known as Fitts bits per second. (d) Path length ratio.

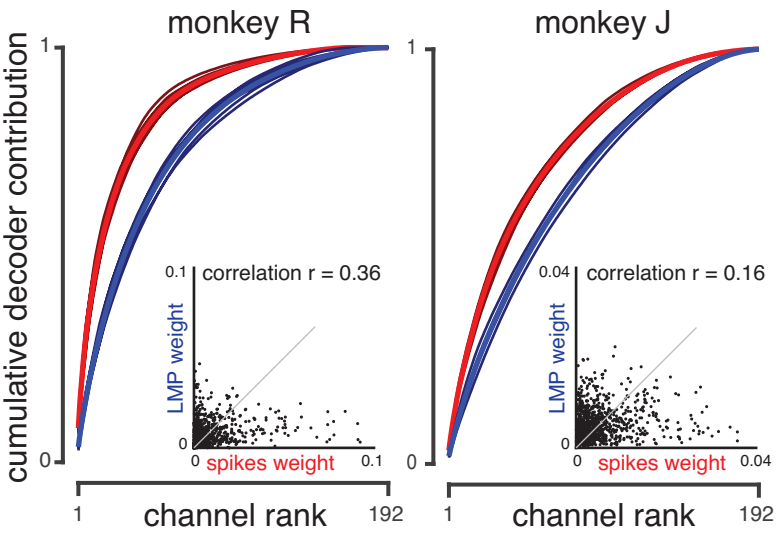


**Supplementary Figure 5**. Cumulative channel contribution to the spikes-only and LMP-only decoders. The contribution of each channel to a decoder is calculated based on that channel's weight in the steady-state velocity Kalman filter multiplied by the average magnitude of the signal from this channel in the training data set. Channels are ranked from greatest to least contribution, and the cumulative decoder contribution is plotted as a function of channel rank, starting from the rank 1 (most contributing) channel. We show the cumulative channel contribution curves for spikes-only (dark red) and LMP-only (dark blue) decoders built from the arm reaching training data of each experiment day in which a full-channel hybrid decoder was evaluated (8 datasets for monkey R, 5 datasets for monkey J). Bright red and bright blue lines show the means of these individual days' contribution curves. On average, half of monkey R's spikes decoder weight was contributed by the top 13 channels, compared to 30 channels for the LMP decoder. For monkey J, on average 31 channels contributed half the spikes decoder weight, compared to 50 LMP channels. Inset scatter plots show the fractional contribution of each channel to the spikes decoder and LMP decoders; each point corresponds to one channel from one session. The mean Pearson's correlation between channels' fractional contribution to the spikes and LMP decoders is summarized in the inset title. This figure shows that LMP decoder weights were more distributed across channels than spikes decoder weights, and that a given channel's importance in the LMP decoder was poorly correlated with its importance in the spikes decoder.

**Supplementary Movie 1**: Example of LFP-only BMI control. This video shows one minute of continuous LMP-driven cursor control while monkey R performs the Continuous Random Target Task with 500 ms target hold time. The monkey controls the white cursor and attempts to acquire the green target. The target turns blue when the cursor is within the 5 x 5 cm target acquisition area. The trials shown are from the second half of the first experiment day shown in Figure 4. This video was re-rendered from recorded experiment data and plays in real time. For the benefit of the viewer, text has been superimposed during re-rendering to show elapsed time, number of trials, and success rate since the start of the movie. Dataset R.2013.09.09.

**Supplementary Movie 2:** Example of spikes-only and hybrid (spikes + LMP) BMI control with both LMP and spikes disabled on 60 channels. One minute of neural cursor control using each decoder are shown side by side to illustrate the benefit of hybrid decoding when channels with good spikes information have been disabled. All trials shown come from the same day in which the hybrid and spikes-only decoders were compared in an interleaved block design. These trials are part of the data used to generate one of the 60 channels removed data points in Figure 8. Dataset R.2013.09.29.