

# Neural control of virtual ego-motion enabled by an opportunistic decoding strategy

Karen E Schroeder\*<sup>1,2</sup>, Sean M Perkins\*<sup>2,3</sup>, Qi Wang<sup>3</sup>, & Mark M Churchland<sup>1,2,4,5, †</sup>

\* Equal contribution

<sup>1</sup>Department of Neuroscience, Columbia University Medical Center, New York, NY

<sup>2</sup>Zuckerman Institute, Columbia University, New York, NY

<sup>3</sup>Department of Biomedical Engineering, Columbia University, New York, NY

<sup>4</sup>Kavli Institute for Brain Science, Columbia University Medical Center, New York, NY

<sup>5</sup>Grossman Center for the Statistics of Mind, Columbia University, New York, NY

<sup>†</sup>Correspondence to: [mc3502@columbia.edu](mailto:mc3502@columbia.edu)

## 1 **Abstract**

2 Brain-machine interfaces (BMIs) for reaching have enjoyed continued performance  
3 improvements. Yet there remains significant need for BMIs that control other types of  
4 movement, including ego-motion through the world. Based on recent scientific findings, the  
5 class of decode algorithms employed by reach-based BMIs seems unlikely to generalize well. To  
6 examine this, we developed an ego-motion BMI based on cortical activity as monkeys cycled a  
7 hand-held pedal to progress along a virtual track. Unlike reaching, strong correlations between  
8 neural activity and kinematics were not present during cycling. We thus employed an  
9 opportunistic decode strategy that abandoned any notion of inverting encoding, and instead  
10 identified and leveraged dominant features of the population response. Online BMI-control  
11 success rates approached those during manual control. We argue that, in retrospect, reach-  
12 based BMIs succeeded by implicitly using opportunistic strategies, and that explicitly  
13 embracing that approach will be essential for expanding the range of BMI-controllable  
14 movements.

## 15 Introduction

16 Brain-machine interfaces (BMIs) interpret neural activity and provide control signals to external  
17 devices such as computers and prosthetic limbs. Intracortical BMIs for reach-like tasks have  
18 proved successful in primates and human clinical trials<sup>1-8</sup>, and more widespread use appears  
19 imminent. Yet at the same time, there exist non-reach-like movements whose restoration is  
20 valuable to patients. As one example, many patients could benefit from a BMI that controls ego-  
21 motion: movement of the self through one's environment. While ego-motion BMIs could be  
22 guided by traditional decoding approaches, other viable strategies exist and remain unexplored.  
23 This highlights a broader issue: a focus on reaching and grasping has produced state-of-the-art  
24 decode algorithms that may not generalize well to non-reaching applications.

25 The early success of reach-based BMIs<sup>9-14</sup> employed the 'biomimetic' strategy of inverting  
26 neural encoding of kinematic variables such as hand velocity / direction. Multiple lines of  
27 evidence now argue against the hypothesis that motor cortex activity literally encodes  
28 kinematic variables<sup>15-20</sup>. Nevertheless, robust correlations between neural activity and direction  
29 allow excellent BMI performance within the confines of a reaching task<sup>9-14</sup>. This core strategy  
30 has thus endured even as decoder sophistication has increased; improvements have derived  
31 largely from better estimating the neural state, thus improving the reach-velocity decode<sup>21-23</sup>.  
32 Decode algorithms for other movement types, including ego motion, have received less  
33 attention. Furthermore, they have typically borrowed the strategies employed by reach-based  
34 BMIs: e.g., decoding a whole-body directional vector<sup>24</sup> or classifying the direction of a joystick  
35 intermediary<sup>25</sup>. Traditional kinematic decoding strategies have also been used for decoding  
36 during treadmill-walking<sup>26,27</sup> and for a brain-spine interface that alleviates gait deficits<sup>28</sup>.

37 Yet there is emerging evidence that BMI decoding may not be best achieved by attempting to  
38 invert encoding. We recently argued that the dominant signals in motor cortex – i.e., those with  
39 the greatest influence on firing rates – exist to ensure noise-robust dynamics and do not encode  
40 any external quantity<sup>16</sup>. In contrast, neural signals that relate directly to outgoing commands  
41 (e.g., downstream muscle activity) are small. In many BMI applications, including controlling  
42 ego-motion, reconstructing the full set of output commands is not necessary to accomplish the  
43 user's goals. This opens the door to strategies that abandon any notion of decoding encoded  
44 variables.

45 We introduce a new approach – opportunistic feature decoding – that seeks to leverage robust  
46 relationships between neural activity and variables that matter to the user, regardless of  
47 whether those relationships are fundamental. Opportunistic decode strategies do not require  
48 knowledge of the true relationship between neural activity and motor output, only an accurate

49 description for the behaviors one wishes to decode. Opportunistic strategies are potentially  
50 more noise robust, as they can leverage the dominant signals in the population activity. In  
51 contrast, a true biomimetic strategy is limited to the smaller output signals.

52 We explore the potential of opportunistic decoding for ego motion using a task in which  
53 monkeys cycle a hand-held pedal to move along a virtual track. Neural population activity  
54 during cycling differed from that during reaching in fundamental ways. ‘Directional’ signals  
55 that correlated with kinematics were only weakly reflected in neural firing rates. Signals related  
56 to muscle activity were of similarly small magnitude. Still, there were three dominant features  
57 of the neural response that possessed robust relationships with desired self-motion. The first  
58 was a translation of the neural population state that indicated whether the monkey was  
59 moving. The second was an elliptical neural trajectory that repeated while cycling. This  
60 trajectory did not reverse with cycling direction (as would a representation of hand velocity),  
61 but instead occupied direction-dependent dimensions. As a result, the direction and magnitude  
62 of self-motion could be estimated by comparing the angular momentum of the neural state  
63 between pairs of dimensions. This feature was robust during ongoing cycling but not at  
64 movement initiation. We thus identified a third feature: at movement initiation, neural activity  
65 corresponding to forward and backward cycling was briefly linearly separable.

66 A decoder leveraging these features provided excellent online control of virtual ego-motion.  
67 Success rates and acquisition times were close to those achieved under manual control. Almost  
68 no training or adaptation time was needed; the low-latency and accuracy of the decoder were  
69 such that monkeys appeared to barely notice transitions from manual control to BMI control.  
70 Although quantitative comparison with previous ego-motion BMIs is challenging due to very  
71 different task designs and very different patterns of underlying activity, our strategy provides  
72 obvious improvements in some domains. For example, BMIs have historically had difficulty  
73 decoding lack of motion, and the tasks used to evaluate previous ego-motion BMIs required  
74 traveling to the target but not stopping. The decode provided by the opportunistic strategy  
75 exhibited virtually no false starts (order of one per day), and allowed subjects to transition from  
76 rapid motion to stopping within ~100 ms of intending to do so.

77 Our results establish that BMIs can be accurately guided by decode strategies that eschew  
78 attempts to invert encoding, and instead leverage the dominant structure of population activity.  
79 Because that dominant structure likely reflects internal dynamics rather than outgoing signals,  
80 it can be very different across different behaviors. Our results argue that focusing on robust  
81 population-level features, rather than correlations with kinematics, will be necessary to expand  
82 BMI control to a broader range of behaviors.

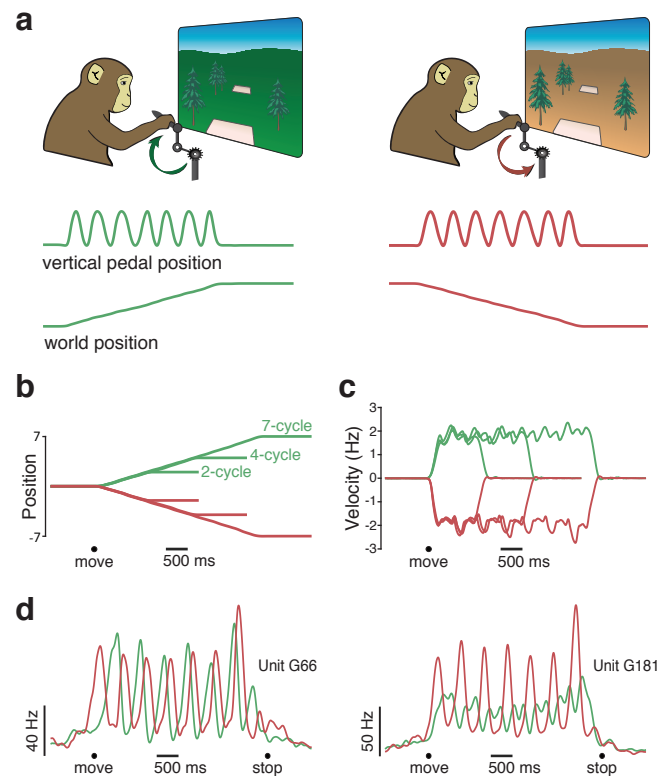
## 83 Results

### 84 *Behavior*

85 We trained two monkeys (G and E) to rotate a hand-held pedal to move through a virtual  
86 environment (**Fig. 1**). All motion was along a linear track – no steering was necessary.  
87 Consistent with this, a single pedal was cycled with the right arm only. Our goal when  
88 decoding was to reconstruct the virtual motion produced by that single pedal. On each trial, a  
89 target appeared in the distance. To acquire that target, monkeys produced virtual velocity in  
90 proportion to the rotational velocity of the pedal. The color of the environment (lush and green  
91 versus desert-like and tan) instructed cycling direction. When the environment was green (**Fig.**  
92 **1a, left**) forward virtual motion was produced by cycling ‘forward’ (i.e., with the hand moving  
93 away from the body at the top of the cycle). When the environment was tan (**Fig. 1a, right**)  
94 forward virtual motion was produced by cycling ‘backward’ (the hand moving toward the  
95 body at the top of the cycle). Cycling in the wrong direction produced motion away from the  
96 target. Trials were presented in blocks of forward or backward trials. Within each block, targets  
97 were separated by a randomized distance of 2, 4 or 7 cycles. Acquisition of a target was  
98 achieved by stopping and remaining stationary ‘on top’ of the virtual target for a specified time.  
99 Reward was then given and the next target appeared.

100 Monkeys performed the task well, moving swiftly between targets, stopping accurately on each  
101 target, and remaining stationary until the next target was shown. Monkeys cycled at a pace that  
102 yielded nearly linear progress through the virtual environment (**Fig. 1b**). Although not  
103 instructed to cycle at any particular angular velocity, monkeys adopted a brisk ~2 Hz rhythm  
104 (**Fig. 1c**). Small ripples in angular velocity were present during steady-state cycling; when  
105 cycling with one hand it is natural for velocity to increase on the downstroke and decrease on  
106 the upstroke. Success rates were high, exceeding 95% in every session (failures typically  
107 involved over- or under-shooting the target location). This excellent performance under manual  
108 control provides a stringent bar by which to judge performance under BMI control.

109 BMI control was introduced after monkeys were adept at performing the task under manual  
110 control. Task structure and the parameters for success were unchanged under BMI control, and  
111 no cue was given regarding the change from manual to BMI control. The switch to BMI control  
112 was made at the beginning of the session, after completion of a block of manual-control trials  
113 (25 forward and 25 backward 7-cycle trials). These manual-control trials were used to train the  
114 decoder. The switch was then made to BMI control for the remainder of the session. For monkey  
115 G, we occasionally included blocks of manual-control trials later in the session to allow



**Figure 1. A cycling task that elicits rhythmic movements.** (a) Monkeys rotated a hand-held pedal forward (*left*, cued by a green background) or backward (*right*, cued by a tan background) to progress through a virtual environment. Traces at bottom plot pedal kinematics (vertical position) and the resulting virtual world position for two example manual-control trials. On both of these trials (one forward and one backward) the monkey progressed from one target to another by cycling seven cycles. (b) Trial-averaged virtual position from a typical manual-control session. Each trace plots the change in virtual position (from a starting position of zero) for one of six conditions: forward or backward for 2, 4, or 7 cycles. Black circle indicates the time of movement onset. Trials were averaged after being aligned to movement onset, and then scaled such that the duration of each trial matched the average duration for that condition. (c) Trial-averaged pedal rotational velocity from the same session, for the same six conditions. (d) Firing rates of two example units. Trial-averaged firing rates (computed after temporally aligning trials) are shown for two conditions: forward (*green*) and backward (*red*) for seven cycles. Black circles indicate the timing of movement onset and offset.

116 comparison between BMI and manual performance. For Monkey E we used separate  
117 (interleaved) sessions to assess manual-control performance.

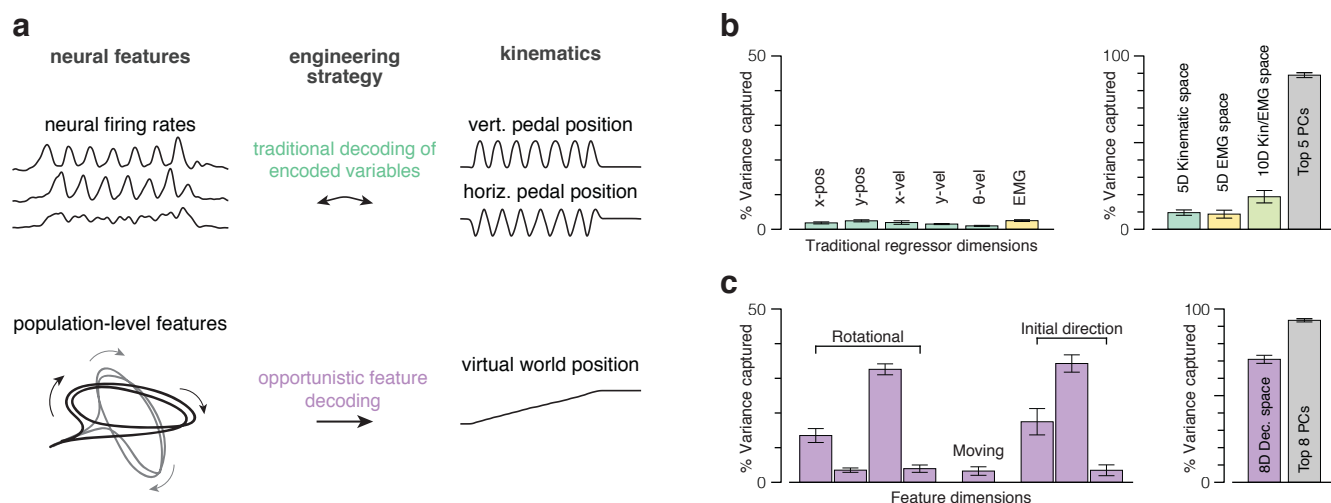
118 During both BMI control and manual control, the monkey's ipsilateral (non-cycling) arm was  
119 restrained. The contralateral (cycling) arm was never restrained. We intentionally did not  
120 dissuade the monkey from continuing to physically cycle during BMI control. Indeed, our goal  
121 was that the transition to BMI control would be sufficiently seamless to be unnoticed by the  
122 monkey, such that he would still believe that he was in manual control. An advantage of this  
123 strategy is that we are decoding neural activity when the subject attempts to actually move, as a  
124 patient presumably would. Had we insisted the arm remain stationary, monkeys would have  
125 needed to actively avoid patterns of neural activity that drive movement – something a patient  
126 would not have to do. Allowing the monkey to continue to move normally allowed us to  
127 extensively quantify the performance of our decoder by comparing decoded with intended (i.e.,  
128 actual) movement. This is often not possible when using other designs. For example, in  
129 Rajangam et. al.<sup>24</sup>, performance could only be assessed via indirect measures (such as time to  
130 target) because what the monkey was actually intending to do at each moment was unclear. We  
131 considered these advantages to outweigh a potential concern: a decoder could potentially  
132 'cheat' by primarily leveraging activity driven by proprioceptive feedback (which would not be  
133 present in a paralyzed patient). This is unlikely to be a large concern. Recordings were made  
134 from motor cortex, where robust neural responses precede movement onset. Furthermore, we  
135 have documented that motor cortex population activity during cycling is quite different from  
136 that within the proprioceptive region of primary somatosensory cortex<sup>16</sup>. Thus, while  
137 proprioceptive activity is certainly present in motor cortex<sup>29-32</sup> (especially during  
138 perturbations<sup>33</sup>) the dominant features of M1 activity, described below, are unlikely to be  
139 primarily proprioceptive.

140 Given our use of healthy animals, we stress that the goal of the present study is to determine  
141 how the dominant structure of neural activity can be leveraged for accurate prosthetic decode.  
142 This follows the successful strategy of BMI studies that leveraged the known structure of  
143 activity during reaching<sup>9,21</sup>. Of course, the nature of the training data used to specify decode  
144 parameters (e.g., the neural dimensions to be used) will necessarily be different for a healthy  
145 animal that cannot understand verbal instructions and an impaired human that can. We thus  
146 stress that our goal is to determine a robust and successful decode strategy that works in real  
147 time during closed-loop performance. We do not attempt to determine the best approach to  
148 parameter specification, which in a patient would necessarily involve intended or imagined  
149 movement.

## 150 *Neural activity and decoding strategy*

151 We recorded motor cortical activity using 96-channel Utah arrays. For monkey G, one array was  
152 implanted in primary motor cortex (M1) and a second in dorsal premotor cortex (PMd). For  
153 monkey E, a single array was implanted in M1. For each channel we recorded times when the  
154 voltage crossed a threshold. Threshold crossings typically reflected individual spikes from a  
155 small handful of neurons (a neural ‘unit’). Spikes from individual neurons could be clearly seen  
156 on many channels, but no attempt was made to spike-sort, as the benefit of doing so is typically  
157 modest when controlling a prosthetic device<sup>34</sup>. Unit activity was strongly modulated during  
158 cycling (**Fig. 1d**). The phase, magnitude, and temporal pattern of activity depended on whether  
159 cycling was forward (*green* traces) or backward (*red* traces). A key question is how these unit-  
160 level features translate into population-level features that might be leveraged to estimate  
161 intended motion through the virtual environment.

162 In traditional decoding approaches (**Fig. 2a, top**) neural activity is hypothesized (usefully if not  
163 literally) to encode kinematic signals, which can be decoded by inverting the encoding scheme.  
164 Although nonlinear methods (including variations of Kalman filtering) are often used to  
165 estimate the neural state, the final conversion to a kinematic command is typically linear or  
166 roughly so. To explore kinematic encoding in the present task, we used linear regression to  
167 identify neural dimensions where activity correlated well with kinematics (including hand  
168 velocity and position). Regression was performed using single trials. Use of single trials  
169 provides a large quantity of training data and is implicitly regularizing: regression must find  
170 signals that are robust in the face of single-trial spiking variability. The regression weights for a  
171 given kinematic parameter define a neural dimension where activity correlates strongly with  
172 that parameter. We computed the neural variance captured by each such dimension. Variance  
173 captured was computed using trial-averaged data, to ensure that values were not diluted by  
174 noise. Despite this, the neural dimensions that best captured kinematic signals captured little  
175 population response variance (**Fig. 2b**, green bars). This was also true of neural dimensions that  
176 captured muscle activity (**Fig. 2b**, yellow bar). This was initially surprising: single-neuron  
177 responses were robustly sinusoidally modulated, as were many kinematic variables. Yet  
178 sinusoidal response features were often superimposed upon other response features (e.g.,  
179 overall shifts in rate when moving versus not moving). Sinusoidal features also did not display  
180 phase relationships, across forward and backward cycling, that were consistent with kinematic  
181 encoding<sup>16</sup>. As a result, the dimensions where activity correlated strongly with kinematics  
182 captured relatively little response variance.



**Figure 2. Different decode strategies leverage neural signals with different magnitudes. (a)**

Two possible decoding strategies. In the first (*top*) neural firing rates are assumed to predominantly encode the key variables. The encoding model is usually assumed to be roughly linear when variables are expressed appropriately. For example, cosine tuning for reach velocity is equivalent to a linear dependence on horizontal and vertical velocity. The goal of decoding is to invert encoding. Thus, decoding dimensions should capture the dominant signals in the neural data (because those are what is encoded). The second strategy (*bottom*) can be applied even if the dominant signals do not have the goal of encoding. This strategy seeks to find neural response features that have a robust relationship with the variable one wishes to decode. That relationship may be complex or even incidental, but is useful if it involves high-variance response features. **(b)** Variance of the neural population response captured by dimensions used to decode kinematic parameters (*green bars*) and muscle activity (*yellow bar*). Data are from three manual-control sessions where units (192 channels per day) and muscles (5-7 channels per day) were recorded simultaneously. Each bar plots the average and standard error across sessions (unless otherwise specified). *Left subpanel*: variance captured for kinematic variables (individual variables shown separately) and muscles (average across 19 recordings, standard error computed across recordings). *Right subpanel*: total variance captured by subspaces spanned by kinematic-decoding dimensions, muscle-decoding dimensions, or both. (These are not the sum of the individual variances as dimensions were not always orthogonal). We had different numbers of EMG recordings per day and thus always selected a subset of five. Variance captured by the top five principal components is shown for comparison. **(c)** Similar plot but for the dimensions upon which our decoder was built. *Left subpanel*: variance captured for each of these eight dimensions. *Right subpanel*: variance captured by the eight-dimensional subspace spanned by those dimensions. Variance captured by the top eight principal components is shown for comparison.



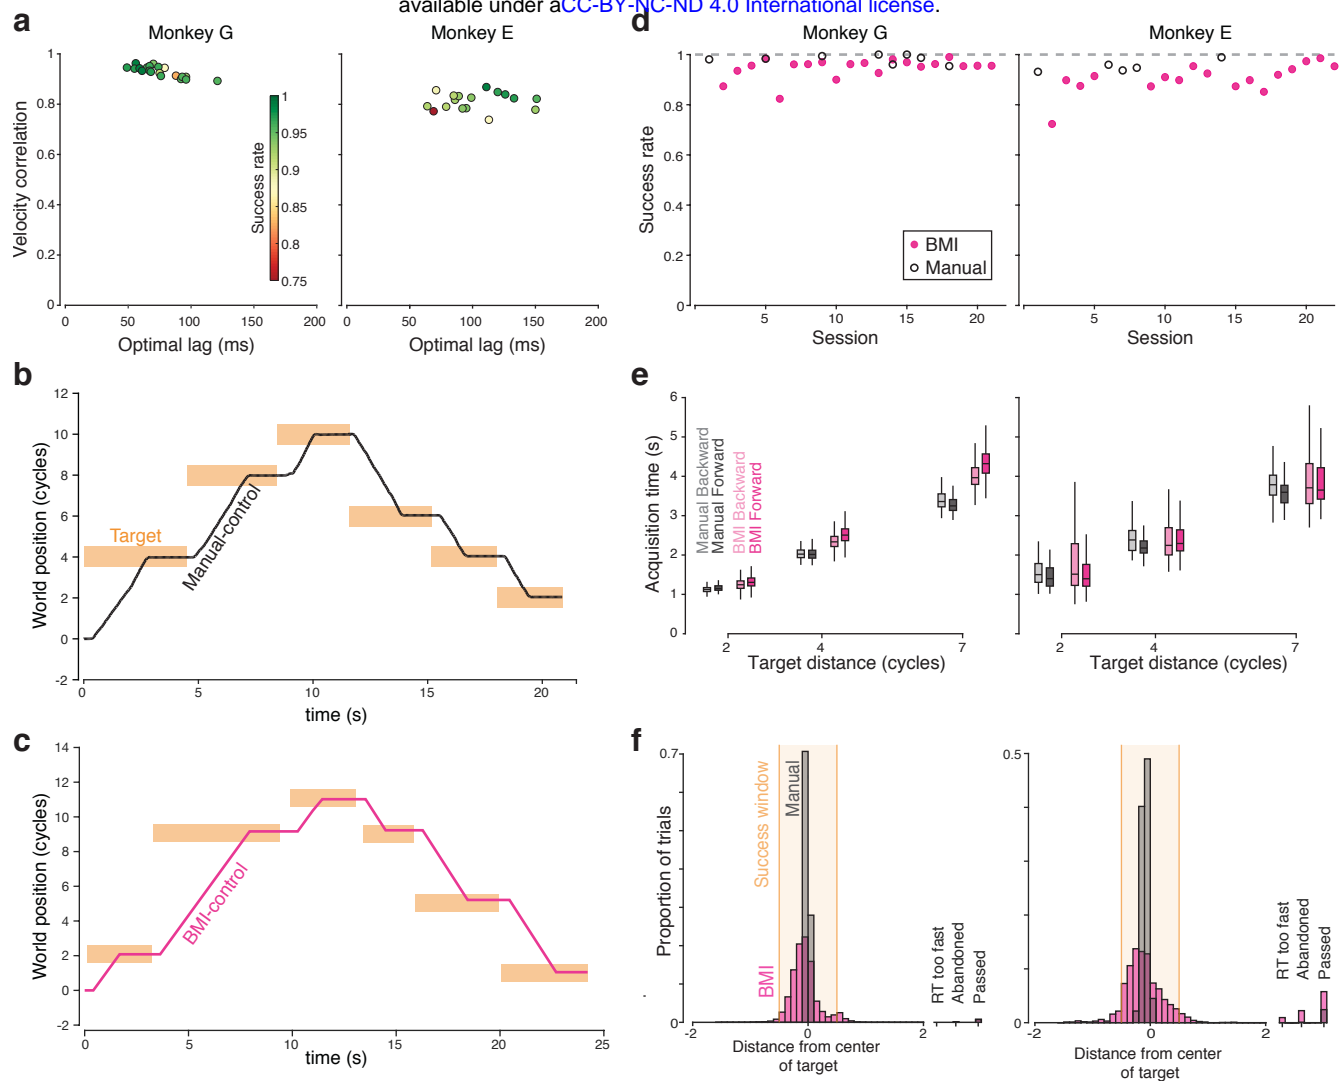
183 Low-variance signals are a poor candidate for decoding intended action; they are likely to be  
184 non-robust with respect to multiple challenges. Some of these challenges (e.g., spiking noise)  
185 can be anticipated and estimated in advance, but others cannot. They include recording  
186 instabilities, changes in strategy or behavior with time, and (outside the laboratory setting)  
187 external sources of noise or variability. Given these challenges, it is worth stressing that there  
188 were two practical reasons why reach-focused BMIs leveraged signals that correlate with hand  
189 velocity. First, such signals are high-variance during reaching – so much so that M1 responses  
190 have often been summarized in terms of a preferred direction<sup>35,36</sup>. Second, movement direction /  
191 velocity are the variables one wishes to decode during reaching. In the present case neither  
192 motivation holds. Signals related to hand kinematics are low-variance, and we most wish to  
193 decode self-motion through the virtual environment. This suggests an alternative strategy (**Fig.**  
194 **2a, bottom**): identifying neural response features that are both robust (high-variance) and relate  
195 reliably to the presence and direction of self-motion.

196 To pursue this strategy, we considered three sets of high-variance dimensions. The first set  
197 included four ‘rotational dimensions’ (two each for forward and backward cycling) which  
198 captured elliptical trajectories present during steady-state cycling<sup>16</sup>. The second set included a  
199 single ‘moving-sensitive’ dimension, in which the neural state distinguished whether the  
200 monkey was stopped or moving regardless of movement direction<sup>37</sup>. The third set was a triplet  
201 of ‘initial-direction’ dimensions. In these dimensions, cycling direction could be transiently but  
202 readily distinguished in the moments after cycling began.

203 In subsequent sections we document the specific features present in these high-variance  
204 dimensions. Here we concentrate on the finding that the space spanned by these eight  
205 dimensions captured  $70.9\% \pm 2.3\%$  of the firing-rate variance (**Fig. 2c**). This was only modestly  
206 less than that captured by the top eight PCs (which capture the most variance possible), and  
207 much more than that captured by spaces spanned by dimensions where activity correlated with  
208 kinematics and/or muscle activity (**Fig. 2b**). We thus based our BMI decode entirely on activity  
209 in these eight high-variance dimensions. Before describing how this was accomplished, we  
210 document the resulting performance.

### 211 *Performance*

212 Monkeys performed the task very well under closed-loop BMI control (**Fig. 3** and **Supp. Movie**  
213 **1**). Monkeys continued to cycle as normal, presumably not realizing that the pedal had been  
214 disconnected from the control system. The illusion that the pedal still controlled the task was  
215 supported by a high similarity between decoded virtual velocity and intended virtual velocity  
216 (i.e., what would have been produced by the pedal were it still controlling the task). The cross-



**Figure 3. Decoder performance.** (a) Summary of the cross-correlation between decoded virtual velocity under BMI control, and the virtual velocity that would have been produced by the pedal (which monkeys continued to manipulate normally). Each symbol corresponds to one BMI-control session, and plots the peak of the cross-correlation versus the lag where that peak occurred. Colors indicate success rate during that session. (b) Example manual-control performance for six consecutive trials, 3 forward and 3 backward. World position is expressed in terms of the number of cycles of the pedal needed to move that distance. For plotting purposes, the position at the beginning of this stretch of behavior was set to zero. Bars indicate the time that targets turned on and off (horizontal span) and the size of the acceptance window (vertical span). (c) Similar plot during BMI control. For ease of comparison, world position is still expressed in terms of the number of physical cycles that would be needed to travel that far, although physical cycling no longer had any impact on virtual velocity. (d) Success rate for both monkeys. Each symbol plots, for one session, the proportion of trials where the monkey successfully moved from the initial target to the final target, stopped within it, and remained stationary until reward delivery. Dashed line at 1 for reference. (e) Target acquisition times for successful trials. Center lines indicate median, the box edges indicate the first and third quartiles, and the whiskers include all non-outlier points (points less than 1.5 times the interquartile range from the box edges). Data are shown separately for the three target distances. (f) Histograms of stopping location from both monkeys. Analysis considers both successful and failed trials. The bar at far right indicates the proportion of trials where the monkey failed for reasons other than stopping accuracy per se. This included trials where monkeys disrespected the reaction time limits, abandoned the trial before approaching the target, or passed through the target without stopping.

217 correlation between these peaked at  $0.93 \pm .02$  and  $0.81 \pm .03$  (monkey G and E, mean  $\pm$  SD) at a  
218 short lag:  $76 \pm 4$  ms and  $102 \pm 7$  ms (**Fig. 3a**). That illusion was also aided by a low rate of false  
219 starts; it was exceedingly rare for decoded motion to be non-zero when the monkey was  
220 attempting to remain stationary on top of a target. False starts occurred on 0.29% and 0.09% of  
221 trials (monkeys G and E), yielding an average of 1.9 and 0.12 occurrences per day. This is  
222 notable because combatting unintended movement is a key challenge for BMI  
223 decoding<sup>2,38,392,38,39</sup>. The above features – high correlation with intended movement, low latency,  
224 and few false starts – led to near-normal performance under BMI control (**Fig. 3b,c**). Success  
225 rates under BMI control (**Fig. 3d**, *magenta symbols*) were almost as high as under manual control  
226 (*open symbols*), and the time to move from target to target was only slightly longer under BMI  
227 control (**Fig. 3e**).

228 The only respect in which BMI control suffered noticeably was accuracy in stopping on the  
229 middle of the target. Under manual control, monkeys stopped very close to the target center  
230 (**Fig. 3f**, *gray histogram*), which always corresponded to the ‘pedal-straight-down’ position.  
231 Stopping was less accurate under BMI control (*magenta histogram*). This was partly due to the  
232 fact that because virtual motion was swift, even small errors in decoded stopping time become  
233 relevant: e.g., a 100 ms error corresponds to  $\sim 0.2$  cycles of physical motion. The average  
234 standard deviation of decoded stopping time (relative to actual stopping time) was 133  
235 (monkey G) and 99 ms (monkey E). Increased stopping error in BMI-control trials was also due  
236 to an incidental advantage of manual control: the target center was aligned with the pedal-  
237 straight-down position, a fact which monkeys leveraged to stop very accurately in that position.  
238 This strategy was not available during BMI control because the correct time to stop rarely  
239 aligned perfectly with the pedal-straight-down position (this occurred only if decoded and  
240 intended virtual velocity matched perfectly when averaged across the cycling bout).

241 Performance was overall modestly better for monkey G versus E. This was likely due to the  
242 implantation of two arrays rather than one. Work ethic may also have been a factor; monkey E  
243 performed fewer trials under both BMI and manual control. Still, both monkeys could use the  
244 BMI successfully starting on the first day, with success rates of 0.87 and 0.74 (monkey G and E).  
245 Monkey G’s performance rapidly approached his manual-control success rate within a few  
246 sessions. Monkey E’s performance also improved quickly, although his manual-control and  
247 BMI-control success rates were mostly lower than Monkey G’s. The last five sessions involved  
248 BMI success rates of 0.97 and 0.96 for the two monkeys. This compares favorably with the  
249 overall averages of 0.98 and 0.95 under manual control. Although this performance  
250 improvement with time may relate to adaptation, the more likely explanation is simply that

251 monkeys learned to not be annoyed or discouraged by the small differences in decoded and  
252 intended velocity.


### 253 *State machine*

254 The performance documented above was achieved using a state-dependent decode (**Fig. 4**).  
255 Features of the neural activity (described more fully in subsequent sections) determined state  
256 transitions and what was decoded in each state. Briefly, state transitions were governed by  
257 activity in the moving-sensitive dimension, which was translated into a probability of moving,  
258  $p_{move}$ . If  $p_{move}$  was low, the STOP state was active and decoded virtual velocity was zero. When  
259  $p_{move}$  became high, the INIT state was entered but decoded velocity remained zero. After 175  
260 ms, the EARLY state was entered and velocity was decoded using the initial-direction  
261 dimensions. After an additional 200 ms, the STEADY state was entered and virtual velocity  
262 depended only on the neural state in the rotational dimensions. Decoded velocity was filtered  
263 to smooth fluctuations during STEADY.

264 Values of  $p_{move} < 0.1$  always produced a transition back to STOP. This typically occurred from  
265 STEADY to STOP, as the movement was successfully ending. However, it could also occur from  
266 the other two states. This was especially helpful if  $p_{move}$  became high very briefly (and  
267 presumably erroneously). In such cases the state could transition from INIT back to STOP with  
268 the decoded velocity never departing from zero. Below we describe how virtual velocity was  
269 estimated while in STEADY, how  $p_{move}$  was derived, and how we decoded the early direction  
270 of movement during EARLY.

### 271 *Direction of steady-state movement inferred from rotational structure*

272 The dominant feature of the neural response during steady-state cycling was a repeating  
273 elliptical trajectory<sup>16</sup>. Our decoder leveraged the fact that forward-cycling and backward-cycling  
274 trajectories occurred in non-identical dimensions. We employed an optimization procedure to  
275 find a two-dimensional ‘forward plane’ that maximized the size of the forward trajectory  
276 relative to the backward trajectory. We similarly found an analogous ‘backward plane’. These  
277 planes were identified based on trial-averaged responses from the 50 trials of training data  
278 collected under manual control (**Fig. 5a**). With the aid of filtering (*Methods*), these planes  
279 continued to capture rotational features on individual trials (**Fig. 5b**). Although forward and  
280 backward trajectories were not orthogonal to one another, the above procedure was still able to  
281 find planes where strongly elliptical trajectories were present for only one cycling direction.



**Figure 4. State machine diagram.** BMI motion was determined by a state machine with four states: STOP, INIT, EARLY, and STEADY, corresponding to the different stages of a typical trial. The output of the state machine at every millisecond was an estimate of decoded velocity through the virtual environment,  $v_{dec}$  which was then smoothed and integrated to compute virtual position. Black arrows indicate the typical path of a successful BMI trial and gray arrows indicate all other possible transitions. State transitions were governed by activity in the moving-sensitive dimension, which was translated into a probability of moving,  $p_{move}$ . While  $p_{move}$  was low, the STOP state was active and decoded velocity was set to zero. When  $p_{move}$  became high, the INIT state was entered but decoded velocity remained zero. If  $p_{move}$  remained high for 175 ms, the EARLY state was entered and velocity was decoded using the initial-direction dimensions. After another 200 ms, the STEADY state was entered and decoded velocity depended on the neural state in the rotational dimensions. If  $p_{move}$  dropped below 0.1 at any point, STOP was reentered. States in which progress is made through the virtual environment are highlighted in blue and states in which BMI motion is held at zero are highlighted in orange.

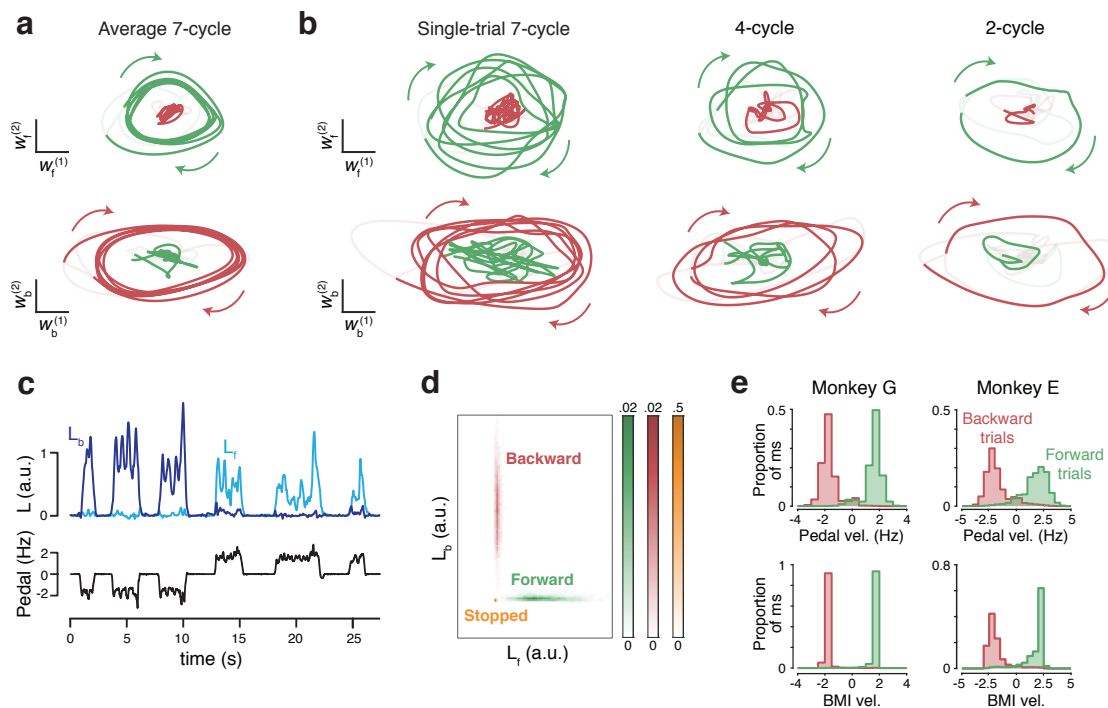
282 A common strategy for reaching prostheses is to linearly transform the neural state into a hand-  
283 velocity command; e.g. a state consistently to the right of zero would result in a consistently  
284 high rightwards velocity. In a given plane (e.g., backwards) the neural state traces a circle, and a  
285 plot of horizontal versus vertical hand velocity would also trace a circle. However, it would not  
286 be desirable to attempt to directly decode velocity. Not only would this require somehow  
287 choosing between planes, but a neural state consistently to the right of zero should not result in  
288 a consistent decode of rightwards hand velocity. A decode of hand position would be  
289 somewhat more natural but still awkward (there are four dimensions rather than two, and  
290 positions near zero are difficult to interpret). We thus chose the strategy of comparing angular  
291 momentum (the cross product of the state vector with its derivative) between the two planes.  
292 When moving backward (first three cycling bouts in **Fig. 5c**) angular momentum was sizeable in  
293 the backward plane (*dark blue*) but not the forward plane (*bright blue*). The opposite was true  
294 when moving forward (subsequent three bouts).

295 Based on training data, we considered the joint distribution of forward-plane and backward-  
296 plane angular momentum. We computed distributions when stopped (**Fig. 5d**, *orange*), when  
297 cycling forward (*green*) and when cycling backward (*red*). These distributions overlapped little,  
298 and we fit a Gaussian to each. During BMI control, we computed the likelihood of the observed  
299 angular momentums under each of the three distributions. If likelihood under the stopped  
300 distribution was high, decoded velocity was zero. Otherwise, decoded velocity was determined  
301 by the relative likelihoods under the forward and backward distributions. These likelihoods  
302 were converted into a virtual velocity that was maximal when one likelihood was much higher  
303 (which was typically the case) and slower when likelihoods were more similar. The maximum  
304 decoded virtual velocity was set to approximate the typical virtual velocity under manual  
305 control, when cycling at ~2 Hz.

306 The above steps were performed when in the STEADY state. Distributions of decoded velocity  
307 under BMI control (**Fig. 5e**, *bottom*) were similar to the distributions of velocity that would have  
308 resulted were the pedal still operative (**Fig. 5e**, *top*). Importantly, distributions overlapped very  
309 little; the direction of decoded motion was almost always correct. Decoded velocity was near  
310 maximal at most times, especially for monkey G. High accuracy and brisk velocities were  
311 responsible for the ability to move between targets almost as rapidly under BMI control as  
312 under manual control.

### 313 *Inferring the probability of moving*

314 Decoders that directly translate neural state to cursor velocity have historically had difficulty  
315 remaining stationary when there is no intended movement. The ability to do so is of even



**Figure 5. Leveraging rotational trajectories to decode velocity.** (a) Trial-averaged population activity, during a manual-control block, projected onto the forward (*top*) and backward (*bottom*) rotational planes. Data are from seven-cycle forward (*green*) and backward (*red*) conditions. By design, the forward plane primarily captures rotational trajectories during forward cycling, and vice versa. Boldly colored portions of each trace highlight rotations during the middle cycles (a period that excludes the first and last half cycle of each movement). Colored arrows indicate rotation direction. Light portion of each trace corresponds to the rest of the trial. In addition to smoothing with a causal filter, neural data have been high-pass filtered to match what was used during BMI control. Data are from monkey G. (b) As in panel (a), but for three example single trials, one for each of the three distances. (c) Example angular momentum ( $L$ ) in the backward plane (*dark blue*) and forward plane (*bright blue*) during six trials of BMI control. Velocity of the pedal is shown in black. Although the pedal was disconnected, this provides a useful indication of how the monkey was intending to move. Data are from the same day shown in panels a and b. (d) Probability densities of angular momenta found from the training dataset collected on the same day. (e) Histograms of BMI-control velocity (*bottom*) and (disconnected) pedal velocity (*top*) for all times the decoder was in the STEADY state, across all BMI-control sessions.

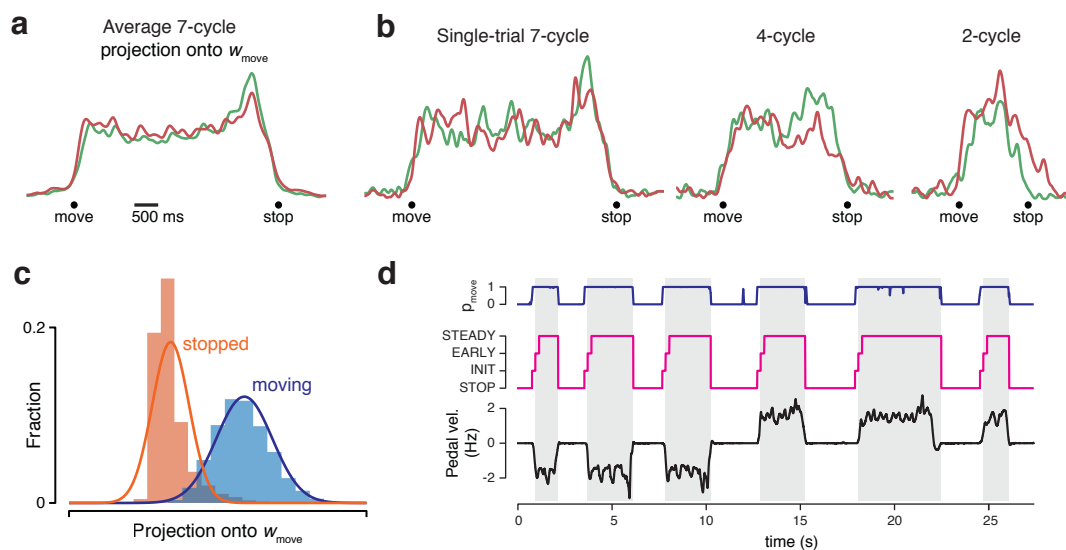
316 greater importance for an ego-motion prosthetic. To meet this challenge, we adopted the  
317 strategy of a state machine with distinct stopped and moving states<sup>38-40</sup>. Transitions between  
318 these states were governed by a probability of moving,  $p_{move}$ , derived from the neural state in  
319 the moving-sensitive dimension.

320 We identified the moving-sensitive dimension by applying linear discriminant analysis to the 50  
321 training-data trials, and finding the direction that best discriminated whether the monkey was  
322 moving versus stopped. Projecting trial-averaged data onto that dimension (**Fig. 6a**) revealed  
323 that activity transitioned suddenly from low to high just before movement onset, and back to  
324 low around the time movement ended. This pattern was remarkably similar regardless of  
325 cycling direction (*red* and *green* traces largely overlap). Activity in this dimension behaved  
326 similarly for single trials (**Fig. 6b**).

327 We used a Hidden Markov Model (HMM)<sup>38,39</sup> to estimate  $p_{move}$ , which allows the current  
328 estimate to depend on all prior observations. Because those observations must be independent,  
329 we did not use filtered rates (which were used for all other aspects of the decode) but instead  
330 considered spike counts in non-overlapping bins, projected onto the moving-sensitive  
331 dimension. Figure 6c plots the resulting distributions when stopped (*orange*) and moving (*blue*).  
332 These overlapped modestly, a result of the narrow (10 ms) bin. The estimate of  $p_{move}$  is robust  
333 to this overlap because the HMM leverages the full history of spike counts; it can ignore brief  
334 weak evidence for moving while still transitioning swiftly given strong evidence. During BMI  
335 control,  $p_{move}$  (**Fig. 6d**, *blue*) was near typically unity during intended movement (i.e., when the  
336 monkey was actually cycling, *black*) and near zero otherwise.

337 State transitions were determined by  $p_{move}$  (**Fig. 3**). Entering a state that produced virtual  
338 movement (EARLY or STEADY) required that  $p_{move}$  exceed 0.9 and remain consistently above  
339 0.1 for 175 ms. This conservative strategy led to a very low rate of false starts (~2 per day for  
340 monkey G and ~1 every ten days for monkey E). The transition to EARLY (**Fig. 6d**, left edge of  
341 *gray regions*) occurred on average 117 and 194 ms after physical movement onset (monkeys G  
342 and E). Trial-to-trial variability around these mean values was modest: standard deviations  
343 were 93 and 138 ms (computed within session and averaged across sessions). As discussed  
344 above, estimated stopping time (when  $p_{move}$  dropped below 0.1) was also decoded with only  
345 modest trial-to-trial variability.





**Figure 6. Leveraging the moving-sensitive dimension to infer probability of moving. (a)** Trial-averaged population activity, during a manual-control block, projected onto the moving-sensitive dimension (same session and trials as Figure 5a). **(b)** As in panel (a), but for three example single trials (same trials as in Figure 5b). **(c)** Histogram of the neural state projected onto the moving-sensitive dimension for training data. The neural state was measured every ten milliseconds, at times when the monkey was stopped within a target (*orange*) or actively cycling (*blue*). Traces show Gaussian fits used to compute  $p_{move}$ . **(d)** Example time-course, during BMI control, of  $p_{move}$  (*blue*) and the active state (*magenta*). Gray regions show times when the decoder produced virtual movement (i.e., when in EARLY or STEADY). These times corresponded well to times when the monkey was intending to move, as indicated by the angular velocity of the disconnected pedal (*black*). Note also that transient inappropriate spikes in  $p_{move}$  (as seen here around 18 s) do not lead to false starts because either they don't exceed 0.9, as was the case here, or they are too brief and the EARLY state is never reached. Same example data as in Figure 5c.


### 346 *Inferring initial movement direction*

347 Angular momentum of the neural state in the forward and backward planes became substantial  
348 a few hundred milliseconds after  $p_{move}$  became high. Thus, the EARLY state became active  
349 before the direction of movement could be inferred from the elliptical trajectories. To overcome  
350 this problem, we identified three dimensions in which the neural state, around the time of  
351 movement onset, distinguished between forward and backward movement. The neural state in  
352 these dimensions (two of which are shown) differed between forward and backward  
353 movements (*green* and *red* traces) both in the average response (**Fig. 7a**) and on individual trials  
354 (**Fig. 7b**). This difference began to grow just prior to physical movement onset (*dark portion of*  
355 *trajectory shows -200 to +175 ms relative to detected movement onset*) and became less  
356 prominent later in the movement (*light portion of trajectory*). We found these dimensions by  
357 performing PCA on training data (*Methods*). For each of the 50 training trials, we considered the  
358 neural state in these dimensions, measured 175 ms after decoded movement onset. We fit  
359 Gaussian distributions separately for forward (**Fig. 7c, green**) and backward (*red*) trials. These  
360 had largely non-overlapping distributions.

361 During BMI control, upon transition from INIT to EARLY, we computed the likelihood of the  
362 neural state under each distribution. A simple winner-take-all computation determined the  
363 direction of virtual velocity during the EARLY state. The inference of movement direction  
364 during EARLY was correct on 94% and 82% of trials (monkeys G and E). After 200 ms, the  
365 STEADY state was entered and virtual velocity was controlled thereafter by activity in the  
366 rotational dimensions. **Figure 7d** illustrates moments (*colored regions*) where the EARLY state  
367 was active and the above strategy was used to decode virtual motion (physical pedal velocity is  
368 shown for reference). These moments were brief, and had a very modest effect on the overall  
369 time to reach the target. However, we still employed this strategy because our goal was to build  
370 a BMI decode that closely tracked intended movement and felt responsive to the subject.

### 371 *Speed control*

372 The excellent performance of the decoder was aided by the relative simplicity of behavior: when  
373 monkeys moved, they did so at a stereotyped speed. This allowed us to concentrate on building  
374 a decode algorithm that decoded intended direction with accurate timing, and remained  
375 stationary if movement was not intended. However, that decode provided only limited control  
376 of movement speed. An obvious extension is to allow finer-grained speed control. This would  
377 presumably be desired by users of an ego-motion prosthetic. Furthermore, speed control  
378 provides one possible way of steering: e.g., by decoding the relative intensity of intended  
379 movement on the two sides of the body. While we do not attempt that here, we still considered



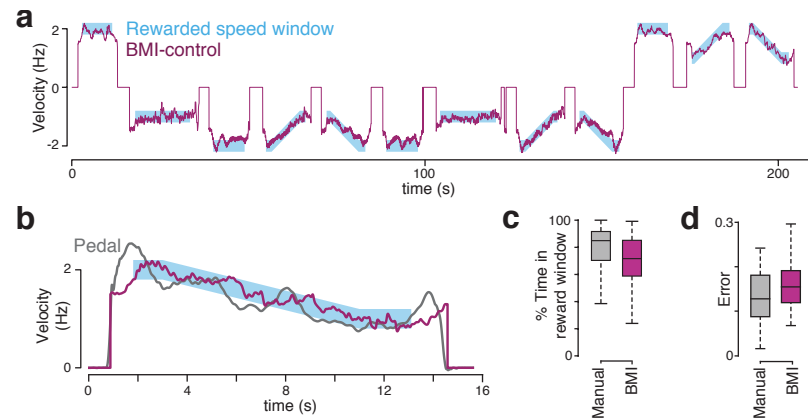
**Figure 7. Leveraging initial-direction dimensions to allow low-latency decoding.** (a) Trial-averaged population activity, during a manual-control block, projected onto two (of three) initial-direction dimensions (same session and trials as Figure 5a and 6a). Boldly colored portions of traces highlight -200 ms to +175 ms relative to physical move onset. Arrows indicate direction of trajectories. (b) As in panel (a), but for three example single trials (same trials as in Figure 5b and 6b). (c) The location of the neural state, for training data, at the time the state-machine (applied post-hoc to that training data) entered the EARLY state. This data (50 total trials) was used to fit two Gaussian distributions. During BMI control, when the EARLY state was entered, virtual direction was determined by which distribution maximized the data likelihood. (d) Example of initial-direction decoding during BMI control. Colored windows show the times in the EARLY state, with red and green indicating decoded direction. Same example data as in Figure 5c and 6d.

380 it important to determine whether the neural features we identified could support speed  
381 control.

382 That assessment required a task where speed control is necessary for success. We thus trained  
383 one monkey to track various speed profiles as he progressed through the virtual environment.  
384 Two floating targets were rendered in the foreground as the monkey cycled. The distance  
385 between them reflected the difference between actual and instructed speed. Obtaining juice  
386 required aligning the two floating targets while progressing towards a final target, on which he  
387 stopped to obtain additional reward. The task was divided into trials, each of which required  
388 moving a distance equivalent to twenty cycles under manual control. We used eight trial-types,  
389 four each for forward and backward cycling. Two of these employed a constant target speed  
390 (equivalent to 1 or 2 Hz cycling) and two involved a ramping speed (from 1 Hz to 2 Hz or vice  
391 versa). As above, the decoder was trained based on a small number of manual-control trials  
392 performed at the beginning of each session. Blocks of manual-control trials were also included  
393 for comparisons between manual and BMI-based performance.

394 Our decode strategy was largely preserved from that described above. However, we used a  
395 modified state machine (**Supp. Fig. 1**) and a slightly different algorithm for transforming  
396 rotations of the neural state into decoded virtual velocity. Direction was determined based on  
397 which distribution (forward or backward) produced the higher likelihood of observing the  
398 measured angular momentums (as in **Fig. 5d**). Once that choice was made, speed was  
399 determined by the angular velocity of the neural state in that plane. Thus, faster rotational  
400 trajectories led to faster decoded virtual velocity. We chose a scaling factor so that a given  
401 neural angular velocity produced the speed that would have been produced by physical cycling  
402 at that angular velocity. Neural angular velocity was exponentially filtered with a time constant  
403 of 500 ms. The filter memory was erased on entry into a movement state (EARLY or STEADY)  
404 from a stopped state (INIT or EXIT) to allow brisk movement onset (see *Methods*).

405 The above strategy allowed smooth BMI control of movement speed. In fact, it tended to give  
406 BMI control an intrinsic advantage over manual control. In manual control, the angular velocity  
407 of the pedal was naturally modulated within each cycle (being higher on the downstroke),  
408 resulting in a fluctuating virtual velocity. Such fluctuations mildly impaired the ability to match  
409 target speed under manual control. To allow a fair comparison, we thus also applied an  
410 exponential filter to virtual velocity under manual control. Filters were chosen separately for  
411 BMI ( $\tau = 500$  ms) and manual control ( $\tau = 1000$  ms) to maximize performance. This was done  
412 informally, in the earliest session, by lengthening the filter until success rate roughly plateaued.  
413 The filter then remained fixed for all further sessions.



**Figure 8. Performance in the modified task requiring speed tracking. (a)** Instructed velocity and BMI-decoded virtual velocity during 12 contiguous trials of BMI control. **(b)** Expanded view of one example trial (the last trial from panel a). The virtual velocity that would have been produced by the pedal is shown in gray for comparison. **(c)** Percentage of time spent in rewarded velocity window for trials in manual-control (2 sessions, 333 trials) and BMI-control (3 sessions, 349 trials). Center lines indicate median, the box edges indicate the first and third quartiles, and the whiskers include all non-outlier points (points less than 1.5 times the interquartile range from the box edges). **(d)** Mean absolute error (MAE) between instructed velocity and virtual velocity for both manual control and BMI control sessions. One mean error was computed per trial. Same format as (c).

414 Under BMI control, decoded virtual speed closely tracked instructed speed. This was true  
415 across trials with different constant speeds, and within trials where speed modulated with time  
416 (**Fig. 8a,b**). To compare BMI with manual control (which were performed on separate days) we  
417 considered all trials where the monkey completed the portion of the trial that required matching  
418 speed (87% of trials in arm control, and 79% in BMI control). The monkey was able to match  
419 instructed speed nearly as accurately under BMI control as under manual control. This was true  
420 judged both by time within the rewarded speed window (**Fig. 8c**) and by the error between  
421 virtual and instructed velocity (**Fig. 8d**).

## 422 Discussion

423 We have argued that the largest signals in motor cortex are not ‘representational’ – they do not  
424 encode variables but are instead essential for noise-robust dynamics. Those dynamics produce  
425 outgoing commands that *are* representational (they covary with the variables they control) but  
426 are low-variance. This perspective argues that decoders should not attempt to invert encoding  
427 unless a great many neurons can be recorded. Instead, decoders should opportunistically  
428 leverage whichever high-variance response features have a robust relationship with the  
429 variables one wishes to decode. In retrospect, traditional reach-based prosthetics can be seen as  
430 implicitly taking this approach; during reaching there exist high-variance neural signals that  
431 correlate (linearly) with a projection of two-dimensional reach velocity onto a ‘preferred  
432 direction’. Similarly, decoding of muscle activity for prosthetic control<sup>6</sup> likely leverages signals  
433 that coincidentally but usefully correlate with muscle force during the task of interest.

434 The need for an opportunistic strategy becomes explicit during cycling, because the neural  
435 signals that linearly correlate with hand kinematics and muscle activity are low variance.  
436 Conversely, there exist high-variance signals with robust (but not linear) relationships with  
437 intended movement. Leveraging those features yielded BMI control that was sufficiently  
438 natural that monkeys appeared not to notice that the task was no longer under manual control.

439 Our approach relates to recent studies that modeled neural dynamics to improve online<sup>21</sup> or  
440 offline<sup>22,41,42</sup> decodes of movement kinematics. A key insight of those studies is that signals that  
441 do not correlate directly with kinematics can be used to infer those that do. For linear decoding,  
442 the value of given variable depends upon the neural state in one dimension: the dimension  
443 defined by the regression weights. Nevertheless, inferring the neural state in that dimension  
444 may benefit from a dynamical model that spans multiple dimensions. Much like the present  
445 approach, this allows the decode to leverage features that are robust, even if they do not directly  
446 correlate with the kinematic parameters of interest. The present approach extends this idea to  
447 situations where there may be no high-variance dimensions that can be linearly decoded,  
448 and/or where the most prominent features are not well-described by linear dynamics.

449 By most measures (success rate, time to target) performance under BMI control was remarkably  
450 close to that under manual control. The only measure by which BMI control was meaningfully  
451 inferior to manual control was detection of movement offset. Yet while not as accurate as under  
452 manual control, stopping accuracy under BMI control was still good: stopping was detected  
453 with ~100 ms precision. Furthermore, BMI control almost never produced movement when it  
454 was not intended. False starts occurred at the rate of ~2 per day for monkey G and ~1 every ten

455 days for monkey E. We consider this a particularly important attribute of any ego-motion  
456 decoding algorithm, due to the potentially large consequences of unintended movement of the  
457 whole body. The first demonstrations of BMI control of ego-motion<sup>24,25</sup>, which leveraged  
458 strategies originally developed for reaching, did not demonstrate accurate stopping; stopping  
459 occurred automatically based on target proximity. The classifier used by Libedinsky et al.<sup>25</sup>  
460 successfully decoded one-of-four commands (turn-left, turn-right, stop, forward) on ~80% of  
461 times. Rajangam et al.<sup>24</sup> employed a whole-body direction decode, which allowed monkeys to  
462 navigate ~2 meters to a target in two-dimensional space (which had to be approached with an  
463 accuracy of +/- ~0.2 meters, or 10% of the distance traveled) in an average of 27-49 seconds  
464 (depending on the monkey and degree of practice). Direct quantitative comparison of  
465 performance of our decode with prior work is essentially impossible as the tasks are so  
466 different. Our task is more challenging in some ways (it is much faster paced and requires  
467 accurate stopping) but less so in others (it did not require turning). The nature of the underlying  
468 neural activity (rhythmic versus not) is also very different. Furthermore, there are no clear  
469 benchmarks for comparison. The major advance of Rajangam et al.<sup>24</sup> is the demonstration of  
470 wheelchair-like control using a physical device, but the unconstrained paths do not afford the  
471 opportunity to ascertain how well decoding tracked intent. What can be said with certainty is  
472 that, for the task we used, the opportunistic approach was both successful (leading to near-  
473 native performance) and necessary (signals that correlate linearly with kinematics were very  
474 weak). Opportunistic decode strategies should thus be considered as BMI control applications  
475 expand beyond reach-like tasks.

476

477 An obvious limitation of the current experiments is that we did not explore strategies for  
478 steering, which would be essential to a real-world ego-motion prosthetic. There exist multiple  
479 candidate strategies for enabling steering. Rajangam et al. used a Wiener filter to decode  
480 angular velocity of the body. While straightforward, this strategy appears to have had limited  
481 success: even during training, the  $R^2$  of their angular velocity decode was 0.16 and 0.12 for the  
482 two monkeys. During online performance, the considerable time to reach the target argues that  
483 steering was not accurate. One alternative strategy would be to apply our decode strategy  
484 bilaterally, and employ a comparison (e.g., between left and right cycling speed) to control  
485 angular velocity. Another strategy would be to control translational velocity using the strategies  
486 developed here, but use a reach-like decode for steering (rather like pedaling a bicycle while  
487 also steering). Which (if any) of these three strategies is preferable remains a question for future  
488 experiments.



489 For convenience, monkeys were trained to control the pedal with their forelimb (allowing them  
490 to be seated in a traditional primate chair) and we thus recorded from the forelimb region of  
491 motor cortex. Recordings during natural locomotion in monkeys reveal broadly similar signals  
492 in both the forelimb<sup>43</sup> and hindlimb<sup>27</sup> regions of motor cortex. These signals are dominated, as in  
493 our task, by elliptical neural trajectories during ongoing locomotion. It is thus likely that  
494 prosthetic ego-motion could be driven by signals derived from either region.

495 Like many proof-of-concept prosthetic systems developed in primates<sup>1,14,44-46</sup>, decoder training  
496 depended on observations of neural activity under manual control. Whether this approach  
497 translates depends upon the assumption that useful patterns of neural activity will emerge  
498 when a paralyzed patient tries to move but can't. A number of existing studies indicate that  
499 motor cortex is active, in reasonably normal ways, when paralyzed patients attempt to move.  
500 Importantly, decode strategies based on a characterization of population activity during normal  
501 reaching in primates have provided successful directional control of a cursor in human patients.  
502 The approach to specifying parameters was of course tailored to the needs of the patients, but  
503 the class of decoder did not need to be altered. It seems likely that the same will be true of  
504 rhythmic neural activity and decoding of intended movement. More broadly, a key point of the  
505 present study is that a fixed decode strategy is unlikely to work well across the different classes  
506 of movement that patients are likely to desire. Instead, decode strategies should leverage  
507 population-level response features that relate robustly to the variables one wishes to control.  
508 Such features are presumably present in paralyzed patients, and will almost certainly be task  
509 dependent.

510 Our results indicate that a nonlinear, yet relatively simple, decode strategy can afford excellent  
511 one-dimensional control of ego-motion. Although other approaches remain possible<sup>47</sup>, our  
512 findings support the idea that cortical control of prosthetic ego-motion is viable and should be  
513 explored further. More broadly, the present results argue that many of the decode strategies  
514 that proved effective for reach-based prostheses are unlikely to generalize across tasks. An  
515 alternative approach is to identify, for each task, the dominant features and determine how they  
516 might be usefully translated into decoded movement. While this approach abandons the elegant  
517 idea of inverting a literal encoding of kinematics, it opens up possibilities for improved  
518 prosthetic control across a variety of contexts.

## 519 **Methods**

### 520 *Subjects and primary task*

521 All procedures were approved by the Columbia University Institutional Animal Care and Use  
522 Committee. Subjects G and E were two adult male macaque monkeys (*Macaca mulatta*).  
523 Monkeys sat in a primate chair facing an LCD monitor (144 Hz refresh rate) that displayed a  
524 virtual environment generated by the Unity engine (Unity Technologies, San Francisco, CA).  
525 The head was restrained via a titanium surgical implant. While the monkey's left arm was  
526 comfortably restrained, the right arm grasped a hand pedal. Cloth tape was used to ensure  
527 consistent placement of the hand on the pedal. The pedal connected via a shaft to a motor  
528 (Applied Motion Products, Watsonville, CA), which contained a rotary encoder that measured  
529 the position of the pedal with a precision of 1/10,000 of the cycle. The motor was also used to  
530 apply forces to the pedal, endowing it with virtual mass and viscosity.

531 Manual-control sessions for the primary cycling task required that the monkey cycle the pedal  
532 in the instructed direction to move through the virtual environment, and stop on top of a  
533 lighted target to collect juice reward. The color of the landscape indicated whether cycling must  
534 be 'forward' (green landscape, the hand moved away from the body at the top of the cycle) or  
535 'backward' (tan landscape, the hand moved toward the body at the top of the cycle). There were  
536 6 total conditions, defined by cycling direction (forward or backward) and target distance (2, 4,  
537 or 7 cycles). Distance conditions were randomized within same-direction blocks (3 trials of each  
538 distance per block), and directional blocks were randomized over the course of each  
539 experiment. Trials began with the monkey stationary on a target. A second target appeared in  
540 the future. To obtain reward, the monkey had to cycle to that target, come to a halt 'on top' of it  
541 (in the first-person perspective of the task) and remain stationary for a hold period of 1000-1500  
542 ms (randomized). A trial was aborted without reward if the monkey began moving before  
543 target onset (or in the 170 ms after, which would indicate attempted anticipation), if the monkey  
544 moved past the target without stopping, or if the monkey moved while awaiting reward. The  
545 next trial began 100 ms after the variable hold period. Monkeys performed until they received  
546 enough liquid reward that they chose to desist. As their motivation waned, they would at times  
547 take short breaks. For both manual-control and BMI-control sessions, we discarded any trials in  
548 which monkeys made no attempt to initiate the trial, and did not count them as 'failed'. These  
549 trials occurred  $2 \pm 2$  times per session (mean and standard deviation, Monkey G, maximum 10)  
550 and  $3 \pm 3$  times per session (Monkey E, maximum 11).

551 In BMI control, trial parameters and failure conditions were the same as in manual control, for  
552 purposes of comparison. The only difference between manual and BMI control was that, in the  
553 latter, position in the virtual environment was controlled by the output of a decoder rather than  
554 the pedal. We did not prevent or discourage the monkey from cycling during BMI-control  
555 blocks, and he continued to do so as normal. In BMI control, monkey G performed an average  
556 of 654 trials/session over 20 sessions and monkey E performed an average of 137 trials/session  
557 over 17 sessions. Manual-control data for monkey G (average of 229 trials/session over 8  
558 sessions) were collected during sessions in which BMI-control data sets were also collected.

559 Manual-control sessions for monkey E (average of 231 trials/session over 5 sessions) were  
560 interleaved with BMI-control sessions on different days. For monkey G, an additional three  
561 manual-control sessions (189, 407, and 394 trials) were employed to record EMG, which was  
562 used for the variance captured analysis (**Fig. 2b,c**). We recorded from 5-7 muscles per session,  
563 yielding a total of 19 recordings. We made one or more recordings from the three heads of the  
564 *deltoid*, the lateral and long heads of *triceps brachii*, the *biceps brachii*, *trapezius*, and *latissimus*  
565 *dorsi*. These muscles were selected due to their clear activations during the cycling task.

#### 566 *Surgery and neural/muscle recordings*

567 Neural activity was recorded using chronic 96-channel Utah arrays (Blackrock Microsystems,  
568 Salt Lake City, UT), implanted in the left hemisphere using standard surgical techniques. In  
569 each monkey, an array was placed in the region of primary motor cortex (M1) corresponding to  
570 the upper arm. In monkey G, a second array was placed in dorsal premotor cortex (PMd), just  
571 anterior to the first array. Array locations were selected based on MRI scans and anatomical  
572 landmarks observed during surgery. Experiments were performed 1-8 months (monkey G) and  
573 3-4 months (monkey E) after surgical implantation. Neural responses both during the task and  
574 during palpation confirmed that arrays were in the proximal-arm region of cortex.

575 Electrode voltages were filtered (band-pass 0.3 Hz – 7.5 kHz) and digitized at 30 kHz using  
576 Digital Headstages, Digital Hubs, and Cerebus Neural Signal Processors from Blackrock  
577 Microsystems. Digitized voltages were high-pass filtered (250 Hz) and spike events were  
578 detected based on threshold crossings. Thresholds were set to between -4.5 and -3 times the  
579 RMS voltage on each channel, depending on the array quality on a given day. On most  
580 channels, threshold crossings included clear action-potential waveforms from one or more  
581 neurons, but no attempt was made to sort action potentials.

582 Intra-muscular EMG recordings were made using pairs of hook-wire electrodes inserted with 30  
583 mm x 27 gauge needles (Natus Neurology, Middleton, WI). Raw voltages were amplified and  
584 filtered (band-pass 10 Hz – 10 kHz) with ISO-DAM 8A modules (World Precision Instruments,  
585 Sarasota, FL), and digitized at 30 kHz with the Cerebus Neural Signal Processors. EMG was  
586 then digitally band-pass filtered (50 Hz – 5 kHz) prior to saving for offline analysis. Offline,  
587 EMG recordings were rectified, low-pass filtered by convolving with a Gaussian (standard  
588 deviation: 25 ms), downsampled to 1 kHz, and then fully normalized such that the maximum  
589 value achieved on each EMG channel was 1.

590 A real-time target computer (Speedgoat, Bern, CH) running Simulink Real-Time environment  
591 (MathWorks, Natick, MA) processed behavioral and neural data and controlled the decoder  
592 output in online experiments. It also streamed variables of interest to another computer that  
593 saved these variables for offline analysis. Stateflow charts were implemented in the Simulink  
594 model to control task state flow as well as the decoder state machine. Real-time control had  
595 millisecond precision.

596 Spike trains were causally converted to firing rates by convolving each spike with a beta kernel.  
597 The beta kernel was defined by temporally scaling a beta distribution (shape parameters:  $\alpha = 3$

598 and  $\beta = 5$ ) to be defined over the interval  $[0, 275]$  ms and normalizing the kernel such that the  
599 firing rates would be in units of spikes/second. The same filtering was applied for online  
600 decoding and offline analyses. Firing rates were also mean centered (subtracting the mean rate  
601 across all times and conditions) and normalized. During online decoding, the mean and  
602 normalization factor were values that had been computed from the training data. We used soft  
603 normalization<sup>16</sup>: the normalization factor was the firing rate range plus a constant (5 spikes/s).

#### 604 *Computing trial-averaged firing rates*

605 Analyses of BMI performance are based on real-time decoding during online performance, with  
606 no need to consider trial-averaged firing rates. However, we still wished to compute trial-  
607 averaged traces of neural activity and kinematics for two purposes. First, some aspects of  
608 decoder training benefited from analyzing trial-averaged firing rates. Second, we employ  
609 analyses that document basic features of single-neuron responses and of the population  
610 response (e.g., **Fig. 1d**, **Fig. 2**, **Fig. 5a**, **Fig. 6a**). These analyses benefit from the denoising that  
611 comes from computing a time-varying firing rate across many trials. Due to the nature of the  
612 task, trials could be quite long (up to 20 cycles in the speed-tracking task), rendering the  
613 traditional approach of aligning all trials to movement onset insufficient for preserving  
614 alignment across all subsequent cycles. It was thus necessary to modestly adjust the time-base  
615 of each individual trial (e.g., stretching time slightly for a trial where cycling was faster than  
616 typical). We employed two alignment methods. Method A is a simplified procedure that was  
617 used prior to parameter fitting when training the decoder before online BMI control. This  
618 method aligns only times during the movement. Method B is a more sophisticated alignment  
619 procedure that was utilized for all offline analyses. This method aligns the entire trial, including  
620 pre- and post-movement data. For visualization, conditions with the same target distance (e.g.,  
621 7 cycles), but different directions, were also aligned to the same time base. Critically, any data  
622 processing that relied on temporal structure was completed in the original, unstretched time  
623 base prior to alignment.

624 **Method A:** The world position for each trial resembles a ramp between movement onset and  
625 offset (**Fig. 1a**). First, we identify the portion of each trial starting  $\frac{1}{4}$  cycle into the movement  
626 and ending  $\frac{1}{4}$  cycle before the end of the movement. We fit a line to the world position in this  
627 period and then extend that line until it intercepts the starting and ending positions. The data  
628 between these two intercepts is considered the movement data for each trial and is extracted.  
629 This movement data is then uniformly stretched in time to match the average trial length for  
630 each trial's associated condition. This approach compresses slower than average movements  
631 and stretches faster than average movements within a condition, such that they can be averaged  
632 while still preserving many of the cycle-specific features of the data.

633 **Method B:** This method consists of a mild, non-uniform stretching of time in order to match  
634 each trial to a condition-specific template. For complete details, see Russo et al. 2018.<sup>16</sup>

#### 635 *Variance captured analysis*

636 Analysis of neural variance captured (**Fig. 2**) was based on successful manual-control trials from  
637 the three sessions with simultaneous neural and muscle recordings. We considered data from  
638 the full duration of each trial, including times before movement onset and after movement  
639 offset. We analyzed the variance captured by neural dimensions of three types. First, neural  
640 dimensions where activity correlated strongly with kinematic features. Second, neural  
641 dimensions where activity correlated strongly with muscle activity. Third, neural dimensions  
642 that captured robust ‘features’ leveraged by our decoder.

643 Dimensions of the third type (were found as detailed below in a dedicated section below.  
644 Dimensions of the first two types were found using the model  $z(r, t) = c + \mathbf{w}^\top \mathbf{y}(r, t)$ , where  
645  $z(r, t)$  is the kinematic or muscle variable at time  $t$  during trial  $r$ , and  $\mathbf{y}(r, t)$  is the  
646 corresponding  $N$ -dimensional vector of neural firing rates. The constant  $c$  and the column  
647 vector  $\mathbf{w}$  were found via regression. The vector  $\mathbf{w}$  defines a direction in neural space where  
648 activity correlates strongly with the variable  $z$ . We found multiple such vectors; e.g.  $\mathbf{w}_{x-vel}$  is a  
649 dimension where neural activity correlates with horizontal velocity and  $\mathbf{w}_{biceps}$  is a dimension  
650 where neural activity correlates with biceps activity. All such vectors were scaled to have unity  
651 norm before computing the neural variance captured by that dimension. Regression was based  
652 on single-trial responses because this was intrinsically regularizing. We wished to encourage  
653 regression to find high-variance dimensions if possible, and the use of single-trial data  
654 encouraged it to do so. Because filtering of neural activity introduces a net lag, this analysis  
655 naturally assumes a  $\sim 100$  ms lag between neural activity and the variables of interest. Results  
656 were extremely similar if we considered longer or shorter lags.

657 We wished to compute, for each dimension, the percentage of neural variance explained – i.e.,  
658 whether that dimension captured large or small signals. We were not interested in whether  
659 dimensions captured stochastic spiking variability, but in whether they captured large features  
660 that were reliable across trials. Thus, variance captured was always computed based on trial-  
661 averaged neural responses. We considered the matrix  $\bar{Y} \in \mathbb{R}^{N \times T}$  where  $T$  is the total number of  
662 time points across all conditions. Each row of  $\bar{Y}$  contains the trial-averaged firing rate of one  
663 neuron. We computed an  $N \times N$  covariance matrix  $\Sigma = \text{cov}(\bar{Y})$  by treating rows of  $\bar{Y}$  as random  
664 variables and columns as observations. The proportion of total neural variance captured by a  
665 given dimension,  $\mathbf{w}$ , is therefore:

$$666 \quad \frac{\mathbf{w}^\top \Sigma \mathbf{w}}{\text{tr}(\Sigma)}$$

667 Some analyses considered the variance captured by a subspace spanned by a set of dimensions.  
668 To do so we took the sum of the variance captured by orthonormal dimensions spanning that  
669 space.

### 670 *Identifying neural dimensions*

671 Although the response features leveraged by the decode algorithm are clearly visible in the top  
672 principal components of the data (when PCA is performed on the full trial-averaged time-series

673 of firing rates across conditions), we sought to choose neural dimensions that would cleanly  
674 isolate particular features. To this end, each feature was isolated using dedicated preprocessing  
675 and dimensionality reduction approaches.

676 We sought a moving-sensitive dimension, the projection onto which would allow an HMM to  
677 estimate the probability of moving,  $p_{move}$ , at each moment. To do so, we first computed binned  
678 spike-counts (10 ms for monkey G, 20 ms for monkey E) and applied a square-root transform to  
679 these counts as this has been shown to improve the Gaussian fit for Poisson data with small  
680 counts<sup>48</sup>. We then aggregated all of these square-rooted binned counts from the training set (25  
681 forward trials, 25 backward trials) and separated them into two classes based on pedaling  
682 speed: ‘moving’ (speed > 1 Hz) and ‘stopped’ (speed < .05 Hz). Samples that didn’t fall into  
683 either of these two classes were discarded. We applied linear discriminant analysis to these two  
684 labeled sets, which yielded a discriminating hyperplane that best separated the two classes. We  
685 defined the moving-sensitive dimension,  $\mathbf{w}_{move}$ , as the vector normal to this hyperplane.

686 In order to decode direction, we sought to isolate four neural dimensions that captured  
687 rotational trajectories during steady-state cycling. Spike time-series were filtered to yield firing  
688 rates (as described above), and then further high-pass filtered (2<sup>nd</sup> order Butterworth, cutoff  
689 frequency: 1 Hz). This removed drift or other low-frequency signals. Single-trial movement-  
690 period responses were then aligned (Method A) and averaged within conditions to generate  
691  $N \times T_c$  matrices  $\bar{Y}_f$  and  $\bar{Y}_b$ . We sought a 4-dimensional projection of these trial-averaged  
692 responses that would maximally capture rotational trajectories while segregating forward and  
693 backward data into different planes. Whereas the standard PCA cost function finds dimensions  
694 that maximize variance captured, we opted instead for a cost function that would maximize the  
695 difference in variance captured between the two conditions:

$$696 \quad J(W) = \text{tr}(W^T \Sigma_f W) - \text{tr}(W^T \Sigma_b W)$$

697 where  $\Sigma_f = \text{cov}(\bar{Y}_f)$ ,  $\Sigma_b = \text{cov}(\bar{Y}_b)$ ,  $W$  is constrained to be orthonormal. Note that this cost  
698 function will be maximized when the projection of the data captures a great deal of variance for  
699 forward trials and very little variance for backward trials. Conversely, this cost function will be  
700 minimized when the projection favors large variances for backward trials and small variances  
701 for forward trials. We thus chose to define our forward rotational plane by the 2D matrix  $W_f =$   
702  $\begin{bmatrix} \mathbf{w}_f^{(1)} & \mathbf{w}_f^{(2)} \end{bmatrix}$  that maximizes  $J(W)$  and our backward rotational plane by the 2D matrix  $W_b =$   
703  $\begin{bmatrix} \mathbf{w}_b^{(1)} & \mathbf{w}_b^{(2)} \end{bmatrix}$  that minimizes  $J(W)$ . An iterative optimization procedure was used to find  $W_f$  and  
704  $W_b$ ; full details of this in<sup>49</sup>.

705 To decode direction during the EARLY state, we found a set of initial-direction dimensions. We  
706 used activity in the moving-sensitive dimension to determine the time,  $t_{init}$ , at which the state  
707 machine would have entered the INIT state during online control. We then considered trial-  
708 averaged neural activity, for each condition, from  $t_{init}$  through  $t_{init} + 200$  ms. We applied PCA  
709 and retained the top three dimensions:  $\mathbf{w}_{dir}^{(1)}$ ,  $\mathbf{w}_{dir}^{(2)}$ , and  $\mathbf{w}_{dir}^{(3)}$ . Such dimensions capture how

710 activity evolves both across that timespan, and how it differs across forwards and backwards  
711 cycling conditions.

712 *Computing probability of moving ( $p_{move}$ )*

713 To compute  $p_{move}$  based on neural activity in the moving-sensitive dimension, an HMM was  
714 used to track two states: ‘moving’ or ‘stopped’<sup>39</sup>. Square-rooted spike counts in the training data  
715 were already separated into ‘moving’ and ‘stopped’ sets for the purposes of identifying  $w_{move}$ .  
716 We projected those counts onto  $w_{move}$  and a fit Gaussian distribution for each state. The  
717 probability,  $p_{move}$ , of being in the ‘moving’ state, given the entire sequence of current and  
718 previously observed square-rooted spike counts, was computed efficiently with a recursive  
719 algorithm that uses the state transition matrix

$$720 \quad \Phi = \begin{bmatrix} p_{move|move} & p_{move|stop} \\ p_{stop|move} & p_{stop|stop} \end{bmatrix}$$

721 and knowledge of the Gaussian distributions.  $\Phi$  encodes prior assumptions about the  
722 probability of transitioning from one state to the next at any given bin. We used a benchmark  
723 set of manual-control training data from each monkey to determine reasonable values for  $\Phi$ ,  
724 which were then used in all experiments. For monkey G, we set  $p_{move|stop} = .0001$  and  
725  $p_{stop|move} = .002$ ; for monkey E, we set  $p_{move|stop} = .0002$  and  $p_{stop|move} = .004$ . The value  
726  $p_{move}$  was used throughout the decoder state machine to control transitions between various  
727 states, effectively dictating the movement onset and offset behavior of the decoder (**Fig. 6d**).

728 *Computing steady-state direction and speed*

729 Projecting single-trial, high-pass filtered firing rates onto the rotational planes spanned by  $W_f$   
730 and  $W_b$  yielded trajectories that differed considerably between forward and backward  
731 conditions. To further denoise these state trajectories we applied a Kalman filter of the form

$$732 \quad \begin{aligned} x_t &= Ax_{t-1} + q_t \\ 733 \quad y_t &= Cx_t + r_t \end{aligned}$$

734 where  $q_t \in \mathcal{N}(0, Q)$ , and  $r_t \in \mathcal{N}(0, R)$ . In these equations,  $x_t$  represents the true underlying  
735 neural state in the rotational dimensions and  $y_t$  are the high-pass filtered firing rates, which we  
736 treat as noisy measurements of that underlying state. We chose to let our measurements be  
737 smooth firing rates, rather than use non-overlapping bins of spikes, for purely opportunistic  
738 reasons: it consistently yielded better performance by our decoder. The parameters of the  
739 Kalman filter were fit to the training data as follows:

$$740 \quad A = \bar{X}_2 \bar{X}_1^T (\bar{X}_1 X_1^T)^{-1}$$

$$741 \quad C = \begin{bmatrix} W_f^T \\ W_b^T \end{bmatrix}^\dagger$$

742 
$$Q = \text{cov}(\bar{X}_2 - A\bar{X}_1)$$

743 
$$R = \text{cov}(Y - C\bar{X})$$

744 where

745 
$$\bar{X}_1 = \begin{bmatrix} W_f^T \\ W_b^T \end{bmatrix} [\bar{Y}_f(:, 1: T_f - 1), \bar{Y}_b(:, 1: T_b - 1)]$$

746 
$$\bar{X}_2 = \begin{bmatrix} W_f^T \\ W_b^T \end{bmatrix} [\bar{Y}_f(:, 2: T_f), \bar{Y}_b(:, 2: T_b)]$$

747 
$$\bar{X} = \begin{bmatrix} W_f^T \\ W_b^T \end{bmatrix} [\bar{Y}_1, \bar{Y}_2, \dots, \bar{Y}_{50}]$$

748 
$$Y = [Y_1, Y_2, \dots, Y_{50}]$$

749 with  $Y_i$  denoting the neural activity (high-pass filtered firing rates) for the  $i$ -th trial in the  
 750 training set,  $\bar{Y}_i$  denoting the trial-averaged activity for the condition that the  $i$ -th trial is an  
 751 instantiation of,  $\dagger$  denoting the Moore-Penrose pseudoinverse, and the colon symbol  
 752 designating how to index matrices (e.g.,  $M(:, a: b)$  refers to the submatrix of  $M$  including all  
 753 rows of  $M$ , but only the columns  $a$  through  $b$ ). Lastly, the initial state parameter  $x_0$  was  
 754 computed by taking the average value of the trial-averaged projections over all times and  
 755 conditions. Online inference of the underlying neural state, which yields an estimate  $\hat{x}_t$  at each  
 756 millisecond  $t$ , was computed recursively using the steady-state form of the Kalman filter<sup>50</sup>.

757 After denoising the neural state in the rotational dimensions via the Kalman filter, angular  
 758 momentum was computed in each plane as the cross product between the estimated neural  
 759 state and its derivative, which (up to a constant scaling) can be written

760 
$$L(t) = \begin{bmatrix} L_f(t) \\ L_b(t) \end{bmatrix} = \begin{bmatrix} \hat{x}_{t-1}^{(1)} \hat{x}_t^{(2)} - \hat{x}_t^{(1)} \hat{x}_{t-1}^{(2)} \\ \hat{x}_{t-1}^{(3)} \hat{x}_t^{(4)} - \hat{x}_t^{(3)} \hat{x}_{t-1}^{(4)} \end{bmatrix}$$

761 where the superscript indexes the elements of  $\hat{x}_t$ . We fit 2D Gaussian distributions to these  
 762 angular momentums for each of three behaviors in the training data: ‘stopped’ (speed < .05 Hz),  
 763 ‘pedaling forward’ (velocity > 1 Hz), and ‘pedaling backward’ (velocity < -1 Hz) (**Fig. 5d**).  
 764 Online, the likelihood of the observed angular momentums with respect to each of these three  
 765 distributions dictated the steady-state estimates of direction and speed. We’ll denote these three  
 766 likelihoods  $f_{stop}$ ,  $f_{forward}$ , and  $f_{backward}$ .

767 In general, one can compute which of these three distributions is most likely by choosing the  
 768 maximizing likelihood and assess confidence in that choice by comparing the relative values of  
 769 the three likelihoods. However, we wanted the decoder to err on the side of withholding  
 770 movement. We therefore set a conservative threshold on  $f_{stop}$  corresponding to the point at  
 771 which  $L$  would have a Mahalanobis distance of 3 to the stopped distribution of angular



772 momentums. If  $f_{stop}$  ever exceeded this threshold, we set  $speed_{steady}$  to zero. If this condition  
773 was not met, we decoded direction and speed as follows:

$$774 \quad direction_{steady}(t) = \text{sgn}\left(f_{forward}(t) - f_{backward}(t)\right)$$

$$775 \quad speed_{steady}(t) = \left| 2 \cdot \frac{f_{forward}(t)}{f_{forward}(t) + f_{backward}(t)} - 1 \right| \beta$$

776 where  $\left| 2 \cdot \frac{f_{forward}(t)}{f_{forward}(t) + f_{backward}(t)} - 1 \right|$  varies between 0 and 1 depending on the relative sizes of  
777 the likelihoods (yielding a slower velocity if the direction decode is uncertain) and  $\beta$  is a  
778 direction-specific constant learned from the training data whose purpose is simply to scale up  
779 the result to match steady-state cycling speed. In practice,  $speed_{steady}$  was frequently very close  
780 to the monkeys' steady-state cycling speeds (**Fig. 5e**).

#### 781 *Computing initial direction and speed*

782 Initial direction and speed were always computed at the moment the EARLY state was entered,  
783  $t_{early}$ . These values then persisted throughout the remainder of the EARLY state. Given that the  
784 decoder state machine doesn't make use of the initial-direction dimensions prior to entering the  
785 EARLY state,  $t_{early}$  can be computed for the training trials. Single-trial firing rates from the  
786 training set were then projected onto the initial-direction dimensions at  $t_{early}$  and 3D Gaussian  
787 distributions were fit to the resulting sets of forward and backward neural states. Online, firing  
788 rates were projected onto the initial-direction dimensions at  $t_{early}$  and likelihoods  $g_{forward}$  and  
789  $g_{backward}$  were computed with respect to each the learned distributions. If the observed neural  
790 state in the initial-direction subspace was not an outlier ( $>10$  Mahalanobis distance units) with  
791 respect to both distributions, then the initial direction and speed were computed as follows:

$$792 \quad direction_{initial}(t_{early}) = \text{sgn}\left(g_{forward}(t_{early}) - g_{backward}(t_{early})\right)$$

$$793 \quad speed_{initial}(t_{early}) = \left| 2 \cdot \frac{g_{forward}(t_{early})}{g_{forward}(t_{early}) + g_{backward}(t_{early})} - 1 \right| \beta$$

794 If the observed neural state was an outlier, initial direction and speed were computed in the  
795 same manner as is done in the STEADY state.

#### 796 *Smoothing of decoded velocity*

797 In the primary experiment, the decoder state machine produced an estimate of velocity,  $v_{dec}$ , at  
798 every millisecond. During the STOP and INIT states, this estimate was zero and the monkey's  
799 position in the virtual environment was held constant. During the EARLY and STEADY states,  
800 this estimate was smoothed with a trailing average:

801 
$$v'_{dec}(t) = \frac{1}{T_{smooth} + 1} \sum_{i=0}^{T_{smooth}} v_{dec}(t - i)$$

802 where  $T_{smooth} = \min(500, t - t_{early})$ , i.e., the trailing average extended in history up to 500 ms  
803 or to the moment the EARLY state was entered, whichever was shorter.  $v'_{dec}$  was integrated  
804 every millisecond to yield decoded position in the virtual environment. In the speed-tracking  
805 experiment (described below) there was no need to smooth of  $v_{dec}$  prior to integration because  
806 the speed estimate had already been smoothed.

### 807 *Speed-tracking task*

808 In addition to the primary task (where the monkey traveled 2-7 cycles between stationary  
809 targets) we employed a speed-tracking task, in which the monkey was required to match his  
810 virtual speed to an instructed speed. Speed was instructed implicitly, via the relative position of  
811 two moving targets. The primary target was located a fixed distance in front of the monkey's  
812 location in virtual space: the secondary target fell 'behind' the first target when cycling was too  
813 slow, and pulled 'ahead' if cycling was too fast. This separation saturated for large errors, but  
814 for small errors was proportional to the difference between the actual and instructed speed.  
815 This provided sufficient feedback to allow the monkey to track the instructed speed even when  
816 it was changing. Because there was no explicit cue regarding the absolute instructed speed,  
817 monkeys began cycling on each trial unaware of the true instructed speed profile and 'honed in'  
818 on that speed over the first ~2 cycles.

819 We quantify instructed speed not in terms of the speed of translation through the virtual  
820 environment (which has arbitrary units) but in terms of the physical cycling velocity necessary  
821 to achieve the desired virtual speed. E.g., an instructed speed of 2 Hz necessitated cycling at an  
822 angular velocity of 2 Hz to ensure maximal reward. Under BMI control, the output of the  
823 decoder had corresponding units. For example, a 2 Hz angular velocity of the neural trajectory  
824 produced movement at the same speed as 2 Hz physical cycling (see '*Neural features for speed-*  
825 *tracking*' for details of decoder). Reward was given throughout the trial so long as the monkey's  
826 speed was within 0.2 Hz of the instructed speed. We employed both constant and ramping  
827 instructed-speed profiles.

828 Constant profiles were at either 1 Hz or 2 Hz. Trials lasted 20 cycles. After 18 cycles, the primary  
829 and secondary targets (described above) disappeared and were replaced by a final stationary  
830 target two cycles in front of the current position. Speed was not instructed during these last two  
831 cycles; the monkey simply had to continue cycling and stop on the final target to receive a large  
832 reward. Analyses of performance (e.g., **Fig. 8c,d**) were based on the ~16 cycle period starting  
833 when the monkey first honed in on the correct speed (within 0.2 Hz of the instructed speed) and  
834 ending when the speed-instructing cues disappeared 2 cycles before the trial's end.

835 Ramping profiles began with three seconds of constant instructed speed to allow the monkey to  
836 hone in on the correct initial speed. Instructed speed then ramped, over 8 seconds, to a new  
837 value, and remained constant thereafter. As for constant profiles, speed-instructing cues

838 disappeared after 18 cycles and the monkey cycled two further cycles before stopping on a final  
839 target. Again, analyses of performance were based on the period from when the monkey first  
840 honed in on the correct speed, to when the speed-instructing cues disappeared. There were two  
841 ramping profiles: one ramping up from 1 to 2 Hz, and one ramping down from 2 to 1 Hz. There  
842 were thus four total speed profiles (two constant and two ramping). These were performed for  
843 both cycling directions (presented in blocks and instructed by color as in the primary task)  
844 yielding eight total conditions. This task was only performed by monkey G, who completed an  
845 average of 166 trials/session over 2 sessions in manual control and an average of 116  
846 trials/session over 3 sessions in BMI control.

847 As will be described below, the speed decoded during BMI control was low-pass filtered to  
848 remove fluctuations due to noise. This had the potential to actually make the task easier under  
849 BMI control, given that changes in instructed speed were slow within a trial (excepting the  
850 onset and offset of movement). We did not wish to provide BMI control with an ‘unfair’  
851 advantage in comparisons with manual control. We therefore also low-pass filtered virtual  
852 speed while under manual control. Filtering (exponential,  $\tau = 1$  second) was applied only when  
853 speed was above 0.2 Hz, so that movement onset and offset could remain brisk. This aided the  
854 monkey’s efforts to track slowly changing speeds under manual control.

855 During training and while under manual control, trials were failed if there was ever a large  
856 discrepancy between actual and instructed speed. This ensured that monkeys tried their best to  
857 consistently match speed at all times. We relaxed this failure mode under BMI control because  
858 we did not wish to mask large failures in decoded speed. Over the course of single sessions, this  
859 did not discourage monkeys from trying their best, but simply allowed us to observe and  
860 quantify decode failures that would otherwise have resulted in aborted trials. This potentially  
861 puts BMI performance – quantified as in **Figure 8c,d** – at a disadvantage relative to manual  
862 control, where large errors could not persist. In practice this was not an issue as large errors  
863 were rare.

#### 864 *Neural features for speed-tracking*

865 Although the speed-tracking experiment leveraged the same dominant neural responses that  
866 were used in the primary experiment, the specific features calculated for the decoder state  
867 machines differed. Details on how the relevant features were calculated in the speed-tracking  
868 experiment are presented in this section.

869 The probability of moving,  $p_{move}$ , was calculated using a different set of parameters for speed-  
870 tracking, largely due to changes in recording quality in the intervening time between data  
871 collection from the primary experiment and data collection for the speed-tracking experiment.  
872 The bin size was increased to 100 ms and the following state transition values were used:  
873  $p_{move|stop} = .0005$  and  $p_{stop|move} = .0005$ . In addition, we observed that the square-root  
874 transform seemed to be having a negligible impact on the quality of the decoder at this bin size,  
875 so we removed it for this task.

876 Several features used in the speed-tracking state machine rely on neural activity in the  
877 rotational dimensions. In the primary experiment, this activity was high-pass filtered (cutoff  
878 frequency: 1 Hz) prior to projection into these dimensions, which helped isolate the rotational  
879 neural trajectories during ~2 Hz cycling. For speed-tracking, we wanted to accommodate a  
880 broader range of cycling speeds (which corresponded to a broader range of periodicities in the  
881 rotational neural trajectories). Thus, we dropped the cutoff frequency from 1 Hz to 0.75 Hz for  
882 this experiment.

883 In computing  $direction_{steady}$ , the same computations were performed as for the primary-  
884 experiment, with one exception: a new direction was not necessarily decoded every millisecond.  
885 In order to decode a new direction, the follow conditions needed to be met: 1) the observed  
886 angular momentums had a Mahalanobis distance of less than 4 to the distribution  
887 corresponding to the decoded direction, 2) the observed angular momentums had a  
888 Mahalanobis distance of greater than 6 to the distribution corresponding to the opposite  
889 direction. These criteria ensured that a new steady-state direction was only decoded when the  
890 angular momentums were highly consistent with a particular direction. When these criteria  
891 were not met, the decoder continued to decode the same direction from the previous time step.

892 Speed was computed identically in the EARLY and STEADY states by decoding directly from  
893 the rotational plane corresponding to the decoded direction. A coarse estimate of speed was  
894 calculated as the derivative of the phase of rotation follows:

$$895 \quad \theta'(t) = \begin{cases} \frac{d\theta_f}{dt}, & direction(t) = +1 \\ \frac{d\theta_b}{dt}, & direction(t) = -1 \end{cases}$$

896 where  $\theta_f(t)$  and  $\theta_b(t)$  are the phases of the two planes in the neural state estimate  $\hat{x}_t$ ,  $direction$   
897 corresponds to  $direction_{early}$  while in the EARLY state and  $direction_{steady}$  while in the  
898 STEADY state, and the derivative  $\theta'$  is computed in units of Hz. The coarse speed estimate,  $\theta'$ ,  
899 was then smoothed with an exponential moving average ( $\tau = 500$  ms) to generate  $speed$ , the  
900 variable that gets used in the decoder state machine. Additional saturation limits were set such  
901 that  $speed$  never dropped below 0.5 Hz or exceeded 3.5 Hz, so as to remain in the range  
902 typically seen during pedaling. On entry into EARLY or STEADY from either INIT or EXIT,  
903 when  $speed$  gets initialized, the output of this exponential moving average was reset to an  
904 initial value of 1.5 Hz, which was the average starting speed across conditions.

905 Lastly, there were two new conditions for decoder state transitions in the speed-tracking  
906 experiment (**Supp. Fig. 1**). First, transitions from INIT to EARLY required that a condition  
907 termed “confident initial direction decode” was obtained. This condition was met when the  
908 Mahalanobis distance from the neural state in the initial-direction subspace to either the  
909 forward or backward distributions dropped below 4. Second, transitions into the EXIT state  
910 required (in addition to a drop in  $p_{move}$ ) that the observed angular momentums,  $L$ , belong to a

911 set termed ‘Stationary’. This set was defined as all  $L$  with a Mahalanobis distance of less than 4  
912 to the ‘stopped’ distribution of angular momentums, which was learned from the training set.

913

#### 914 **Acknowledgements**

915 We thank Y. Pavlova for expert animal care, A. Russo for sharing code and data for preliminary  
916 analysis, and E. Oby for surgical expertise and assistance. This work was supported by NINDS  
917 1DP2NS083037, NIH CRCNS R01NS100066, NINDS 1U19NS104649, the Simons Foundation  
918 (SCGB#325233 and SCGB#542957), the Grossman Center for the Statistics of Mind, the  
919 McKnight Foundation, P30 EY019007, a Klingenstein-Simons Fellowship, and the Searle  
920 Scholars Program.

921

#### 922 **Author Contributions**

923 M.M.C. conceived the study. K.E.S., S.M.P., and M.M.C. designed experiments. S.M.P., K.E.S.,  
924 M.M.C., and Q.W. created the decoding algorithms. K.E.S. and S.M.P. collected and analyzed  
925 datasets. K.E.S., S.M.P., and M.M.C. wrote the paper. All authors contributed to editing.

926

#### 927 **References**

- 928 1. Gilja, V. *et al.* A high-performance neural prosthesis enabled by control algorithm design.  
929 *Nat. Neurosci.* **15**, 1752–1757 (2012).
- 930 2. Gilja, V. *et al.* Clinical translation of a high-performance neural prosthesis. *Nat. Med.* **21**,  
931 1142–1145 (2015).
- 932 3. Shanechi, M. M. *et al.* Rapid control and feedback rates enhance neuroprosthetic control.  
933 *Nat. Commun.* **8**, 13825 (2017).
- 934 4. Shenoy, K. V. & Carmena, J. M. Combining Decoder Design and Neural Adaptation in  
935 Brain-Machine Interfaces. *Neuron* **84**, 665–680 (2014).
- 936 5. Ajiboye, A. B. *et al.* Restoration of reaching and grasping movements through brain-  
937 controlled muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration.  
938 *The Lancet* **0**, (2017).
- 939 6. Ethier, C., Oby, E. R., Bauman, M. J. & Miller, L. E. Restoration of grasp following paralysis  
940 through brain-controlled stimulation of muscles. *Nature* **485**, 368–371 (2012).

- 941 7. Collinger, J. L. *et al.* 7 degree-of-freedom neuroprosthetic control by an individual with  
942 tetraplegia. *Lancet* **381**, 557–564 (2013).
- 943 8. Wodlinger, B. *et al.* Ten-dimensional anthropomorphic arm control in a human  
944 brain–machine interface: difficulties, solutions, and limitations. *J. Neural Eng.* **12**, 016011  
945 (2015).
- 946 9. Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S. & Schwartz, A. B. Cortical control of a  
947 prosthetic arm for self-feeding. *Nature* **453**, 1098–1101 (2008).
- 948 10. Chapin, J. K., Moxon, K. A., Markowitz, R. S. & Nicolelis, M. A. L. Real-time control of a  
949 robot arm using simultaneously recorded neurons in the motor cortex. *Nat. Neurosci.* **2**,  
950 664–670 (1999).
- 951 11. Wessberg, J. *et al.* Real-time prediction of hand trajectory by ensembles of cortical neurons  
952 in primates. *Nature* **408**, 361–365 (2000).
- 953 12. Carmena, J. M. *et al.* Learning to Control a Brain–Machine Interface for Reaching and  
954 Grasping by Primates. *PLoS Biol* **1**, e42 (2003).
- 955 13. Taylor, D. M., Tillery, S. I. H. & Schwartz, A. B. Direct Cortical Control of 3D Neuroprosthetic  
956 Devices. *Science* **296**, 1829–1832 (2002).
- 957 14. Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R. & Donoghue, J. P. Brain-  
958 machine interface: Instant neural control of a movement signal. *Nature* **416**, 141–142  
959 (2002).
- 960 15. Scott, S. H., Gribble, P. L., Graham, K. M. & Cabel, D. W. Dissociation between hand  
961 motion and population vectors from neural activity in motor cortex. *Nature* **413**, 161–165  
962 (2001).
- 963 16. Russo, A. A. *et al.* Motor Cortex Embeds Muscle-like Commands in an Untangled  
964 Population Response. *Neuron* **97**, 953-966.e8 (2018).
- 965 17. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56  
966 (2012).

- 967 18. Michaels, J. A., Dann, B. & Scherberger, H. Neural Population Dynamics during Reaching  
968 Are Better Explained by a Dynamical System than Representational Tuning. *PLOS Comput.*  
969 *Biol.* **12**, e1005175 (2016).
- 970 19. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: A  
971 dynamical systems perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
- 972 20. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds  
973 a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* **18**, 1025–1033  
974 (2015).
- 975 21. Kao, J. C. *et al.* Single-trial dynamics of motor cortex and their applications to brain-machine  
976 interfaces. *Nat. Commun.* **6**, 1–12 (2015).
- 977 22. Aghagolzadeh, M. & Truccolo, W. Inference and Decoding of Motor Cortex Low-  
978 Dimensional Dynamics via Latent State-Space Models. *IEEE Trans. Neural Syst. Rehabil.*  
979 *Eng.* **24**, 272–282 (2016).
- 980 23. Makin, J. G., O’Doherty, J. E., Cardoso, M. M. B. & Sabes, P. N. Superior arm-movement  
981 decoding from cortex with a new, unsupervised-learning algorithm. *J. Neural Eng.* **15**,  
982 026010 (2018).
- 983 24. Rajangam, S. *et al.* Wireless Cortical Brain-Machine Interface for Whole-Body Navigation in  
984 Primates. *Sci. Rep.* **6**, (2016).
- 985 25. Libedinsky, C. *et al.* Independent Mobility Achieved through a Wireless Brain-Machine  
986 Interface. *PLOS ONE* **11**, e0165773 (2016).
- 987 26. Fitzsimmons, N., Lebedev, M., Peikon, I. & Nicolelis, M. A. L. Extracting kinematic  
988 parameters for monkey bipedal walking from cortical neuronal ensemble activity. *Front.*  
989 *Integr. Neurosci.* **3**, (2009).
- 990 27. Xing, D., Aghagolzadeh, M., Truccolo, W. & Borton, D. Low-Dimensional Motor Cortex  
991 Dynamics Preserve Kinematics Information During Unconstrained Locomotion in Nonhuman  
992 Primates. *Front. Neurosci.* **13**, (2019).

- 993 28. Capogrosso, M. *et al.* A brain–spine interface alleviating gait deficits after spinal cord injury  
994 in primates. *Nature* **539**, 284–288 (2016).
- 995 29. Fetz, E. E., Finocchio, D. V., Baker, M. A. & Soso, M. J. Sensory and motor responses of  
996 precentral cortex cells during comparable passive and active joint movements. *J.*  
997 *Neurophysiol.* **43**, 1070–1089 (1980).
- 998 30. Lemon, R. N., Hanby, J. A. & Porter, R. Relationship between the activity of precentral  
999 neurones during active and passive movements in conscious monkeys. *Proc. R. Soc. Lond.*  
1000 *B Biol. Sci.* **194**, 341–373 (1976).
- 1001 31. Suminski, A. J., Tkach, D. C. & Hatsopoulos, N. G. Exploiting multiple sensory modalities in  
1002 brain-machine interfaces. *Neural Netw.* **22**, 1224–1234 (2009).
- 1003 32. Schroeder, K. E. *et al.* Robust tactile sensory responses in finger area of primate motor  
1004 cortex relevant to prosthetic control. *J. Neural Eng.* **14**, 046016 (2017).
- 1005 33. Pruszynski, J. A. *et al.* Primary motor cortex underlies multi-joint integration for fast  
1006 feedback control. *Nature* **478**, 387–390 (2011).
- 1007 34. Christie, B. P. *et al.* Comparison of spike sorting and thresholding of voltage waveforms for  
1008 intracortical brain–machine interface performance. *J. Neural Eng.* **12**, 016009 (2014).
- 1009 35. Moran, D. W. & Schwartz, A. B. Motor Cortical Representation of Speed and Direction  
1010 During Reaching. *J. Neurophysiol.* **82**, 2676–2692 (1999).
- 1011 36. Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. Neuronal population coding of  
1012 movement direction. *Science* **233**, 1416–1419 (1986).
- 1013 37. Kaufman, M. T. *et al.* The Largest Response Component in the Motor Cortex Reflects  
1014 Movement Timing but Not Movement Type. *eNeuro* **3**, (2016).
- 1015 38. Kemere, C. *et al.* Detecting Neural-State Transitions Using Hidden Markov Models for Motor  
1016 Cortical Prostheses. *J. Neurophysiol.* **100**, 2441–2452 (2008).



- 1017 39. Kao, J. C., Nuyujukian, P., Ryu, S. I. & Shenoy, K. V. A High-Performance Neural  
1018 Prosthesis Incorporating Discrete State Selection With Hidden Markov Models. *IEEE Trans.*  
1019 *Biomed. Eng.* **64**, 935–945 (2017).
- 1020 40. Ethier, C., Sachs, N. A. & Miller, L. E. Continuous state-dependent decoders for brain  
1021 machine interfaces. in *2011 5th International IEEE/EMBS Conference on Neural*  
1022 *Engineering* 473–477 (2011). doi:10.1109/NER.2011.5910589.
- 1023 41. Kao, J. C., Ryu, S. I. & Shenoy, K. V. Leveraging neural dynamics to extend functional  
1024 lifetime of brain-machine interfaces. *Sci. Rep.* **7**, (2017).
- 1025 42. Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A. & Miller, L. E. Long-term  
1026 stability of cortical population dynamics underlying consistent behavior. *Nat. Neurosci.* **23**,  
1027 260–270 (2020).
- 1028 43. Foster, J. D. *et al.* A freely-moving monkey treadmill model. *J. Neural Eng.* **11**, 046020  
1029 (2014).
- 1030 44. Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A. & Shenoy, K. V. A high-performance brain-  
1031 computer interface. *Nature* **442**, 195–198 (2006).
- 1032 45. Sadtler, P. T., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M. & Batista, A. P. Brain–computer  
1033 interface control along instructed paths. *J. Neural Eng.* **12**, 016015 (2015).
- 1034 46. Nuyujukian, P., Kao, J. C., Ryu, S. I. & Shenoy, K. V. A Nonhuman Primate Brain–Computer  
1035 Typing Interface. *Proc. IEEE* **105**, 66–72 (2017).
- 1036 47. Fernández-Rodríguez, Á., Velasco-Álvarez, F. & Ron-Angevin, R. Review of real brain-  
1037 controlled wheelchairs. *J. Neural Eng.* **13**, 061001 (2016).
- 1038 48. Thacker, N. A. & Bromiley, P. A. *The Effects of a Square Root Transform on a Poisson*  
1039 *Distributed Quantity*. Technical Report [http://www.tina-vision.net/tina-knoppix/tina-](http://www.tina-vision.net/tina-knoppix/tina-memo/2001-010.pdf)  
1040 [memo/2001-010.pdf](http://www.tina-vision.net/tina-knoppix/tina-memo/2001-010.pdf) (2001).
- 1041 49. Cunningham, J. P. & Ghahramani, Z. Linear Dimensionality Reduction: Survey, Insights,  
1042 and Generalizations. *J Mach Learn Res* **16**, 2859–2900 (2015).

- 1043 50. Malik, W. Q., Truccolo, W., Brown, E. N. & Hochberg, L. R. Efficient Decoding With Steady-  
1044 State Kalman Filter in Neural Interface Systems. *IEEE Trans. Neural Syst. Rehabil. Eng.* **19**,  
1045 25–34 (2011).