Supporting Information

Hauschild et al. 10.1073/pnas.1215092109

SI Results

Sensitivity of the Decoding Performance to Visual Stimuli. We performed an additional ridge decode analysis with all visual cue onset intervals (0–220 ms after target appearance) removed and compared the result to the original ridge decoding results obtained without cue onset removal. Cue-onset removal shortened the data set by 16%/20% (monkey R/G) but without eliminating much of the hand trajectory toward the target because the reaction time (the duration from target appearance to movement onset) is known to be significantly longer than 220 ms (1). In monkey R, the average decoding performance with visual cue–onset intervals removed was $R^2 = 0.45 \pm 0.05$, and it was $R^2 = 0.45 \pm 0.05$ without cue-onset removal. The difference in R^2 is not significant (two-sided sign test; $P = 0.46$). In monkey G, the average decoding performance with visual cue–onset intervals removed was $R^2 = 0.38 \pm 0.06$, whereas it was $R^2 = 0.36 \pm 0.06$ with cue onsets not removed. Although the difference is significant in monkey G (two-sided sign test; $P = 1.31 \times 10^{-7}$), the small magnitude of the difference still is negligible. Because the data segments after cue-onset removal are shorter than the original segments before cue onset removal, the two performance results cannot be compared directly. Despite this limitation, the negligible difference found suggests that the cue onset does not have a strong detrimental effect on decoding performance. Because the abrupt appearance of the visual target in an otherwise completely dark workspace represents a stronger visual stimulus than typically occurring in natural environments, it is likely that the control signals extracted from PPC neural populations are immune to changes occurring in the visual workspace under realistic contrast conditions.

Comparing Brain-Control Performance with Motor Cortex Results. Because, to our knowledge, this is the only study that relied on PPC signals for continuous 3D cursor brain-control, it would be instructive to compare the results obtained here with a 3D brain-control study relying on motor cortical populations of neurons (2). The overall success rates reported in the two studies were in the same range. For example, the "coadaptive task" of Taylor et al., with a 2-cm target radius (table 2 in ref. 2), was reasonably close to our experimental design, with similarities in workspace configuration and accuracy requirements. Both tasks were performed in a cubic 3D workspace of comparable dimensions (coadaptive task, cube edge length: 12.25 cm; this study: 10 cm) and required the monkey to move the cursor center to within 3 cm of the target center. Under these conditions, Taylor et al. reported success rates of 86 ± 8% and 47 ± 21% (monkey M, O), whereas we report comparable 89 ± 9% and 54 ± 9% (monkey R, G; mean ± SD) average performance. Notice, however, that numerous significant differences between the tasks remain: task timing, decoding algorithms used, and the neuronal population available for decoding were different, and, in addition, the study by Taylor et al. was based on a center-out task, whereas this study relied on a point-to-point reaching task. These differences in experimental variables make the quantitative comparison of M1 and PPC in brain-control tasks difficult. Nevertheless, the results reported from PPC are qualitatively similar to M1.

SI Materials and Methods

Model Assessment. The offline reconstruction performance of each decoding algorithm was quantified using the coefficient of determination, R^2. Separate R^2 values for position, velocity, and acceleration in x, y, and z directions were determined, and results for the three degrees of freedom (df) were combined to provide composite R^2 values for position, velocity, and acceleration, respectively.

Optimal Lag Time. The OLT (the temporal offset between instantaneous movement state and instantaneous neural representation where tuning is maximal) can vary considerably for PPC neurons, ranging from sensory (38, 39), where the behavior leads the neural representation, to instantaneous, to motor (23), where the neural representation leads the movement. For a motor prosthetic, only causal neural representations generating motor behavior are of interest, but because the precise OLT for the neurons in the population is unknown, one single 90-ms binning interval (Kalman filter) may not be sufficient to capture a neuron's firing at its optimal lead time. To optimize the decoding success, the optimal temporal location of the binning interval was, therefore, determined by shifting it iteratively from 45 ms (first possible causal 90ms bin) to 125 ms in 10-ms increments. The OLT resulting in the most accurate trajectory reconstruction was chosen for decoding. The analysis was performed for the entire population.

VR Task Details. All experiments were conducted in a VR environment (3) providing closed-loop, real-time visual feedback. The virtual workspace scaled 1:1 with the monkey's physical workspace: the stereoscopic rendering considered viewing distance and eye separation to provide realistic visual disparity, object size scaled with depth, and transparency allowed occluded or intersecting objects to be perceived. The head-fixed monkey sat in a dark room facing a mirror oriented at 45° presenting the image to him from an overhead mounted downward-facing monitor. The mirror prevented the monkey from observing his own limbs (Fig. 1*A*). Hand position was recorded at 100-Hz sampling rate using an infrared 3D motion capture system (OPTOTRAK 3020; Northern Digital). Three-dimensional shutter glasses (NuVision 60GX; MacNaughton) and a CRT monitor were used for 3D stereoscopic visualization. Eye position was monitored at a 120-Hz sampling rate (ETL200; ISCAN).

The virtual workspace consisted of a $20 \times 20 \times 20$ cm cube with the possible reach-target locations arranged in a $3 \times 3 \times 3$ cubic grid (edge length 10 cm), centered within the workspace. Targets were presented by a green, semitransparent sphere (diameter, 32 mm), and the monkey's cursor was represented by a solid white sphere (diameter, 31 mm).

Ridge Regression Filter Details. The least-squares solution for β yields the minimum variance, unbiased estimator. Zero-bias estimators, however, often suffer from high mean squared error (MSE) because of a large variance component of the error. The ridge regression, a variant of the ordinary multiple linear regression, was, therefore, used to estimate the movement state. The ridge regression reduces the variance component of the error, while allowing a small increase in the bias (4) by adding a complexity term to the optimization problem. The complexity term penalizes coefficients in β for having large weights (5), such that:

$$\hat{\beta}^{ridge} = \arg\min_{\beta} \left\{ \underbrace{\sum_{k}^{M} \left(x(k) - \beta_0 - \sum_{j}^{N} \beta_j r(k)_j \right)^2}_{\text{Least-squares Term}} + \underbrace{\lambda \sum_{j}^{N} \beta_j^2}_{\text{Complexity Term}} \right\}, \quad \textbf{[S1]}$$

where M is the number of training samples used in a session. The regularization parameter, λ, was chosen iteratively to minimize the MSE. A gradient descent algorithm was used to find the optimal λ iteratively. For a given value of λ, a unique solution for the ridge coefficients can be expressed in matrix notation, when estimating the 3D cursor position:

$$\hat{\beta}^{ridge} = \left(R^T R + \lambda I\right)^{-1} R^T X, \quad \textbf{[S2]}$$

where $R \in \Re^{M \times N}$ is the standardized firing rate matrix sampled at four lag time steps, $X \in \Re^{M \times 3}$ is the mean-subtracted 3D position matrix, and $\beta \in \Re^{N \times 3}$ are the model coefficients unique to a particular λ.

1. Hwang EJ, Andersen RA (2011) Effects of visual stimulation on LFPs, spikes, and LFP-spike relations in PRR. *J Neurophysiol* 105:1850–1860.
2. Taylor DM, Tillery SI, Schwartz AB (2002) Direct cortical control of 3D neuroprosthetic devices. *Science* 296:1829–1832.
3. Hauschild M, Davoodi R, Loeb GE (2007) A virtual reality environment for designing and fitting neural prosthetic limbs. *IEEE Trans Neural Syst Rehabil Eng* 15:9–15.
4. Hoerl AE, Kennard RW (1970) Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 12:55–67.
5. Hastie T, Tibshirani R, Friedman J (2001) *The Elements of Statistical Learning* (Springer, New York).

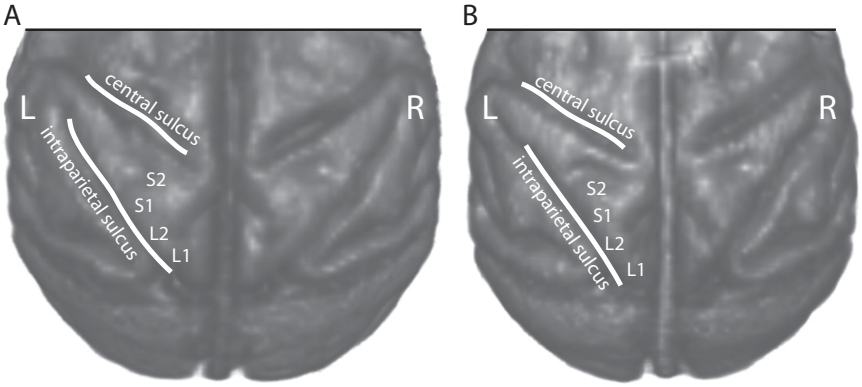

Fig. S1. PPC recording array arrangement. Four arrays, each containing 32 electrodes, were implanted stereotaxically using MRI to guide the implantation. In both animals, all arrays were in the left hemisphere, i.e., contralateral to the hand used. The short electrode arrays (S1, S2) were designed to record neurons from the cortical surface (electrode lengths are between 1.2 and 1.8mm). The long electrode arrays (L1, L2) contained a limited number of long electrodes allowing to record neural activity from PRR in the medial bank of the intraparietal sulcus (electrode lengths: 16 electrodes between 4.7 and 7.1 mm, and 16 electrodes between 1.5 and 3.9 mm). Seventy-five percent of all electrodes, therefore, targeted the cortical surface. (*A*) Three-dimensional reconstruction of monkey R's brain. The arrays are shown at the approximate surface locations. The stereotaxic coordinates were: L1: −5.7 lateral (L)/−9.0 posterior (P)/oriented 5° toward midline; L2: −8.5 L/−5.7 P/oriented 30° toward midline; S1: −10.0 L/−1.9 P/oriented surface-normal; S2: −8.0 L/1.4 anterior (A)/oriented surface-normal. The insertion angles for arrays L1 and L2 were chosen to target PRR in the medial bank of the intraparietal sulcus. (*B*) Three-dimensional reconstruction of monkey G's brain with approximate array implant locations. The stereotaxic coordinates were L1: −4.1 L/−10.0 P/oriented 4° toward midline; L2: −7.0 L/−6.6 P/oriented 18° toward midline; S1: −8.0 L/−1.5 P/oriented surface-normal; S2: −8.4 L/4.1 A/oriented surface-normal. (All units are mm.)

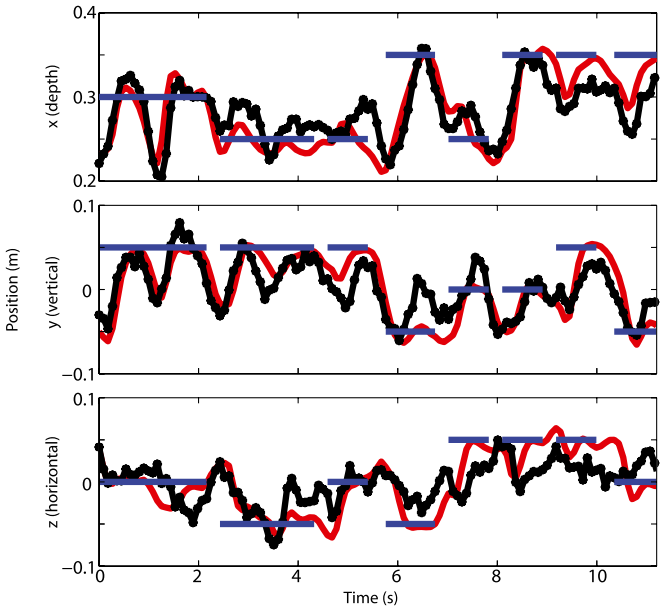

Fig. S2. Offline ridge filter 3D trajectory reconstruction under free gaze. Position reconstruction (black) of a previously recorded sequence of reaches (red) to 8 targets (blue). ●, discrete reconstruction points resulting from the 90-ms sampling rate used. In comparison with Kalman reconstructions of trajectories (e.g., Fig. 2*A*), the trajectories appear to be slightly less accurate and noisier, as confirmed by the R^2 results in Tables 1 and 2.

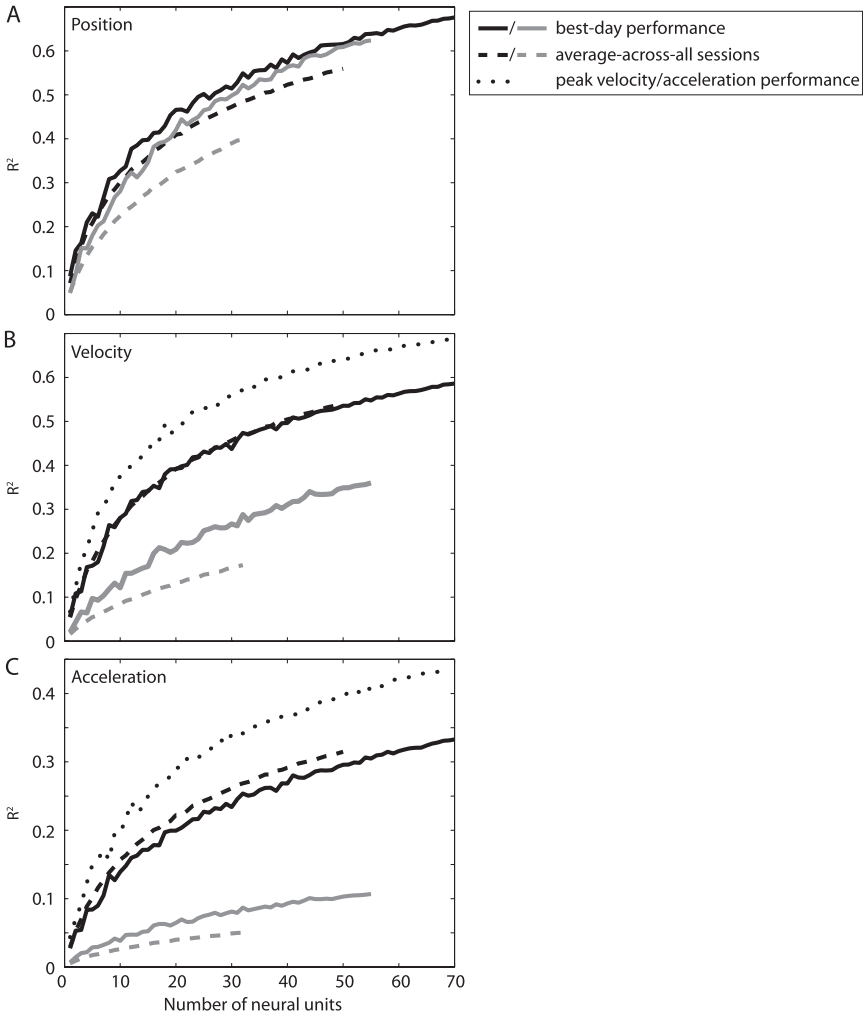

Fig. S3. To quantify the decoding efficiency, we calculated neuron-dropping curves describing how the decoding performance varied as a function of the PPC ensemble size. The Kalman filter decode was used to construct the curves. (*A*) Neuron-dropping curves, which plot R^2 for decoding cursor position as a function of ensemble size for monkey R (black) and G (gray), for best-day performance (solid) and average performance across all sessions (dashed). (*B* and *C*) Neuron-dropping curves for velocity and acceleration, respectively. The solid curves represent velocity and acceleration performances obtained on the day of the best position-decoding performance, whereas the dashed curves show the average across all recording days (the curve for monkey R partially occluded by the best-day curve). The dotted curves for velocity and acceleration (monkey R only) represent the best achievable velocity and acceleration-decoding accuracies, which were obtained during recording sessions different from the session when best position-decoding performance was reached. In monkey G, the best position-, velocity-, and acceleration-decoding accuracy was achieved on the same day. Note that the maximum number of neural units considered to calculate average performance was constrained to 50 (monkey R) and 32 (monkey G), the smallest ensemble size of any of the sessions reported.

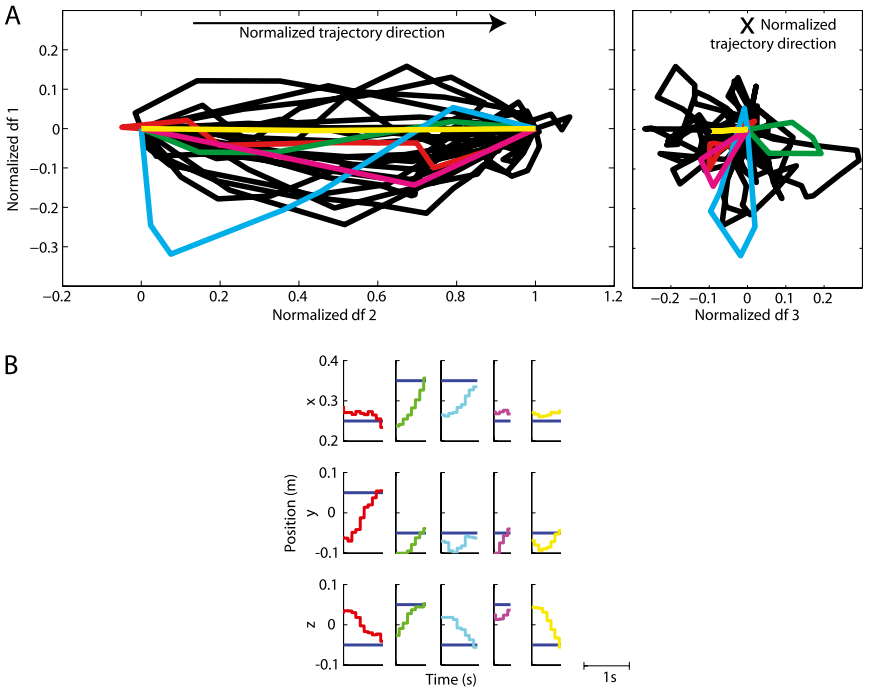

Fig. S4. Online 3D brain-control trajectory summary. (*A*) Orthographic summary plot of trajectories to all 27 targets, allowing a qualitative assessment of the trajectories. Unlike in a center-out task, where reaches are limited to relatively few directions and to a fixed amplitude, trajectories could point in many directions and had varying amplitudes. To summarize a set of such diverse trajectories in a single plot, all brain-control reaches were normalized for direction and amplitude. All trajectories were first reoriented such that the beginning and endpoint were located on a common axis: the direction of all trajectories in the left graph is to the right. To show the third df of the 3-dimensional trajectories, we then rotated the view by 90° such that all trajectories pointed into the page, resulting in the right graph with the third df showing, instead of the second. (*B*) For comparison, the five color-coded trajectories in *A* are plotted against time.

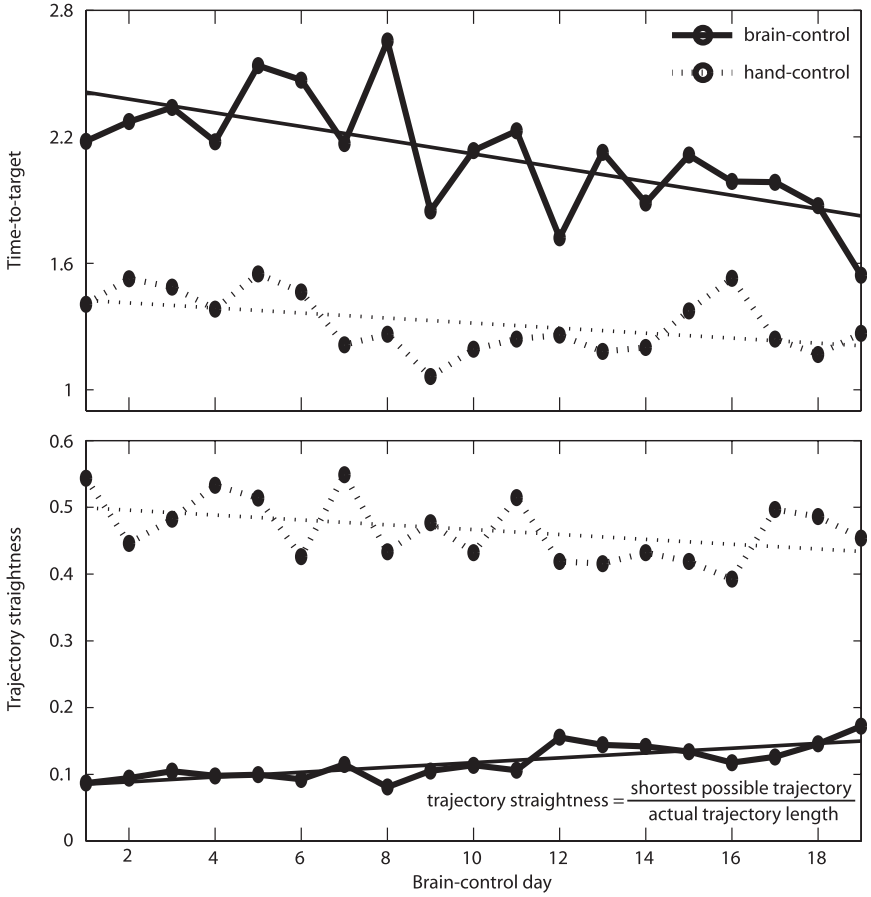

Fig. S5. Comparison of same-day hand-control and brain-control performance. In contrast to Fig. 4*A*, where trajectory straightness was normalized to facilitate the comparison of brain-control performance between animals R and G, this figure shows the nonnormalized behavioral metrics, allowing us to quantitatively benchmark the degree of task proficiency under brain control relative to hand control, which serves as the most natural performance reference on which a prosthetic device can be measured. In monkey R, the time to target (*Upper*) decreased more rapidly under brain control ($m = -0.03261$; 95% CI: −0.05157/−0.01365) than under hand control ($m = -0.01195$; 95% CI: $-0.02355/-3.471 \times 10^{-4}$). On the first day, monkey R acquired targets on average ∼1.55 times slower under brain control than under hand control, whereas on day 19, he was only 1.22 times slower, a 21% reduction of the time-to-target ratio for brain control vs. hand control. Simultaneously, trajectory straightness (*Lower*) increased significantly ($m = 0.003599$; 95% CI: 0.002229/0.004968). In contrast, trajectory straightness under hand control did not change significantly ($m = -0.003634$; 95% CI: $-0.00759/3.23 \times 10^{-4}$). The trajectory straightness ratio for brain control vs. hand control improved from 0.1600 measured on the first day to 0.3796 measured on day 19, a >100% improvement. To perform a fair comparison between point-to-point hand-control reaches and individually rewarded brain-control reaches, the first reach in all hand-control sequences was isolated and used for comparison with the individually rewarded brain-control reaches.

Movie S1. Three-dimensional reach control. Each monkey guided a cursor (white) in a 3D VR display to a reach target (green). In reach-control mode, the monkeys used their hand to control cursor movement. The movie faithfully reproduces the visual workspace as perceived by the monkey while performing two successful sequences of reaches to six targets each. Although the monkey was performing in a stereoscopic true 3D workspace, the 2D movie limits workspace presentation to a monoscopic view, thus making size scaling with depth the only readily perceivable depth cue in the movie. The inset shows a side view of the monkey's limb with the infrared motion tracking device attached to the hand. The recorded reach sequences and simultaneously obtained neural signals from PPC were used to identify the decoding algorithm used for subsequent brain-control reaches.

Movie S1

Movie S2. Brain control. The decoding algorithm identified from a set of reach-control sequences was used for brain control where cursor movement was governed by the monkeys' cortical signals. Under brain control, the monkeys were rewarded for individual reaches. The movie shows a set of 12 consecutive successful brain-control reaches from the monkey's perspective, performed by monkey R using the Kalman filter algorithm. Notice that although the cursor was under direct cortical control, the monkey performed limb movements (small movie inset). Monkey R chose to move his limb even during the reward break when no reach target was present. Importantly, during this phase, cursor movement appeared to remain synchronized with hand movement, suggesting that PPC provided accurate control signals not only when the monkey was fully engaged in the task but also during moments of distraction, such as reward.

Movie S2

Movie S3. Brain control in the absence of limb movement. To test whether PPC can generate prosthetic control signals in absence of sensory feedback from the limb, the animal was encouraged to perform brain-control reaches without moving the limb. The movie shows a set of 10 selected successful brain-control reaches, demonstrating that PPC neural activity does not depend on sensory feedback from the limb to generate cursor control commands. This result suggests that the PPC of patients suffering from paralysis will be able to provide neural signals suitable for control of a prosthetic assist device.

Movie S3