

Cortical control of virtual self-motion using task-specific subspaces

Karen E Schroeder*^{1,2}, Sean M Perkins*^{2,3}, Qi Wang³, & Mark M Churchland^{1,2,4,5, †}

* Equal contribution

¹ Department of Neuroscience, Columbia University Medical Center, New York, NY

² Zuckerman Institute, Columbia University, New York, NY

³ Department of Biomedical Engineering, Columbia University, New York, NY

⁴ Kavli Institute for Brain Science, Columbia University Medical Center, New York, NY

⁵ Grossman Center for the Statistics of Mind, Columbia University, New York, NY

† Correspondence to: 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 movement
4 classes. The question of how to decode neural activity is inexorably linked with the intrinsic
5 covariance structure of that activity, which may depend strongly upon movement class. Here,
6 we develop a self-motion BMI based on cortical activity as monkeys cycle a hand-held pedal
7 to progress along a virtual track. Unlike during reaching, there were no high-variance
8 dimensions that directly correlated with to-be-decoded variables. Yet this challenge yielded an
9 opportunity: we could decode a single variable – self-motion – by non-linearly leveraging
10 structure that spanned many high-variance neural dimensions. Online BMI-control success
11 rates approached those during manual control. Our results argue that decoding can and
12 should be task-specific, and suggest a broad principle: even when the decoded output is low-
13 dimensional, it can be beneficial to leverage a multi-dimensional high-variance subspace.

14 **Introduction**

15 Intracortical brain-machine interfaces (BMIs) that support reach-like tasks have proved
16 successful in primates and human clinical trials (Ajiboye et al., 2017; Collinger et al., 2013;
17 Ethier et al., 2012; Gilja et al., 2015, 2012; Shanechi et al., 2017; Shenoy and Carmena, 2014;
18 Wodlinger et al., 2015). Yet it is unclear whether current decode algorithms will generalize well
19 to non-reaching applications. Early reach-based BMIs (Carmena et al., 2003; Chapin et al.,
20 1999; Serruya et al., 2002; Taylor et al., 2002; Velliste et al., 2008; Wessberg et al., 2000)
21 employed a strategy of inverting the ostensible neural encoding of kinematic variables,
22 primarily hand velocity and direction. Despite evidence against literal kinematic encoding
23 (Churchland et al., 2012; Michaels et al., 2016; Russo et al., 2018; Scott et al., 2001; Sergio et
24 al., 2005; Sussillo et al., 2015), studies continue to exploit the robust correlation between
25 neural activity and movement direction to achieve excellent BMI performance. Improvements
26 have derived from honing this core strategy and/or better estimating the neural state
27 (Aghagolzadeh and Truccolo, 2016; Gilja et al., 2012; Kao et al., 2015; Makin et al., 2018).
28 Viewed this way, current decode strategies successfully leverage a statistical regularity
29 (aspects of neural activity correlate with velocity) that is sizeable during reach-like movements.
30 Should non-reach-like movements leverage the same or different statistical regularities?

31 This question is intertwined with our evolving understanding of motor cortex covariance
32 structure. During a given task, neural covariance is typically low-dimensional, in the sense that
33 a handful of neural dimensions captures considerable variance (Churchland et al., 2007;
34 Gallego et al., 2017; Sadtler et al., 2014). Nevertheless, dimensions with non-negligible
35 variance greatly outnumber to-be-decoded variables (Churchland and Shenoy, 2007; Shenoy
36 et al., 2013). The traditional solution is to identify a subspace where population activity is both
37 reasonably high-variance and correlates strongly with variables one wishes to infer. For
38 example, regressing two-dimensional cursor velocity against neural activity identifies a two-
39 dimensional neural subspace where activity resembles velocity. A few studies have
40 successfully leveraged additional dimensions (Degenhart et al., 2020; Gallego et al., 2020; Kao
41 et al., 2017b; Sussillo et al., 2016, 2012), but the typical scenario is that decoding N variables
42 leverages an N -dimensional subspace chosen to maximize correlations.

43 Within a given task, neural covariance structure remains surprisingly fixed even when the
44 decoder is altered (Golub et al., 2018; Sadtler et al., 2014) and can remain similar across

45 related tasks (Gallego et al., 2018). Yet covariance is often not stable. Covariance changes
46 dramatically when cycling forward versus backward (Russo et al., 2018), when using one arm
47 versus another (Ames and Churchland, 2019), when preparing versus moving (Elsayed et al.,
48 2016; Kaufman et al., 2014), and when reaching versus walking (Miri et al., 2017). Thus, there
49 may be no high-variance task-independent subspace that consistently exhibits correlations
50 with to-be-decoded variables. This would agree with the view that the highest-variance signals
51 in motor cortex ensure noise-robust task-specific dynamics and do not reliably encode any
52 external quantity (Russo et al., 2018). This possibility creates challenges, but also
53 opportunities. Task-specific subspaces may be of high utility even when there is no linear
54 correlation between activity and a decoded variable. For example, one could decode forward
55 self-motion from the angular momentum of neural activity in a two-dimensional subspace. To
56 guide backward self-motion, a completely different neural subspace could be employed. Thus,
57 a one-dimensional output could potentially leverage activity spanning many dimensions.

58 To explore such strategies, we employed a simple task in which monkeys cycle a hand-held
59 pedal to move along a virtual track. The class of neural activity evoked by this task could
60 potentially be leveraged by future BMI devices that guide self-motion. Even more critically,
61 cycling is overtly different from reach- or joystick-based tasks. Prior explorations of self-
62 motion-decoding borrowed from the strategies employed by reach-based BMIs, and decoded
63 a whole-body directional vector (Rajangam et al., 2016) or classified the direction of a joystick
64 intermediary (Libedinsky et al., 2016). This is a promising approach, but we wished to explore
65 whether a task-specific approach might also be promising.

66 During cycling, neural activity that correlated with kinematics and/or muscle activity was
67 relegated to low-variance subspaces. In contrast, there existed high-variance subspaces
68 where neural activity had reliable non-linear relationships with intended self-motion. A decoder
69 leveraging those multi-dimensional subspaces provided excellent online control of a one-
70 dimensional command for virtual self-motion. Almost no training was needed; decoder latency
71 and accuracy were such that monkeys appeared to barely notice transitions from manual
72 control to BMI control. Success rates and acquisition times were close to those achieved
73 under manual control. Our results establish that BMIs can be accurately guided by strategies
74 that leverage task-specific subspaces containing structure that, while not directly correlating
75 with kinematic variables, bears a reliable relationship with desired movement.

76

77 **Results**

78 ***Behavior***

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

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

103 BMI control was introduced after monkeys were adept at performing the task under manual
104 control. Task structure and the parameters for success were unchanged under BMI control,
105 and no cue was given regarding the change from manual to BMI control. The switch to BMI
106 control was made after completion of a block of manual-control trials (25 forward and 25
107 backward 7-cycle trials). These manual-control trials were used to train the decoder, at which

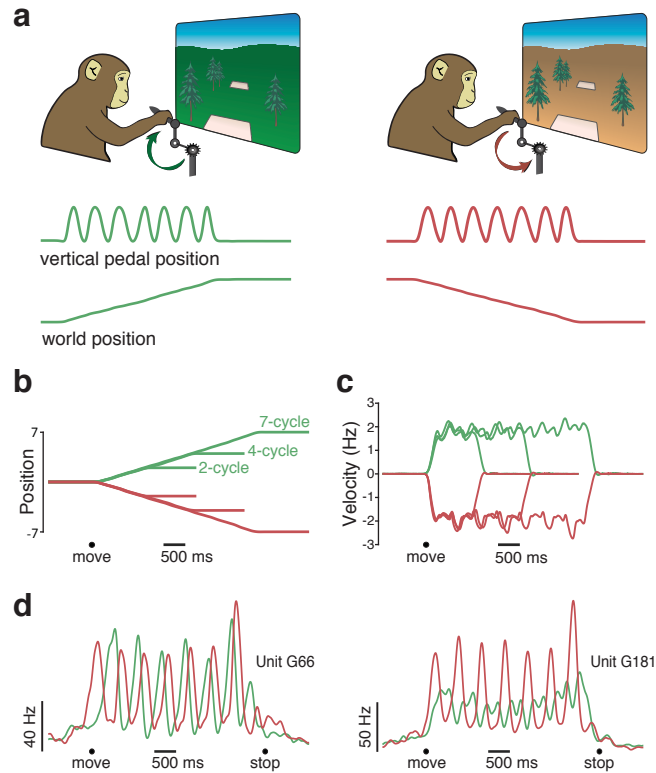


Figure 1. A cycling task that elicits rhythmic movements. (a) Monkeys rotated a hand-held pedal forward (*left*, cued by a green background) or backward (*right*, cued by a tan background) to progress through a virtual environment. Traces at bottom plot pedal kinematics (vertical position) and the resulting 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.

108 point the switch was made to BMI control for the remainder of the session. For monkey G, we
109 occasionally included blocks of manual-control trials later in the session to allow comparison
110 between BMI and manual performance. For Monkey E we used separate (interleaved) sessions
111 to assess manual-control performance.

112 During both BMI control and manual control, the monkey's ipsilateral (non-cycling) arm was
113 restrained. The contralateral (cycling) arm was never restrained. We intentionally did not
114 dissuade the monkey from continuing to physically cycle during BMI control. Indeed, our goal
115 was that the transition to BMI control would be sufficiently seamless to be unnoticed by the
116 monkey, such that he would still believe that he was in manual control. An advantage of this
117 strategy is that we are decoding neural activity when the subject attempts to actually move, as
118 a patient presumably would. Had we insisted the arm remain stationary, monkeys would have
119 needed to actively avoid patterns of neural activity that drive movement – something a patient
120 would not have to do. Allowing the monkey to continue to move normally also allowed us to
121 quantify decoder performance via direct comparisons with intended (i.e., actual) movement.
122 This is often not possible when using other designs. For example, in Rajangam et. al.
123 (Rajangam et al., 2016), performance could only be assessed via indirect measures (e.g., time
124 to target) because what the monkey was actually intending to do at each moment was unclear.
125 We considered these advantages to outweigh a potential concern: a decoder could potentially
126 'cheat' by primarily leveraging activity driven by proprioceptive feedback (which would not be
127 present in a paralyzed patient). This is unlikely to be a large concern. Recordings were made
128 from motor cortex, where robust neural responses precede movement onset. Furthermore, we
129 have documented that motor cortex population activity during cycling is quite different from
130 that within the proprioceptive region of primary somatosensory cortex (Russo et al., 2018).
131 Thus, while proprioceptive activity is certainly present in motor cortex (Fetz et al., 1980; Lemon
132 et al., 1976; Schroeder et al., 2017; Suminski et al., 2009) especially during perturbations
133 (Pruszynski et al., 2011), the dominant features of M1 activity that we leverage are unlikely to
134 be primarily proprioceptive.

135 Given our use of healthy animals, we stress that the goal of the present study is to determine
136 strategies for leveraging the dominant structure of neural activity. This follows the successful
137 strategy of BMI studies that leveraged the well-characterized structure of activity during
138 reaching. Of course, the nature of the training data used to specify decode parameters (e.g.,
139 the weights that define the neural dimensions to be used) will necessarily be different for a

140 healthy animal that cannot understand verbal instructions and an impaired human that can. We
141 thus stress that our goal is to determine a robust and successful decode strategy that works in
142 real time during closed-loop performance. We do not attempt to determine the best approach
143 to parameter specification, which in a patient would necessarily involve intended or imagined
144 movement.

145 ***Neural activity and decoding strategy***

146 We recorded motor cortical activity using 96-channel Utah arrays. For monkey G, one array
147 was implanted in primary motor cortex (M1) and a second in dorsal premotor cortex (PMd). For
148 monkey E, a single array was implanted in M1. For each channel we recorded times when the
149 voltage crossed a threshold. Threshold crossings typically reflected individual spikes from a
150 small handful of neurons (a neural ‘unit’). Spikes from individual neurons could be clearly seen
151 on many channels, but no attempt was made to spike-sort. The benefit of sorting is typically
152 modest when controlling a prosthetic device (Christie et al., 2014), and reduced-dimension
153 projections of motor cortex population activity are similar whether based on single or multi-
154 units (Trautmann et al., 2019). Unit activity was strongly modulated during cycling (**Fig. 1d**). The
155 phase, magnitude, and temporal pattern of activity depended on whether cycling was stopped,
156 moving forward (*green* traces) or moving backward (*red* traces). A key question is how these
157 unit-level features translate into population-level features that might be leveraged to estimate
158 intended motion through the virtual environment.

159 In traditional decoding approaches (**Fig. 2a**) neural activity is hypothesized (usefully if not
160 literally) to encode kinematic signals, which can be decoded by inverting the encoding scheme.
161 Although nonlinear methods (including variations of Kalman filtering) are often used to estimate
162 the neural state, the final conversion to a kinematic command is typically linear or roughly so.
163 To explore kinematic encoding in the present task, we used linear regression to identify neural
164 dimensions where activity correlated well with kinematics (including hand velocity and
165 position). Regression was performed using single trials. Use of single trials provides a large
166 quantity of training data and is implicitly regularizing: regression must find signals that are
167 robust in the face of single-trial spiking variability. The regression weights for a given kinematic
168 parameter define a neural dimension where activity correlates strongly with that parameter. We
169 computed the neural variance captured by each such dimension. Variance captured was
170 computed using trial-averaged data, to ensure that values were not diluted by noise. Despite

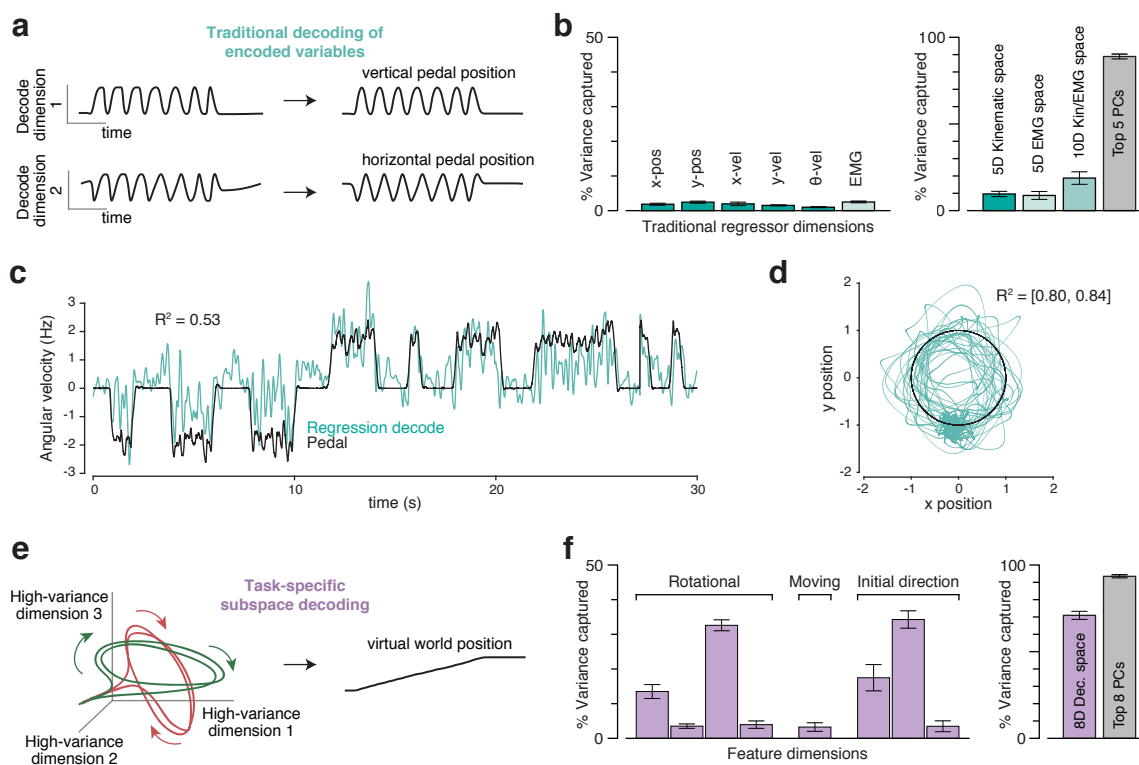


Figure 2. Different decode strategies leverage neural signals with different magnitudes.

(a) In the traditional decoding strategy, 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). **(b)** Variance of the neural population response captured by dimensions used to decode kinematic parameters (*teal bars*) and muscle activity (*light teal 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)** Example cross-validated regression performance for offline decoding of angular velocity. R^2 is the coefficient of determination for the segment of data shown. **(d)** Example cross-validated regression performance for offline decoding of horizontal and vertical pedal position. R^2 is the coefficient of determination for the segment of data shown, same time period as (c). **(e)** A new strategy that 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. **(f)** Similar plot to (b) 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.

171 this, the neural dimensions that best captured kinematic signals captured little population
172 response variance (**Fig. 2b**, *teal* bars). This was also true of neural dimensions that captured
173 muscle activity (**Fig. 2b**, *light teal* bar). This was initially surprising: single-neuron responses
174 were robustly sinusoidally modulated, as were many kinematic variables. Yet sinusoidal
175 response features were often superimposed upon other response features (e.g., overall shifts in
176 rate when moving versus not moving). Sinusoidal features also did not display phase
177 relationships, across forward and backward cycling, consistent with kinematic encoding
178 (Russo et al., 2018). As a result, the dimensions where activity correlated strongly with
179 kinematics captured relatively little response variance.

180 The fact that kinematic-encoding neural dimensions are low-variance makes them a potentially
181 challenging basis for decoding. For example, we identified a dimension where the projection of
182 trial-averaged neural activity correlated with angular velocity, which is conveniently proportional
183 to the quantity we wish to decode (velocity of virtual self-motion). However, because that
184 dimension captured relatively little variance ($1.0\% \pm 0.2\%$ of the overall population variance;
185 standard deviation across three sessions) the relevant signal was variable on single trials (**Fig.**
186 **2c**) and the correlation with angular velocity was poor ($R^2 = 0.52 \pm 0.14$). Any decode based on
187 this signal would result in many false starts and false stops. The observed single-trial variability
188 is expected for any low-variance signal estimated from a limited number of neurons.

189 We considered an alternative kinematic decoding strategy that more closely mimics prior
190 successful reach-based approaches: decoding horizontal and vertical hand position. These
191 were captured in slightly higher-variance dimensions, and could be decoded somewhat better
192 on single trials ($R^2 = [0.81 \pm 0.03, 0.85 \pm .04]$, averaged across three sessions). A reasonable
193 strategy would be to convert this two-dimensional neural signal (**Fig. 2d**) into a one-dimensional
194 self-motion command. Possible approaches for doing so include deciphering where on the circle
195 the hand is most likely to be at each moment, computing the angular velocity of the neural state,
196 or computing its angular momentum. However, consideration of such approaches raises a
197 deeper question. If we are willing to non-linearly decode a one-dimensional quantity (self-
198 motion) from a two-dimensional subspace, why stop at two dimensions? Furthermore, why not
199 leverage dimensions that capture most of the variance in the population response, rather than
200 dimensions that capture a small minority? Projections onto high-variance dimensions have
201 proportionally less contribution from spiking variability. They are also less likely to be impacted
202 by electrical interference or recording instabilities. While of little concern in a controlled

203 laboratory environment, this is relevant to the clinical goals of prosthetic devices. One should
204 thus wish to leverage multiple dimensions that jointly capture as much variance as possible, so
205 long as the signals captured by those dimensions can be successfully leveraged.

206 We term this approach (**Fig. 2e**) task-specific subspace decoding, because it seeks one or more
207 subspaces that capture robust response features as the subject attempts to move within the
208 context of a particular task (in the present case, while generating a rhythmic movement to
209 produce self-motion). Both the population-level response features, and the subspaces
210 themselves, may be specific to that task. To pursue this strategy, we identified three high-
211 variance subspaces. The first was spanned by four ‘rotational dimensions’ (two each for
212 forward and backward cycling) which captured elliptical trajectories present during steady-
213 state cycling (Russo et al., 2018). The second was a single ‘moving-sensitive’ dimension, in
214 which the neural state distinguished whether the monkey was stopped or moving regardless of
215 movement direction (Kaufman et al., 2016). The third was a triplet of ‘initial-direction’
216 dimensions, in which cycling direction could be transiently distinguished in the moments after
217 cycling began.

218 In subsequent sections we document the specific features present in these high-variance
219 subspaces. Here we concentrate on the finding that the eight-dimensional space spanned by
220 these eight dimensions captured $70.9\% \pm 2.3\%$ of the firing-rate variance (**Fig. 2f**). Note that
221 because dimensions are not orthogonal, the total captured variance is not the sum of that for
222 each dimension. The ‘initial direction’ dimensions, for example, overlap considerably with the
223 ‘rotational dimensions.’ Nevertheless, the total captured variance was only modestly less than
224 that captured by the top eight PCs (which capture the most variance possible), and much
225 greater than that captured by spaces spanned by dimensions where activity correlated with
226 kinematics and/or muscle activity (**Fig. 2b**). We thus based our BMI decode on activity in these
227 eight high-variance dimensions.

228 ***Direction of steady-state movement inferred from rotational structure***

229 The dominant feature of the population neural response during steady-state cycling was a
230 repeating elliptical trajectory (Russo et al., 2018). Thus, the core of our decoder was built upon
231 this feature, leveraging the fact that forward-cycling and backward-cycling trajectories
232 occurred in non-identical subspaces (**Fig. 3a,b**). We employed an optimization procedure to
233 find a two-dimensional ‘forward plane’ that maximized the size of the forward trajectory relative

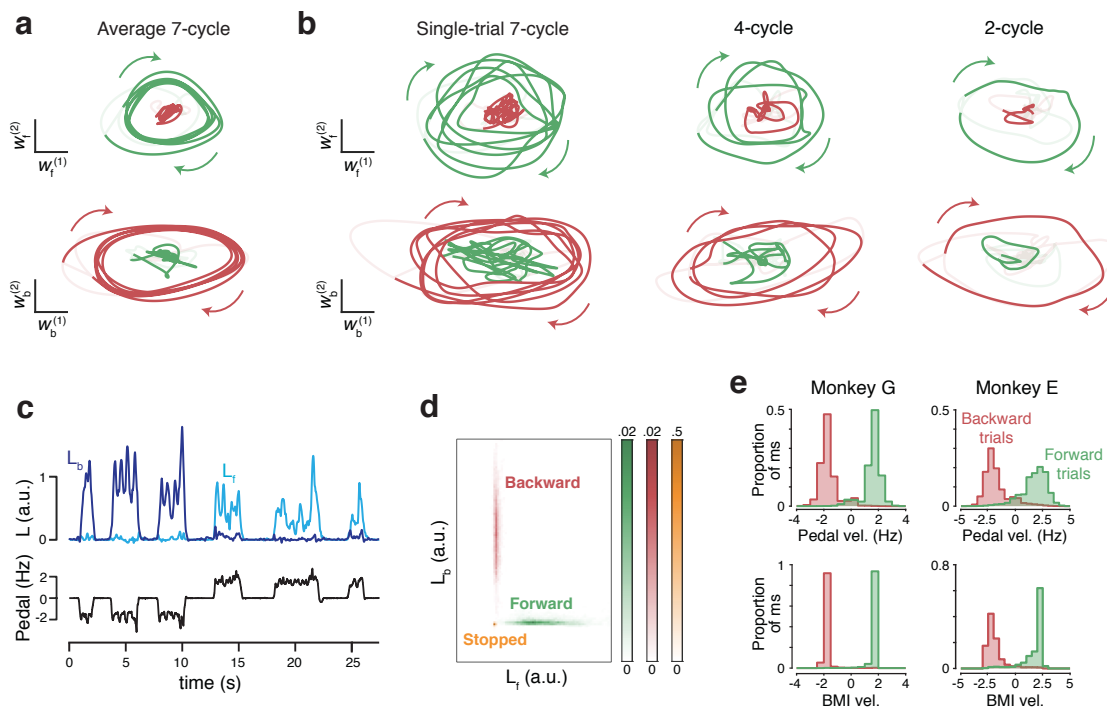


Figure 3. Leveraging rotational trajectories to decode velocity. (a) Trial-averaged population activity, during a manual-control block, projected onto the forward (*top*) and backward (*bottom*) rotational planes. Data are from seven-cycle forward (*green*) and backward (*red*) conditions. By design, 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 momentums 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 (see Methods), across all BMI-control sessions.

234 to the backward trajectory. We similarly found an analogous ‘backward plane’. These planes
235 were identified based on trial-averaged responses from the 50 trials of training data collected
236 under manual control (**Fig. 3a**). With the aid of filtering (*Methods*), these planes continued to
237 capture rotational features on individual trials (**Fig. 3b**). Although forward and backward
238 trajectories were not orthogonal to one another, the above procedure was still able to find
239 planes where strongly elliptical trajectories were present for only one cycling direction. To
240 convert the four-dimensional neural state to a one-dimensional decode of self-motion, we
241 compared angular momentum (the cross product of the state vector with its derivative)
242 between the two planes. When moving backward (first three cycling bouts in **Fig. 3c**) angular
243 momentum was sizeable in the backward plane (*dark blue*) but not the forward plane (*bright*
244 *blue*). The opposite was true when moving forward (subsequent three bouts).

245 Based on training data, we considered the joint distribution of forward-plane and backward-
246 plane angular momentum. We computed distributions when stopped (**Fig. 3d, orange**), when
247 cycling forward (*green*) and when cycling backward (*red*). These distributions overlapped little,
248 and we fit a Gaussian to each. During BMI control, we computed the likelihood of the observed
249 angular momentums under each of the three distributions. If likelihood under the stopped
250 distribution was high, decoded velocity was zero. Otherwise, decoded velocity was determined
251 by the relative likelihoods under the forward and backward distributions. These likelihoods
252 were converted into a virtual velocity that was maximal when one likelihood was much higher
253 and slower when likelihoods were more similar. The maximum decoded virtual velocity was set
254 to approximate the typical virtual velocity under manual control, when cycling at ~2 Hz.

255 Distributions of decoded velocity when moving under BMI control (**Fig. 3e, bottom**) were
256 similar to the distributions of velocity that would have resulted were the pedal still operative
257 (**Fig. 3e, top**). Importantly, distributions overlapped very little; the direction of decoded motion
258 was almost always correct. Decoded velocity was near maximal at most times, especially for
259 monkey G. The decoded velocity obtained from these four dimensions (out of the eight total we
260 used) constituted the core of our decoder. We document the performance of that decoder in
261 the next section. Later sections describe how aspects of the decoder were fine-tuned by
262 leveraging the remaining four dimensions.

263 **Performance**

264 Monkeys performed the task very well under closed-loop BMI control (**Fig. 4** and **Supp. Movie**
265 **1**). Monkeys continued to cycle as normal, presumably not realizing that the pedal had been
266 disconnected from the control system. The illusion that the pedal still controlled the task was
267 supported by a high similarity between decoded virtual velocity and intended virtual velocity
268 (i.e., what would have been produced by the pedal were it still controlling the task). The cross-
269 correlation between these peaked at $0.93 \pm .02$ and $0.81 \pm .03$ (monkey G and E, mean \pm SD
270 across sessions) at a short lag: 76 ± 4 ms and 102 ± 7 ms (**Fig. 4a**). That illusion was also
271 aided by a low rate of false starts; it was exceedingly rare for decoded motion to be non-zero
272 when the monkey was attempting to remain stationary on top of a target. False starts occurred
273 on 0.29% and 0.09% of trials (monkeys G and E), yielding an average of 1.9 and 0.12
274 occurrences per day. This is notable because combatting unintended movement is a key
275 challenge for BMI decoding. The above features – high correlation with intended movement,
276 low latency, and few false starts – led to near-normal performance under BMI control (**Fig.**
277 **4b,c**). Success rates under BMI control (**Fig. 4d**, *magenta symbols*) were almost as high as
278 under manual control (*open symbols*), and the time to move from target to target was only
279 slightly greater under BMI control (**Fig. 4e**).

280 The only respect in which BMI control suffered noticeably was accuracy in stopping on the
281 middle of the target. Under manual control, monkeys stopped very close to the target center
282 (**Fig. 4f**, *gray histogram*), which always corresponded to the ‘pedal-straight-down’ position.
283 Stopping was less accurate under BMI control (*magenta histogram*). This was partly due to the
284 fact that because virtual self-motion was swift, small errors in decoded stopping time became
285 relevant. E.g., a 100 ms error corresponded to ~ 0.2 cycles of physical motion. The average
286 standard deviation of decoded stopping time (relative to actual stopping time) was 133
287 (monkey G) and 99 ms (monkey E). The relatively larger error in BMI-control trials was also due
288 to an incidental advantage of manual control: the target center was aligned with the pedal-
289 straight-down position, a fact which monkeys leveraged to stop very accurately in that
290 position. This strategy was not available during BMI control because the correct moment to
291 stop rarely aligned perfectly with the pedal-straight-down position (this occurred only if
292 decoded and intended virtual velocity matched perfectly when averaged across the cycling
293 bout).

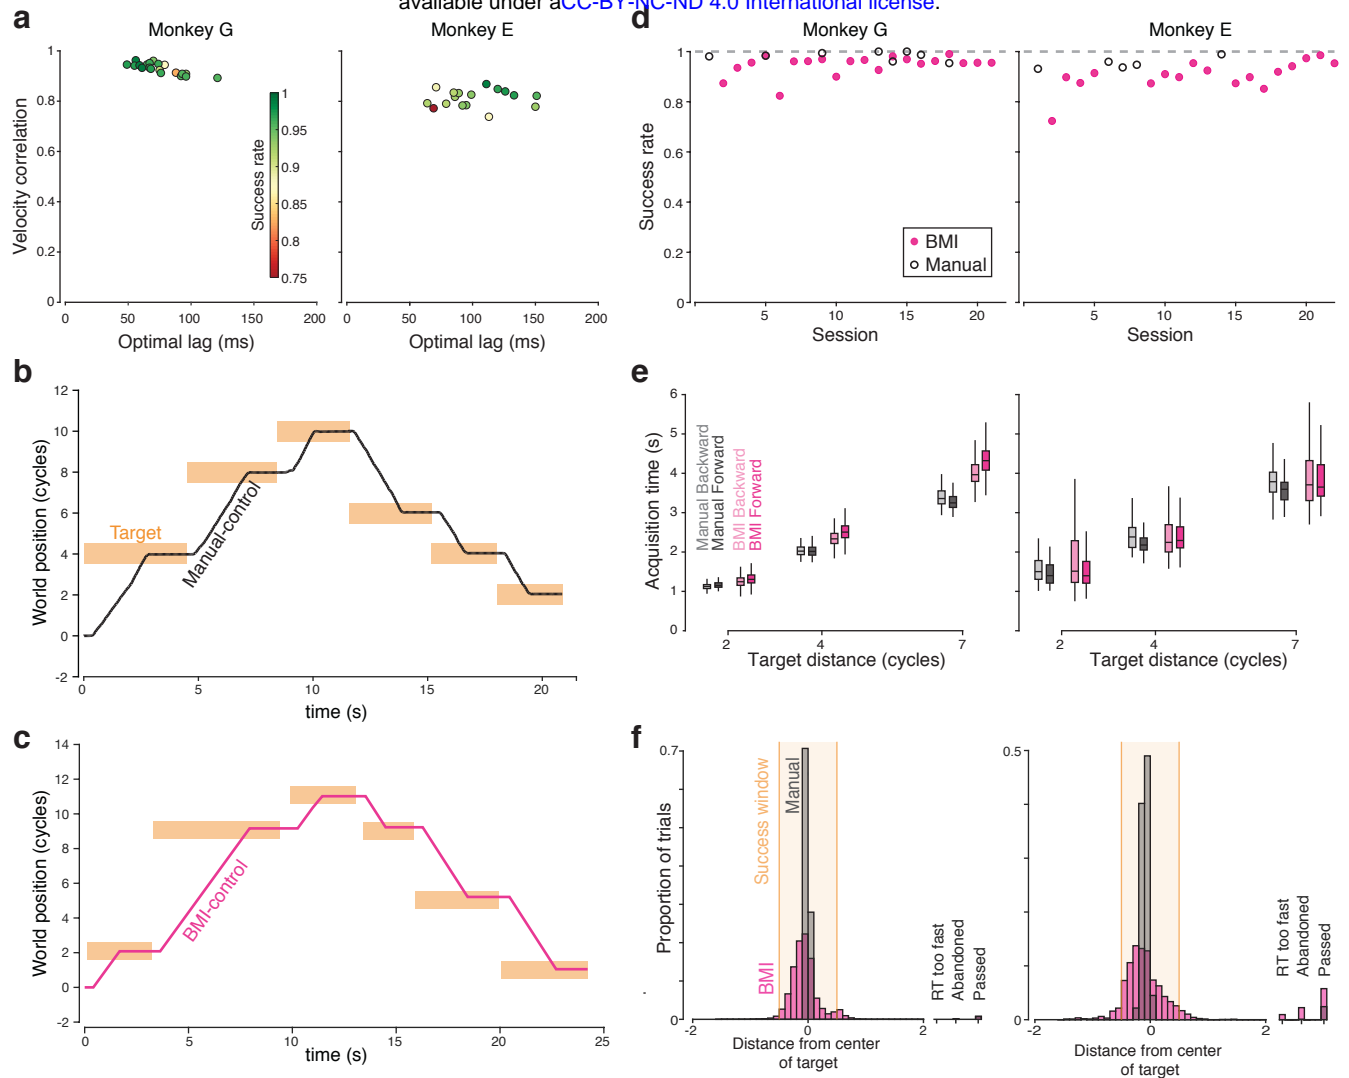


Figure 4. Decoder performance. **(a)** Summary of the cross-correlation between decoded virtual velocity under BMI control, and the virtual velocity that would have been produced by the pedal (which monkeys continued to manipulate normally). Each symbol corresponds to one BMI-control session, and plots 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.

294 Performance was modestly better for monkey G versus E. This was likely due to the
295 implantation of two arrays rather than one. Work ethic may also have been a factor; monkey E
296 performed fewer trials under both BMI and manual control. Still, both monkeys could use the
297 BMI successfully starting on the first day, with success rates of 0.87 and 0.74 (monkey G and
298 E). Monkey G's performance rapidly approached his manual-control success rate within a few
299 sessions. Monkey E's performance also improved quickly, although his manual-control and
300 BMI-control success rates were mostly lower than Monkey G's. The last five sessions involved
301 BMI success rates of 0.97 and 0.96 for the two monkeys. This compares favorably with the
302 overall averages of 0.98 and 0.95 under manual control. Although this performance
303 improvement with time may relate to adaptation, the more likely explanation is simply that
304 monkeys learned to not be annoyed or discouraged by the small differences in decoded and
305 intended velocity.

306 ***State machine***

307 The performance documented above was achieved using a state-dependent decode (**Fig. 5**).
308 The angular-momentum-based strategy determined virtual self-motion (as described above)
309 when in the "STEADY" state. Thus, self-motion could be forward, backward, or stopped while
310 in STEADY. The use of other states was not strictly necessary but helped fine-tune
311 performance. State transitions were governed by activity in the moving-sensitive dimension,
312 which was translated into a probability of moving, p_{move} as described in the next section. If
313 p_{move} was low, the STOP state was active and decoded virtual velocity was enforced to be
314 zero. For those time-points when the STOP state was active, the angular-momentum-based
315 strategy would have typically estimated zero velocity even if it were not enforced.
316 Nevertheless, the use of an explicit STOP state was useful in nearly eliminating false starts. It
317 was also useful in detecting transitions from stopping to moving, allowing us to leverage the
318 initial-direction dimensions.

319 When p_{move} became high, the INIT state was entered but decoded velocity remained zero.
320 After 175 ms, the EARLY state was entered and velocity was decoded using the initial-direction
321 dimensions (see below). After an additional 200 ms, the STEADY state was entered. The
322 decode then depended only on the four rotational dimensions. Values of $p_{move} < 0.1$ always
323 produced a transition back to STOP. This typically occurred from STEADY to STOP, as the
324 movement was successfully ending. However, it could also occur from the other two states.




Figure 5. 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.

325 This was especially helpful if p_{move} became high very briefly (and presumably erroneously). In
326 such cases the state could transition from INIT back to STOP with the decoded velocity never
327 departing from zero.

328 ***Inferring the probability of moving***

329 Decoders that directly translate neural state to cursor velocity have historically had difficulty
330 remaining stationary when there is no intended movement. The ability to successfully decode
331 stationarity is of even greater importance for self-motion. This was central to our motivation for
332 using a state machine with distinct stopped and moving states (Ethier et al., 2011; Kao et al.,
333 2017a; Kemere et al., 2008). To govern state transitions, we identified the moving-sensitive
334 dimension by applying linear discriminant analysis to the 50 training-data trials, and finding the
335 direction that best discriminated whether the monkey was moving versus stopped. Projecting
336 trial-averaged data onto that dimension (**Fig. 6a**) revealed that activity transitioned suddenly
337 from low to high just before movement onset, and back to low around the time movement
338 ended. This pattern was remarkably similar regardless of cycling direction (*red* and *green*
339 traces largely overlap). Activity in this dimension behaved similarly for single trials (**Fig. 6b**).

340 We used a Hidden Markov Model (HMM) (Kao et al., 2017a; Kemere et al., 2008) to estimate
341 p_{move} , which allows the current estimate to depend on all prior observations. Because those
342 observations must be independent, we did not use filtered rates (which were used for all other
343 aspects of the decode) but instead considered spike counts in non-overlapping bins, projected
344 onto the moving-sensitive dimension. Figure 6c plots the resulting distributions when stopped
345 (*orange*) and moving (*blue*). These overlapped modestly, a result of the narrow (10 ms) bin. The
346 estimate of p_{move} is robust to this overlap because the HMM leverages the full history of spike
347 counts; it can ignore brief weak evidence for moving while still transitioning swiftly given strong
348 evidence. During BMI control, p_{move} (**Fig. 6d, blue**) was typically near unity during intended
349 movement (i.e., when the monkey was actually cycling, *black*) and near zero otherwise.

350 Transitions both out of STOP and into STOP were determined by p_{move} (**Fig. 5**). We employed
351 a conservative design; entering a state that produced virtual movement required that p_{move}
352 exceed 0.9 and remain consistently above 0.1 for 175 ms. This led to a very low rate of false
353 starts (~2 per day for monkey G and ~1 every ten days for monkey E). The transition to EARLY
354 (**Fig. 6d, left edge of gray regions**) occurred on average 117 and 194 ms after physical
355 movement onset (monkeys G and E). Trial-to-trial variability around these mean values was

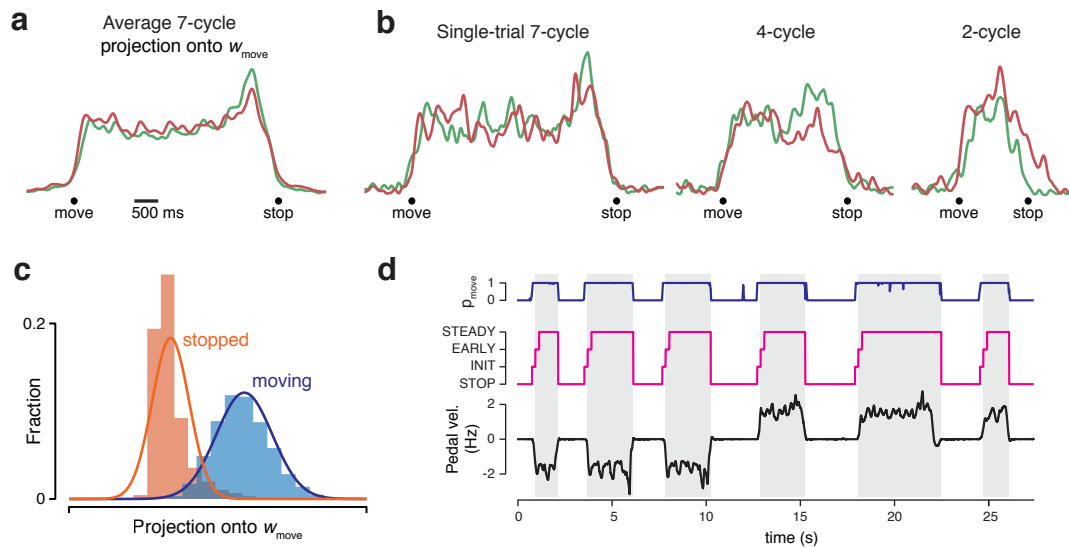


Figure 6. Leveraging the moving-sensitive dimension to infer probability of moving. (a) Trial-averaged population activity, during a manual-control block, projected onto the moving-sensitive dimension (same session and trials as Figure 3a). **(b)** As in panel (a), but for three example single trials (same trials as in Figure 3b). **(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 3c.

356 modest: standard deviations were 93 and 138 ms (computed within session and averaged
357 across sessions). As discussed above, estimated stopping time (when p_{move} dropped below
358 0.1) was also decoded with only modest trial-to-trial variability.

359 ***Inferring initial movement direction***

360 Angular momentum of the neural state in the forward and backward planes became substantial
361 a few hundred milliseconds after p_{move} became high. Thus, the EARLY state became active
362 before the direction of movement could be inferred from the elliptical trajectories. To overcome
363 this problem, we leveraged the three-dimensional initial-direction subspace. The initial-
364 direction subspace comprised the top three PCs found from portions of the training data
365 surrounding movement onsets (a 200 ms segment from each trial, beginning at the time the
366 decoder would enter the INIT state, see *Methods*). The difference between the trajectory on
367 forward and backward trials began to grow just prior to physical movement onset, both on
368 average (**Fig. 7a**) and on individual trials (**Fig. 7b**; *solid* trajectory segment shows -200 to +175
369 ms relative to movement onset). For each of the 50 training trials, we considered the neural
370 state in these dimensions, measured 175 ms after decoded movement onset (**Fig. 7c**). We fit
371 Gaussian distributions separately for forward and backward trials.

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

382 ***Speed control***

383 The excellent performance of the decoder was aided by the relative simplicity of behavior:
384 when monkeys moved, they did so at a stereotyped speed. This allowed us to concentrate on
385 building a decode algorithm that decoded intended direction with accurate timing, and




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 3a 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 3b 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 3c and 6d.

386 remained stationary if movement was not intended. However, that decode provided only
387 limited control of movement speed. An obvious extension is to allow finer-grained speed
388 control. This would presumably be desired by users of a self-motion prosthetic. Furthermore,
389 speed control provides one possible way of steering: e.g., by decoding the relative intensity of
390 intended movement on the two sides of the body. While we do not attempt that here, we still
391 considered it important to determine whether the neural features we identified could support
392 speed control.

393 That assessment required a task where speed control is necessary for success. We thus
394 trained one monkey to track various speed profiles as he progressed through the virtual
395 environment. Two floating targets were rendered in the foreground as the monkey cycled. The
396 distance between them reflected the difference between actual and instructed speed.
397 Obtaining juice required aligning the two floating targets while progressing towards a final
398 target, on which he stopped to obtain additional reward. The task was divided into trials, each
399 of which required moving a distance equivalent to twenty cycles under manual control. We
400 used eight trial-types, four each for forward and backward cycling. Two of these employed a
401 constant target speed (equivalent to 1 or 2 Hz cycling) and two involved a ramping speed (from
402 1 Hz to 2 Hz or vice versa). As above, the decoder was trained based on a small number of
403 manual-control trials performed at the beginning of each session. Blocks of manual-control
404 trials were also included for comparisons between manual and BMI-based performance.

405 Our decode strategy was largely preserved from that described above. However, we used a
406 modified state machine (**Supp. Fig. 1**) and a slightly different algorithm for transforming
407 rotations of the neural state into decoded virtual velocity. Direction was determined based on
408 which distribution (forward or backward) produced the higher likelihood of observing the
409 measured angular momentums (as in **Fig. 3d**). Once that choice was made, speed was
410 determined by the angular velocity of the neural state in that plane. Thus, faster rotational
411 trajectories led to faster decoded virtual velocity. We chose a scaling factor so that a given
412 neural angular velocity produced the speed that would have been produced by physical
413 cycling at that angular velocity. Neural angular velocity was exponentially filtered with a time
414 constant of 500 ms. The filter memory was erased on entry into a movement state (EARLY or
415 STEADY) from a stopped state (INIT or EXIT) to allow brisk movement onset (see *Methods*).

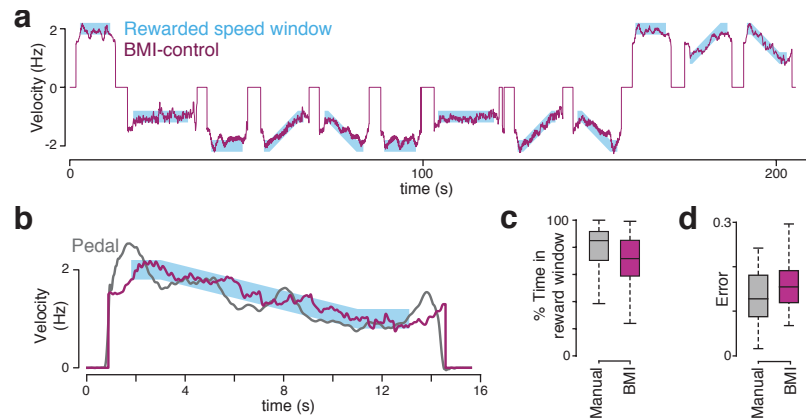


Figure 8. Performance in the modified task requiring speed tracking. (a) Instructed velocity and BMI-decoded virtual velocity during 12 contiguous trials of BMI control. **(b)** Expanded view of one example trial (the last trial from panel a). The virtual velocity that would have been produced by the 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).

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

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

433 **Discussion**

434 ***Task-specific subspace decoding***

435 We have previously argued that the largest signals in motor cortex are not ‘representational’ –
436 they do not encode variables but are instead essential for noise-robust dynamics. Those
437 dynamics produce outgoing commands that *are* representational (they covary with the
438 variables they control) but are low-variance. Because the dominant neural signals must be
439 produced by noise-robust dynamics, they obey the property of low trajectory tangling: similar
440 neural states avoid being associated with dissimilar derivatives (Russo et al., 2018). There is
441 little guarantee that the way in which low tangling is achieved will be the same across tasks –
442 on the contrary, the task often determines the most natural ways to keep tangling low. Thus,
443 the identity and number of high-variance dimensions will likely vary across tasks, as will the
444 basic response features captured by those dimensions. Put differently, both the intrinsic
445 covariance structure of the neural activity, and the nature of its correlations with kinematics,
446 will likely be task specific.

447 In this study we demonstrated that overtly task-specific decoding can be successful. This was
448 true even though we identified subspaces very differently from most reach-based BMI
449 approaches, and converted neural activity to movement using different (and often more non-
450 linear) approaches. These differences were critical. Unlike reaching we found no high-variance
451 dimensions where neural activity correlated with kinematics. Yet by using noncoding features
452 in high-variance dimensions, we achieved BMI control that was sufficiently natural that
453 monkeys appeared not to notice that the task was no longer under manual control. By most
454 measures (success rate, time to target) performance under BMI control was remarkably close
455 to that under manual control. The main limitation of BMI control was stopping accuracy.
456 Although our algorithm detected stopping with ~ 0.1 second precision, even small
457 discrepancies could lead to the target being over or undershot by a noticeable amount. A
458 beneficial feature of our BMI decode is that it almost never produced movement when it was
459 not intended. With rare exceptions, truly zero velocity was decoded when the monkey was
460 intending to remain stopped on the target. We consider this a particularly important attribute of
461 any self-motion decoding algorithm, due to the potentially large consequences of unintended
462 movement of the whole body.

463 ***Population dynamics in decoding***

464 The present approach relates to recent studies that modeled neural dynamics to improve
465 online (Kao et al., 2015) or offline (Aghagolzadeh and Truccolo, 2016; Gallego et al., 2020; Kao
466 et al., 2017b) decoding of movement kinematics. A key insight of those studies is that signals
467 that do not correlate directly with kinematics can be used to infer those that do. For linear
468 decoding, the value of a given variable depends upon the neural state in one dimension: the
469 dimension defined by the regression weights. Nevertheless, inferring the neural state in that
470 dimension may benefit from a dynamical model that spans multiple dimensions. Much like the
471 present approach, this allows the decode to leverage features that are robust, even if they do
472 not directly correlate with the kinematic parameters of interest. The present approach extends
473 this idea to situations where there may be no high-variance dimensions that can be linearly
474 decoded, and/or where the most prominent features are not well-described by linear dynamics.

475 Another way of leveraging additional dimensions – or even the full neural state – would be to
476 use a trained recurrent neural network (RNN) to decode parameters of interest (Sussillo et al.,
477 2012; Willett et al., 2020). This is a promising approach, but also has known drawbacks. For
478 safety in clinical implementation, it may be crucial to know which features the decoder is

479 actually exploiting, and to understand how it uses them. We therefore decided to work with
480 hand-selected features. While these features will not generalize to different tasks, that is in
481 some sense the point: the presence of task-specific covariance structure essentially ensures
482 that any decoding approach will need to be task-specific.

483 ***Self-motion BMIs***

484 In choosing an alternative to traditional reach-based BMI tasks, we had three reasons for
485 focusing on a BMI for virtual self-motion. First, our recently developed cycling task naturally
486 lends itself to this application. Second, BMI-controlled self-motion is likely to be desired by a
487 large patient population (potentially much larger than the population that desires BMI-
488 controlled cursors or robot arms). Third, prior work has demonstrated that BMI control of self-
489 motion is viable (Libedinsky et al., 2016; Rajangam et al., 2016). In particular, Rajangam et al.
490 (Rajangam et al., 2016) demonstrated BMI control of a physical wheelchair based on neural
491 activity recorded from monkey motor and somatosensory cortex. Our work supports their
492 conclusion that BMI-controlled self-motion is possible, and demonstrates the feasibility of an
493 alternative decode strategy. Rajangam et al. employed a traditional decode strategy: linear
494 filters transformed neural activity into the key variables: translational and angular velocity. That
495 strategy allowed monkeys to navigate ~2 meters to a target (which had to be approached with
496 an accuracy of +/- ~0.2 meters, or 10% of the distance traveled) in an average of 27-49
497 seconds (depending on the monkey and degree of practice). In our task, monkeys had to stop
498 with similar relative accuracy: +/- 0.5 cycles, or 7% of the distance traveled for a seven-cycle
499 movement. They traversed those seven cycles in ~4 seconds under BMI control (averages of
500 4.3 and 3.7 seconds for monkey G and E). While this is roughly tenfold faster, we stress that
501 movement durations are not directly comparable between our task and theirs. Success
502 requirements differed in multiple ways. For example, Rajangam et al. required that monkeys
503 turn *en route* (which adds considerable challenge) but did not require them to stop on the
504 target location. Yet while direct comparison is not possible, a tenfold improvement in time-to-
505 target argues that task-specific subspace decode strategies can be effective and should be
506 explored further.

507 An obvious limitation of the current study is that we did not explore strategies for steering,
508 which would be essential to a real-world self-motion prosthetic. There exist multiple candidate
509 strategies for enabling steering. Rajangam et al. used a Wiener filter to decode angular velocity
510 of the body. While straightforward, this strategy appears to have had limited success: even

511 during training, the R^2 of their angular velocity decode was 0.16 and 0.12 for the two monkeys.
512 One alternative strategy would be to apply our decode strategy bilaterally, and employ a
513 comparison (e.g., between left and right cycling speed) to control angular velocity. This
514 strategy should be viable even with a unilateral implant; recent work has shown that
515 information about both forelimbs can be decoded equally well from either hemisphere (Ames
516 and Churchland, 2019; Heming et al., 2019). Another strategy would be to control translational
517 velocity using the strategies developed here, but use a reach-like decode for steering (rather
518 like pedaling a bicycle while also steering). Which (if any) of these three strategies is preferable
519 remains a question for future experiments.

520 **Methods**

521 *Subjects and primary task*

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

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

552 In BMI control, trial parameters and failure conditions were the same as in manual control, for
553 purposes of comparison. The only difference between manual and BMI control was that, in the
554 latter, position in the virtual environment was controlled by the output of a decoder rather than
555 the pedal. We did not prevent or discourage the monkey from cycling during BMI-control
556 blocks, and he continued to do so as normal. In BMI control, monkey G performed an average

557 of 654 trials/session over 20 sessions and monkey E performed an average of 137
558 trials/session over 17 sessions. Manual-control data for monkey G (average of 229
559 trials/session over 8 sessions) were collected during sessions in which BMI-control data sets
560 were also collected. Manual-control sessions for monkey E (average of 231 trials/session over
561 5 sessions) were interleaved with BMI-control sessions on different days. For monkey G, an
562 additional three manual-control sessions (189, 407, and 394 trials) were employed to record
563 EMG, which was used for the variance captured analysis (**Fig. 2**). We recorded from 5-7
564 muscles per session, yielding a total of 19 recordings. We made one or more recordings from
565 the three heads of the *deltoid*, the lateral and long heads of *triceps brachii*, the *biceps brachii*,
566 *trapezius*, and *latissimus dorsi*. These muscles were selected due to their clear activations
567 during the cycling task.

568 *Surgery and neural/muscle recordings*

569 Neural activity was recorded using chronic 96-channel Utah arrays (Blackrock Microsystems,
570 Salt Lake City, UT), implanted in the left hemisphere using standard surgical techniques. In
571 each monkey, an array was placed in the region of primary motor cortex (M1) corresponding to
572 the upper arm. In monkey G, a second array was placed in dorsal premotor cortex (PMd), just
573 anterior to the first array. Array locations were selected based on MRI scans and anatomical
574 landmarks observed during surgery. Experiments were performed 1-8 months (monkey G) and
575 3-4 months (monkey E) after surgical implantation. Neural responses both during the task and
576 during palpation confirmed that arrays were in the proximal-arm region of cortex.

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

584 Intra-muscular EMG recordings were made using pairs of hook-wire electrodes inserted with
585 30 mm x 27 gauge needles (Natus Neurology, Middleton, WI). Raw voltages were amplified and
586 filtered (band-pass 10 Hz – 10 kHz) with ISO-DAM 8A modules (World Precision Instruments,
587 Sarasota, FL), and digitized at 30 kHz with the Cerebus Neural Signal Processors. EMG was
588 then digitally band-pass filtered (50 Hz – 5 kHz) prior to saving for offline analysis. Offline, EMG
589 recordings were rectified, low-pass filtered by convolving with a Gaussian (standard deviation:
590 25 ms), downsampled to 1 kHz, and then fully normalized such that the maximum value
591 achieved on each EMG channel was 1.

592 A real-time target computer (Speedgoat, Bern, CH) running Simulink Real-Time environment
593 (MathWorks, Natick, MA) processed behavioral and neural data and controlled the decoder

594 output in online experiments. It also streamed variables of interest to another computer that
595 saved these variables for offline analysis. Stateflow charts were implemented in the Simulink
596 model to control task state flow as well as the decoder state machine. Real-time control had
597 millisecond precision.

598 Spike trains were causally converted to firing rates by convolving each spike with a beta kernel.
599 The beta kernel was defined by temporally scaling a beta distribution (shape parameters: $\alpha = 3$
600 and $\beta = 5$) to be defined over the interval $[0, 275]$ ms and normalizing the kernel such that the
601 firing rates would be in units of spikes/second. The same filtering was applied for online
602 decoding and offline analyses. Firing rates were also mean centered (subtracting the mean rate
603 across all times and conditions) and normalized. During online decoding, the mean and
604 normalization factor were values that had been computed from the training data. We used soft
605 normalization (Russo et al., 2018): the normalization factor was the firing rate range plus a
606 constant (5 spikes/s).

607 *Computing trial-averaged firing rates*

608 Analyses of BMI performance are based on real-time decoding during online performance, with
609 no need to consider trial-averaged firing rates. However, we still wished to compute trial-
610 averaged traces of neural activity and kinematics for two purposes. First, some aspects of
611 decoder training benefited from analyzing trial-averaged firing rates. Second, we employ
612 analyses that document basic features of single-neuron responses and of the population
613 response (e.g., **Fig. 1d**, **Fig. 3a**, **Fig. 6a**, **Fig. 7a**). These analyses benefit from the denoising
614 that comes from computing a time-varying firing rate across many trials. Due to the nature of
615 the task, trials could be quite long (up to 20 cycles in the speed-tracking task), rendering the
616 traditional approach of aligning all trials to movement onset insufficient for preserving
617 alignment across all subsequent cycles. It was thus necessary to modestly adjust the time-
618 base of each individual trial (e.g., stretching time slightly for a trial where cycling was faster
619 than typical). We employed two alignment methods. Method A is a simplified procedure that
620 was used prior to parameter fitting when training the decoder before online BMI control. This
621 method aligns only times during the movement. Method B is a more sophisticated alignment
622 procedure that was utilized for all offline analyses. This method aligns the entire trial, including
623 pre- and post-movement data. For visualization, conditions with the same target distance (e.g.,
624 7 cycles), but different directions, were also aligned to the same time base. Critically, any data
625 processing that relied on temporal structure was completed in the original, unstretched time
626 base prior to alignment.

627 **Method A:** The world position for each trial resembles a ramp between movement onset and
628 offset (**Fig. 1a**). First, we identify the portion of each trial starting $\frac{1}{4}$ cycle into the movement
629 and ending $\frac{1}{4}$ cycle before the end of the movement. We fit a line to the world position in this
630 period and then extend that line until it intercepts the starting and ending positions. The data

631 between these two intercepts is considered the movement data for each trial and is extracted.
632 This movement data is then uniformly stretched in time to match the average trial length for
633 each trial's associated condition. This approach compresses slower than average movements
634 and stretches faster than average movements within a condition, such that they can be
635 averaged while still preserving many of the cycle-specific features of the data.

636 **Method B:** This method consists of a mild, non-uniform stretching of time in order to match
637 each trial to a condition-specific template. For complete details, see Russo et al. 2018. (Russo
638 et al., 2018)

639 *Variance captured analysis*

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

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

661 We wished to compute, for each dimension, the percentage of neural variance explained – i.e.,
662 whether that dimension captured large or small signals. We were not interested in whether
663 dimensions captured stochastic spiking variability, but in whether they captured large features
664 that were reliable across trials. Thus, variance captured was always computed based on trial-
665 averaged neural responses. We considered the matrix $\bar{Y} \in \mathbb{R}^{N \times T}$ where T is the total number of
666 time points across all conditions. Each row of \bar{Y} contains the trial-averaged firing rate of one

667 neuron. We computed an $N \times N$ covariance matrix $\Sigma = \text{cov}(\bar{Y})$ by treating rows of \bar{Y} as random
668 variables and columns as observations. The proportion of total neural variance captured by a
669 given dimension, \mathbf{w} , is therefore:

$$670 \quad \frac{\mathbf{w}^T \Sigma \mathbf{w}}{\text{tr}(\Sigma)}$$

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

674 *Decoding from kinematic dimensions*

675 Offline decodes (**Fig. 2c,d**) were performed using the same arm-control datasets described in
676 the previous section, but in this analysis the data were split into training and testing sets. The
677 training set for a session comprised the first 25 forward and 25 backward trials of that day (the
678 same training set size used in our online decodes). Regression dimensions were found for
679 these training sets as described above, and coefficients of determination were calculated on
680 the test sets.

681 *Identifying neural dimensions*

682 Although the response features leveraged by the decode algorithm are clearly visible in the top
683 principal components of the data (when PCA is performed on the full trial-averaged time-series
684 of firing rates across conditions), we sought to choose neural dimensions that would cleanly
685 isolate particular features. To this end, each feature was isolated using dedicated
686 preprocessing and dimensionality reduction approaches.

687 We sought a moving-sensitive dimension, the projection onto which would allow an HMM to
688 estimate the probability of moving, p_{move} , at each moment. To do so, we first computed binned
689 spike-counts (10 ms for monkey G, 20 ms for monkey E) and applied a square-root transform
690 to these counts as this has been shown to improve the Gaussian fit for Poisson data with small
691 counts (Thacker and Bromiley, 2001). We then aggregated all of these square-rooted binned
692 counts from the training set (25 forward trials, 25 backward trials) and separated them into two
693 classes based on pedaling speed: ‘moving’ (speed > 1 Hz) and ‘stopped’ (speed < .05 Hz).
694 Samples that didn’t fall into either of these two classes were discarded. We applied linear
695 discriminant analysis to these two labeled sets, which yielded a discriminating hyperplane that
696 best separated the two classes. We defined the moving-sensitive dimension, \mathbf{w}_{move} , as the
697 vector normal to this hyperplane.

698 In order to decode direction, we sought to isolate four neural dimensions that captured
699 rotational trajectories during steady-state cycling. Spike time-series were filtered to yield firing
700 rates (as described above), and then further high-pass filtered (2nd order Butterworth, cutoff

701 frequency: 1 Hz). This removed drift or other low-frequency signals. Single-trial movement-
702 period responses were then aligned (Method A) and averaged within conditions to generate
703 $N \times T_c$ matrices \bar{Y}_f and \bar{Y}_b . We sought a 4-dimensional projection of these trial-averaged
704 responses that would maximally capture rotational trajectories while segregating forward and
705 backward data into different planes. Whereas the standard PCA cost function finds dimensions
706 that maximize variance captured, we opted instead for a cost function that would maximize the
707 difference in variance captured between the two conditions:

$$708 \quad J(W) = \text{tr}(W^T \Sigma_f W) - \text{tr}(W^T \Sigma_b W)$$

709 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
710 function will be maximized when the projection of the data captures a great deal of variance for
711 forward trials and very little variance for backward trials. Conversely, this cost function will be
712 minimized when the projection favors large variances for backward trials and small variances
713 for forward trials. We thus chose to define our forward rotational plane by the 2D matrix $W_f =$
714 $\begin{bmatrix} \mathbf{w}_f^{(1)} & \mathbf{w}_f^{(2)} \end{bmatrix}$ that maximizes $J(W)$ and our backward rotational plane by the 2D matrix $W_b =$
715 $\begin{bmatrix} \mathbf{w}_b^{(1)} & \mathbf{w}_b^{(2)} \end{bmatrix}$ that minimizes $J(W)$. An iterative optimization procedure was used to find W_f and
716 W_b with the Manopt toolbox (Boumal et al., 2014); full details of this in (Cunningham and
717 Ghahramani, 2015).

718 To decode direction during the EARLY state, we found a set of initial-direction dimensions. We
719 used activity in the moving-sensitive dimension to determine the time, t_{init} , at which the state
720 machine would have entered the INIT state during online control. We then considered trial-
721 averaged neural activity, for each condition, from t_{init} through $t_{init} + 200$ ms. We applied PCA
722 and retained the top three dimensions: $\mathbf{w}_{dir}^{(1)}$, $\mathbf{w}_{dir}^{(2)}$, and $\mathbf{w}_{dir}^{(3)}$. Such dimensions capture how
723 activity evolves both across that timespan, and how it differs across forwards and backwards
724 cycling conditions.

725 *Computing probability of moving (p_{move})*

726 To compute p_{move} based on neural activity in the moving-sensitive dimension, an HMM was
727 used to track two states: ‘moving’ or ‘stopped’ (Kao et al., 2017a). Square-rooted spike counts
728 in the training data were already separated into ‘moving’ and ‘stopped’ sets for the purposes of
729 identifying \mathbf{w}_{move} . We projected those counts onto \mathbf{w}_{move} and a fit Gaussian distribution for
730 each state. The probability, p_{move} , of being in the ‘moving’ state, given the entire sequence of
731 current and previously observed square-rooted spike counts, was computed efficiently with a
732 recursive algorithm that uses the state transition matrix

$$733 \quad \Phi = \begin{bmatrix} p_{move|move} & p_{move|stop} \\ p_{stop|move} & p_{stop|stop} \end{bmatrix}$$

734 and knowledge of the Gaussian distributions. Φ encodes prior assumptions about the
 735 probability of transitioning from one state to the next at any given bin. We used a benchmark
 736 set of manual-control training data from each monkey to determine reasonable values for Φ ,
 737 which were then used in all experiments. For monkey G, we set $p_{move|stop} = .0001$ and
 738 $p_{stop|move} = .002$; for monkey E, we set $p_{move|stop} = .0002$ and $p_{stop|move} = .004$. The value
 739 p_{move} was used throughout the decoder state machine to control transitions between various
 740 states, effectively dictating the movement onset and offset behavior of the decoder (**Fig. 6d**).

741 *Computing steady-state direction and speed*

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

$$745 \quad x_t = Ax_{t-1} + q_t$$

$$746 \quad y_t = Cx_t + r_t$$

747 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
 748 neural state in the rotational dimensions and y_t are the high-pass filtered firing rates, which we
 749 treat as noisy measurements of that underlying state. We chose to let our measurements be
 750 smooth firing rates, rather than use non-overlapping bins of spikes, for purely opportunistic
 751 reasons: it consistently yielded better performance by our decoder. The parameters of the
 752 Kalman filter were fit to the training data as follows:

$$753 \quad A = \bar{X}_2 \bar{X}_1^T (\bar{X}_1 X_1^T)^{-1}$$

$$754 \quad C = \begin{bmatrix} W_f^T \\ W_b^T \end{bmatrix}^\dagger$$

$$755 \quad Q = \text{cov}(\bar{X}_2 - A\bar{X}_1)$$

$$756 \quad R = \text{cov}(Y - C\bar{X})$$

757 where

$$758 \quad \bar{X}_1 = \begin{bmatrix} W_f^T \\ W_b^T \end{bmatrix} [\bar{Y}_f(:, 1: T_f - 1), \bar{Y}_b(:, 1: T_b - 1)]$$

$$759 \quad \bar{X}_2 = \begin{bmatrix} W_f^T \\ W_b^T \end{bmatrix} [\bar{Y}_f(:, 2: T_f), \bar{Y}_b(:, 2: T_b)]$$

$$760 \quad \bar{X} = \begin{bmatrix} W_f^T \\ W_b^T \end{bmatrix} [\bar{Y}_1, \bar{Y}_2, \dots, \bar{Y}_{50}]$$

$$761 \quad Y = [Y_1, Y_2, \dots, Y_{50}]$$

762 with Y_i denoting the neural activity (high-pass filtered firing rates) for the i -th trial in the training
 763 set, \bar{Y}_i denoting the trial-averaged activity for the condition that the i -th trial is an instantiation
 764 of, \dagger denoting the Moore-Penrose pseudoinverse, and the colon symbol designating how to
 765 index matrices (e.g., $M(:, a:b)$ refers to the submatrix of M including all rows of M , but only the
 766 columns a through b). Lastly, the initial state parameter x_0 was computed by taking the
 767 average value of the trial-averaged projections over all times and conditions. Online inference
 768 of the underlying neural state, which yields an estimate \hat{x}_t at each millisecond t , was computed
 769 recursively using the steady-state form of the Kalman filter (Malik et al., 2011).

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

$$773 \quad L(t) = \begin{bmatrix} L_f(t) \\ L_b(t) \end{bmatrix} = \begin{bmatrix} \hat{x}_{t-1}^{(1)} \hat{x}_t^{(2)} - \hat{x}_t^{(1)} \hat{x}_{t-1}^{(2)} \\ \hat{x}_{t-1}^{(3)} \hat{x}_t^{(4)} - \hat{x}_t^{(3)} \hat{x}_{t-1}^{(4)} \end{bmatrix}$$

774 where the superscript indexes the elements of \hat{x}_t . We fit 2D Gaussian distributions to these
 775 angular momentums for each of three behaviors in the training data: ‘stopped’ (speed < .05
 776 Hz), ‘pedaling forward’ (velocity > 1 Hz), and ‘pedaling backward’ (velocity < -1 Hz) (**Fig. 3d**).
 777 Online, the likelihood of the observed angular momentums with respect to each of these three
 778 distributions dictated the steady-state estimates of direction and speed. We’ll denote these
 779 three likelihoods f_{stop} , $f_{forward}$, and $f_{backward}$.

780 In general, one can compute which of these three distributions is most likely by choosing the
 781 maximizing likelihood and assess confidence in that choice by comparing the relative values of
 782 the three likelihoods. However, we wanted the decoder to err on the side of withholding
 783 movement. We therefore set a conservative threshold on f_{stop} corresponding to the point at
 784 which L would have a Mahalanobis distance of 3 to the stopped distribution of angular
 785 momentums. If f_{stop} ever exceeded this threshold, we set $speed_{steady}$ to zero. If this condition
 786 was not met, we decoded direction and speed as follows:

$$787 \quad direction_{steady}(t) = \text{sgn} \left(f_{forward}(t) - f_{backward}(t) \right)$$

$$788 \quad speed_{steady}(t) = \left| 2 \cdot \frac{f_{forward}(t)}{f_{forward}(t) + f_{backward}(t)} - 1 \right| \beta$$

789 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
 790 the likelihoods (yielding a slower velocity if the direction decode is uncertain) and β is a
 791 direction-specific constant learned from the training data whose purpose is simply to scale up
 792 the result to match steady-state cycling speed. In practice, $speed_{steady}$ was frequently very
 793 close to the monkeys’ steady-state cycling speeds (**Fig. 3e**).

794 *Computing initial direction and speed*

795 Initial direction and speed were always computed at the moment the EARLY state was entered,
796 t_{early} . These values then persisted throughout the remainder of the EARLY state. Given that
797 the decoder state machine doesn't make use of the initial-direction dimensions prior to
798 entering the EARLY state, t_{early} can be computed for the training trials. Single-trial firing rates
799 from the training set were then projected onto the initial-direction dimensions at t_{early} and 3D
800 Gaussian distributions were fit to the resulting sets of forward and backward neural states.
801 Online, firing rates were projected onto the initial-direction dimensions at t_{early} and likelihoods
802 $g_{forward}$ and $g_{backward}$ were computed with respect to each the learned distributions. If the
803 observed neural state in the initial-direction subspace was not an outlier (>10 Mahalanobis
804 distance units) with respect to both distributions, then the initial direction and speed were
805 computed as follows:

$$806 \quad direction_{initial}(t_{early}) = \text{sgn}\left(g_{forward}(t_{early}) - g_{backward}(t_{early})\right)$$

$$807 \quad speed_{initial}(t_{early}) = \left| 2 \cdot \frac{g_{forward}(t_{early})}{g_{forward}(t_{early}) + g_{backward}(t_{early})} - 1 \right| \beta$$

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

810 *Smoothing of decoded velocity*

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

$$815 \quad v'_{dec}(t) = \frac{1}{T_{smooth} + 1} \sum_{i=0}^{T_{smooth}} v_{dec}(t - i)$$

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

821 *Speed-tracking task*

822 In addition to the primary task (where the monkey traveled 2-7 cycles between stationary
823 targets) we employed a speed-tracking task, in which the monkey was required to match his
824 virtual speed to an instructed speed. Speed was instructed implicitly, via the relative position of

825 two moving targets. The primary target was located a fixed distance in front of the monkey's
826 location in virtual space: the secondary target fell 'behind' the first target when cycling was too
827 slow, and pulled 'ahead' if cycling was too fast. This separation saturated for large errors, but
828 for small errors was proportional to the difference between the actual and instructed speed.
829 This provided sufficient feedback to allow the monkey to track the instructed speed even when
830 it was changing. Because there was no explicit cue regarding the absolute instructed speed,
831 monkeys began cycling on each trial unaware of the true instructed speed profile and 'honed
832 in' on that speed over the first ~2 cycles.

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

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

850 Ramping profiles began with three seconds of constant instructed speed to allow the monkey
851 to hone in on the correct initial speed. Instructed speed then ramped, over 8 seconds, to a new
852 value, and remained constant thereafter. As for constant profiles, speed-instructing cues
853 disappeared after 18 cycles and the monkey cycled two further cycles before stopping on a
854 final target. Again, analyses of performance were based on the period from when the monkey
855 first honed in on the correct speed, to when the speed-instructing cues disappeared. There
856 were two ramping profiles: one ramping up from 1 to 2 Hz, and one ramping down from 2 to 1
857 Hz. There were thus four total speed profiles (two constant and two ramping). These were
858 performed for both cycling directions (presented in blocks and instructed by color as in the
859 primary task) yielding eight total conditions. This task was only performed by monkey G, who
860 completed an average of 166 trials/session over 2 sessions in manual control and an average
861 of 116 trials/session over 3 sessions in BMI control.

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

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

879 *Neural features for speed-tracking*

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

884 The probability of moving, p_{move} , was calculated using a different set of parameters for speed-
885 tracking, largely due to changes in recording quality in the intervening time between data
886 collection from the primary experiment and data collection for the speed-tracking experiment.
887 The bin size was increased to 100 ms and the following state transition values were used:
888 $p_{move|stop} = .0005$ and $p_{stop|move} = .0005$. In addition, we observed that the square-root
889 transform seemed to be having a negligible impact on the quality of the decoder at this bin
890 size, so we removed it for this task.

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

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

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

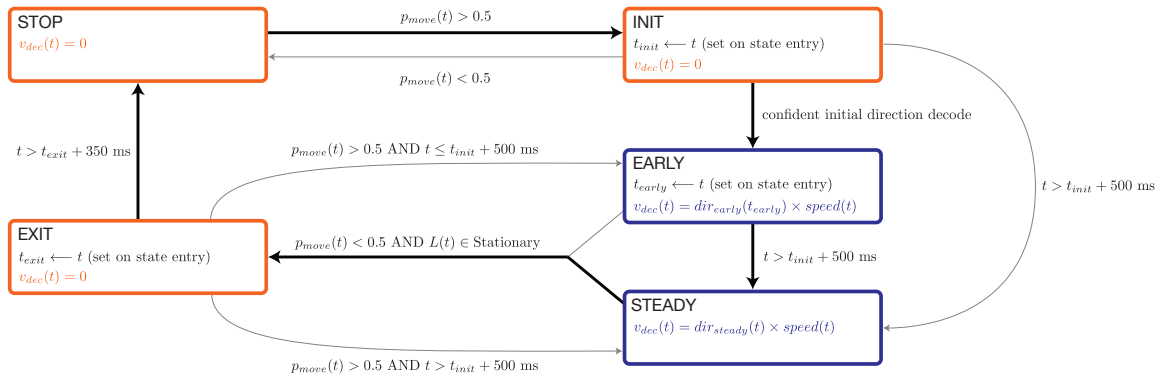
$$911 \quad \theta'(t) = \begin{cases} \frac{d\theta_f}{dt}, & direction(t) = +1 \\ \frac{d\theta_b}{dt}, & direction(t) = -1 \end{cases}$$

912 where $\theta_f(t)$ and $\theta_b(t)$ are the phases of the two planes in the neural state estimate \hat{x}_t ,
913 $direction$ corresponds to $direction_{early}$ while in the EARLY state and $direction_{steady}$ while in
914 the STEADY state, and the derivative θ' is computed in units of Hz. The coarse speed estimate,
915 θ' , was then smoothed with an exponential moving average ($\tau = 500$ ms) to generate $speed$, the
916 variable that gets used in the decoder state machine. Additional saturation limits were set such
917 that $speed$ never dropped below 0.5 Hz or exceeded 3.5 Hz, so as to remain in the range
918 typically seen during pedaling. On entry into EARLY or STEADY from either INIT or EXIT, when
919 $speed$ gets initialized, the output of this exponential moving average was reset to an initial
920 value of 1.5 Hz, which was the average starting speed across conditions.

921 Lastly, there were two new conditions for decoder state transitions in the speed-tracking
922 experiment (**Supp. Fig. 1**). First, transitions from INIT to EARLY required that a condition
923 termed “confident initial direction decode” was obtained. This condition was met when the
924 Mahalanobis distance from the neural state in the initial-direction subspace to either the
925 forward or backward distributions dropped below 4. Second, transitions into the EXIT state
926 required (in addition to a drop in p_{move}) that the observed angular momentums, L , belong to a
927 set termed ‘Stationary’. This set was defined as all L with a Mahalanobis distance of less than
928 4 to the ‘stopped’ distribution of angular momentums, which was learned from the training set.

929

930



Supplementary Figure 1. State machine diagram for speed-control experiments. BMI motion for the speed-tracking experiment was determined by a state machine with five states: STOP, INIT, EARLY, STEADY, and EXIT. Each of these states describes how to compute virtual velocity, v_{dec} , which gets integrated every millisecond into virtual position. Following the conventions from Figure 4, black arrows indicate the typical path of a successful BMI trial (gray arrows indicate all other transitions) and colors differentiate states in which there is progress through the virtual environment (blue) from states in which BMI motion is zero (orange). In the STOP, INIT, and EXIT states, virtual velocity is simply set to zero. During the EARLY and STEADY states, virtual velocity is computed as the product of decoded direction and decoded speed, which are computed based on activity in the rotational dimensions. Transitions between the states were determined by p_{move} (computed from the moving-sensitive dimension), angular momentum vector L (computed from the rotational planes), and the neural state in the initial-direction dimensions. Details on computing each of the neural features used in this state machine are provided in the section ‘Neural features for speed-tracking’. While in the STOP state, virtual velocity remains zero. If p_{move} goes high, the INIT state is entered (time of entry is denoted t_{init}), but virtual velocity remains zero. From the INIT state, one of three things can happen: 1) p_{move} goes low and the movement aborts with a transition back to the STOP state, 2) the neural state in the initial-direction dimensions becomes sufficiently close to the learned distributions of neural states for forward or backward conditions, triggering a transition to the EARLY state, or 3) no transition occurs within 500 ms of t_{init} , triggering a transition to the STEADY state. The INIT state serves to withhold BMI motion in situations where there is a transient spike in p_{move} and to delay movement when the early estimate of decoded direction is only weakly confident. The EARLY state dictates how virtual velocity is computed in the early portion of movement (when less than 500 ms have elapsed since t_{init}). The STEADY state determines virtual velocity for the remainder of decoded movement. While in the EARLY or STEADY states, if p_{move} goes low and the angular momentums in the rotational planes are consistent with a stationary pedal, a transition will occur to the EXIT state and virtual velocity will again go to zero. After 350 consecutive ms in the EXIT state, the state machine transitions back to the STOP state. However, if p_{move} goes high while in the EXIT state, a transition back to either EARLY or STEADY will occur. The transitions from EXIT back to EARLY and STEADY

exist to allow a rapid return to decoded movement if neural activity briefly, erroneously reflects that the monkey is ceasing movement, but then returns to exhibiting robust activity consistent with an ongoing movement.

Supplementary Movie 1. MI and manual control of primary task (monkey G, 10/4/2018). Both manual (left) and BMI control (right) sessions were performed on the same day, which was the fourth day of BMI control. Manual and BMI control success rates were very similar on this day (see Figure 3d). The task required that the monkey turn the hand pedal in the instructed direction to move through the virtual environment, stop on top of a lighted target, and remain still while collecting juice reward. The color of the landscape indicated whether cycling must be 'forward' (green landscape) or 'backward' (tan landscape). There were 6 total conditions, defined by cycling direction and target distance (2, 4, or 7 cycles).

Supplementary Movie 2. BMI and manual control of speed-tracking task (monkey G, 4/18/2019 and 3/27/2019, respectively). These examples were selected as representative of performance in this task. In this version of the task, the monkey was required to track various instructed speed profiles (constant, ramping up, and ramping down) as he progressed through the environment. A red target was rendered a short distance in front of the monkey, and remained at that distance ahead throughout the trial. A second, white target drifted ahead or behind the red target to indicate the error between the avatar's speed and the instructed speed. Thus, successful speed tracking required aligning these targets on top of one another. Whenever the error was within the rewarded tolerance window, the white target would light up and juice reward would be dispensed. This speed-tracking epoch lasted until the monkey was within two cycles of the final turquoise stopping target (total of 20 cycles from starting position).

931 **Acknowledgements**

932 We thank Y. Pavlova for expert animal care, A. Russo for sharing code and data for preliminary
933 analysis, and E. Oby for surgical expertise and assistance. This work was supported by NINDS
934 1DP2NS083037, NIH CRCNS R01NS100066, NINDS 1U19NS104649, the Simons Foundation
935 (SCGB#325233 and SCGB#542957), the Grossman Center for the Statistics of Mind, the
936 McKnight Foundation, P30 EY019007, a Klingenstein-Simons Fellowship, and the Searle
937 Scholars Program.

938

939 **Author Contributions**

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

944

945 **References**

- 946 Aghagolzadeh M, Truccolo W. 2016. Inference and Decoding of Motor Cortex Low-Dimensional
947 Dynamics via Latent State-Space Models. *IEEE Trans Neural Syst Rehabil Eng* **24**:272–
948 282. doi:10.1109/TNSRE.2015.2470527
- 949 Ajiboye AB, Willett FR, Young DR, Memberg WD, Murphy BA, Miller JP, Walter BL, Sweet JA,
950 Hoyen HA, Keith MW, Peckham PH, Simeral JD, Donoghue JP, Hochberg LR, Kirsch
951 RF. 2017. Restoration of reaching and grasping movements through brain-controlled
952 muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration. *The*
953 *Lancet* **0**. doi:10.1016/S0140-6736(17)30601-3
- 954 Ames KC, Churchland MM. 2019. Motor cortex signals for each arm are mixed across
955 hemispheres and neurons yet partitioned within the population response. *eLife* **8**.
956 doi:10.7554/eLife.46159
- 957 Boumal N, Mishra B, Absil P-A, Sepulchre R. 2014. Manopt, a Matlab Toolbox for Optimization
958 on Manifolds. *J Mach Learn Res* **15**:1455–1459.
- 959 Carmena JM, Lebedev MA, Crist RE, O’Doherty JE, Santucci DM, Dimitrov DF, Patil PG,
960 Henriquez CS, Nicolelis MAL. 2003. Learning to Control a Brain–Machine Interface for
961 Reaching and Grasping by Primates. *PLoS Biol* **1**:e42. doi:10.1371/journal.pbio.0000042
- 962 Chapin JK, Moxon KA, Markowitz RS, Nicolelis MAL. 1999. Real-time control of a robot arm
963 using simultaneously recorded neurons in the motor cortex. *Nat Neurosci* **2**:664–670.
964 doi:10.1038/10223
- 965 Christie BP, Tat DM, Irwin ZT, Gilja V, Nuyujukian P, Foster JD, Ryu SI, Shenoy KV, Thompson
966 DE, Chestek CA. 2014. Comparison of spike sorting and thresholding of voltage
967 waveforms for intracortical brain–machine interface performance. *J Neural Eng*
968 **12**:016009. doi:10.1088/1741-2560/12/1/016009
- 969 Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV.
970 2012. Neural population dynamics during reaching. *Nature* **487**:51–56.
971 doi:10.1038/nature11129
- 972 Churchland MM, Shenoy KV. 2007. Temporal complexity and heterogeneity of single-neuron
973 activity in premotor and motor cortex. *J Neurophysiol* **97**:4235–4257.
974 doi:10.1152/jn.00095.2007

- 975 Churchland MM, Yu BM, Sahani M, Shenoy KV. 2007. Techniques for extracting single-trial
976 activity patterns from large-scale neural recordings. *Curr Opin Neurobiol*, Neuronal and
977 glial cell biology / New technologies **17**:609–618. doi:10.1016/j.conb.2007.11.001
- 978 Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJ,
979 Velliste M, Boninger ML, Schwartz AB. 2013. 7 degree-of-freedom neuroprosthetic
980 control by an individual with tetraplegia. *Lancet* **381**:557–564. doi:10.1016/S0140-
981 6736(12)61816-9
- 982 Cunningham JP, Ghahramani Z. 2015. Linear Dimensionality Reduction: Survey, Insights, and
983 Generalizations. *J Mach Learn Res* **16**:2859–2900.
- 984 Degenhart AD, Bishop WE, Oby ER, Tyler-Kabara EC, Chase SM, Batista AP, Yu BM. 2020.
985 Stabilization of a brain–computer interface via the alignment of low-dimensional spaces
986 of neural activity. *Nat Biomed Eng* **4**:672–685. doi:10.1038/s41551-020-0542-9
- 987 Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP. 2016. Reorganization
988 between preparatory and movement population responses in motor cortex. *Nat Commun*
989 **7**. doi:10.1038/ncomms13239
- 990 Ethier C, Oby ER, Bauman MJ, Miller LE. 2012. Restoration of grasp following paralysis through
991 brain-controlled stimulation of muscles. *Nature* **485**:368–371. doi:10.1038/nature10987
- 992 Ethier C, Sachs NA, Miller LE. 2011. Continuous state-dependent decoders for brain machine
993 interfaces2011 5th International IEEE/EMBS Conference on Neural Engineering.
994 Presented at the 2011 5th International IEEE/EMBS Conference on Neural Engineering.
995 pp. 473–477. doi:10.1109/NER.2011.5910589
- 996 Fernández-Rodríguez Á, Velasco-Álvarez F, Ron-Angevin R. 2016. Review of real brain-
997 controlled wheelchairs. *J Neural Eng* **13**:061001. doi:10.1088/1741-2560/13/6/061001
- 998 Fetz EE, Finocchio DV, Baker MA, Soso MJ. 1980. Sensory and motor responses of precentral
999 cortex cells during comparable passive and active joint movements. *J Neurophysiol*
1000 **43**:1070–1089.
- 1001 Gallego JA, Perich MG, Chowdhury RH, Solla SA, Miller LE. 2020. Long-term stability of cortical
1002 population dynamics underlying consistent behavior. *Nat Neurosci* **23**:260–270.
1003 doi:10.1038/s41593-019-0555-4
- 1004 Gallego JA, Perich MG, Miller LE, Solla SA. 2017. Neural Manifolds for the Control of
1005 Movement. *Neuron* **94**:978–984. doi:10.1016/j.neuron.2017.05.025
- 1006 Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE. 2018. Cortical population
1007 activity within a preserved neural manifold underlies multiple motor behaviors. *Nat*
1008 *Commun* **9**:4233. doi:10.1038/s41467-018-06560-z
- 1009 Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Churchland MM,
1010 Kaufman MT, Kao JC, Ryu SI, Shenoy KV. 2012. A high-performance neural prosthesis
1011 enabled by control algorithm design. *Nat Neurosci* **15**:1752–1757. doi:10.1038/nn.3265
- 1012 Gilja V, Pandarinath C, Blabe CH, Nuyujukian P, Simeral JD, Sarma AA, Sorice BL, Perge JA,
1013 Jarosiewicz B, Hochberg LR, Shenoy KV, Henderson JM. 2015. Clinical translation of a
1014 high-performance neural prosthesis. *Nat Med* **21**:1142–1145. doi:10.1038/nm.3953
- 1015 Golub MD, Sadtler PT, Oby ER, Quick KM, Ryu SI, Tyler-Kabara EC, Batista AP, Chase SM, Yu
1016 BM. 2018. Learning by neural reassociation. *Nat Neurosci* **1**. doi:10.1038/s41593-018-
1017 0095-3
- 1018 Heming EA, Cross KP, Takei T, Cook DJ, Scott SH. 2019. Independent representations of
1019 ipsilateral and contralateral limbs in primary motor cortex. *eLife* **8**:e48190.
1020 doi:10.7554/eLife.48190
- 1021 Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, Shenoy KV. 2015. Single-trial
1022 dynamics of motor cortex and their applications to brain-machine interfaces. *Nat*
1023 *Commun* **6**:1–12. doi:10.1038/ncomms8759

- 1024 Kao JC, Nuyujukian P, Ryu SI, Shenoy KV. 2017a. A High-Performance Neural Prosthesis
1025 Incorporating Discrete State Selection With Hidden Markov Models. *IEEE Trans Biomed*
1026 *Eng* **64**:935–945. doi:10.1109/TBME.2016.2582691
- 1027 Kao JC, Ryu SI, Shenoy KV. 2017b. Leveraging neural dynamics to extend functional lifetime of
1028 brain-machine interfaces. *Sci Rep* **7**. doi:10.1038/s41598-017-06029-x
- 1029 Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. 2014. Cortical activity in the null space:
1030 permitting preparation without movement. *Nat Neurosci* **17**:440–448.
1031 doi:10.1038/nn.3643
- 1032 Kaufman MT, Seely JS, Sussillo D, Ryu SI, Shenoy KV, Churchland MM. 2016. The Largest
1033 Response Component in the Motor Cortex Reflects Movement Timing but Not
1034 Movement Type. *eNeuro* **3**. doi:10.1523/ENEURO.0085-16.2016
- 1035 Kemere C, Santhanam G, Yu BM, Afshar A, Ryu SI, Meng TH, Shenoy KV. 2008. Detecting
1036 Neural-State Transitions Using Hidden Markov Models for Motor Cortical Prostheses. *J*
1037 *Neurophysiol* **100**:2441–2452. doi:10.1152/jn.00924.2007
- 1038 Lemon RN, Hanby JA, Porter R. 1976. Relationship between the activity of precentral neurones
1039 during active and passive movements in conscious monkeys. *Proc R Soc Lond B Biol*
1040 *Sci* **194**:341–373. doi:10.1098/rspb.1976.0083
- 1041 Libedinsky C, So R, Xu Z, Kyar TK, Ho D, Lim C, Chan L, Chua Y, Yao L, Cheong JH, Lee JH,
1042 Vishal KV, Guo Y, Chen ZN, Lim LK, Li P, Liu L, Zou X, Ang KK, Gao Y, Ng WH, Han
1043 BS, Chng K, Guan C, Je M, Yen S-C. 2016. Independent Mobility Achieved through a
1044 Wireless Brain-Machine Interface. *PLOS ONE* **11**:e0165773.
1045 doi:10.1371/journal.pone.0165773
- 1046 Makin JG, O’Doherty JE, Cardoso MMB, Sabes PN. 2018. Superior arm-movement decoding
1047 from cortex with a new, unsupervised-learning algorithm. *J Neural Eng* **15**:026010.
1048 doi:10.1088/1741-2552/aa9e95
- 1049 Malik WQ, Truccolo W, Brown EN, Hochberg LR. 2011. Efficient Decoding With Steady-State
1050 Kalman Filter in Neural Interface Systems. *IEEE Trans Neural Syst Rehabil Eng* **19**:25–
1051 34. doi:10.1109/TNSRE.2010.2092443
- 1052 Michaels JA, Dann B, Scherberger H. 2016. Neural Population Dynamics during Reaching Are
1053 Better Explained by a Dynamical System than Representational Tuning. *PLOS Comput*
1054 *Biol* **12**:e1005175. doi:10.1371/journal.pcbi.1005175
- 1055 Miri A, Warriner CL, Seely JS, Elsayed GF, Cunningham JP, Churchland MM, Jessell TM. 2017.
1056 Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor
1057 Cortex. *Neuron* **95**. doi:10.1016/j.neuron.2017.06.042
- 1058 Pruszynski JA, Kurtzer I, Nashed JY, Omrani M, Brouwer B, Scott SH. 2011. Primary motor
1059 cortex underlies multi-joint integration for fast feedback control. *Nature* **478**:387–390.
1060 doi:10.1038/nature10436
- 1061 Rajangam S, Tseng P-H, Yin A, Lehew G, Schwarz D, Lebedev MA, Nicolelis MAL. 2016.
1062 Wireless Cortical Brain-Machine Interface for Whole-Body Navigation in Primates. *Sci*
1063 *Rep* **6**. doi:10.1038/srep22170
- 1064 Russo AA, Bittner SR, Perkins SM, Seely JS, London BM, Lara AH, Miri A, Marshall NJ, Kohn
1065 A, Jessell TM, Abbott LF, Cunningham JP, Churchland MM. 2018. Motor Cortex Embeds
1066 Muscle-like Commands in an Untangled Population Response. *Neuron* **97**:953-966.e8.
1067 doi:10.1016/j.neuron.2018.01.004
- 1068 Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP.
1069 2014. Neural constraints on learning. *Nature* **512**:423–426. doi:10.1038/nature13665
- 1070 Schroeder KE, Irwin ZT, Bullard AJ, Thompson DE, Bentley JN, Stacey WC, Patil PG, Chestek
1071 CA. 2017. Robust tactile sensory responses in finger area of primate motor cortex
1072 relevant to prosthetic control. *J Neural Eng* **14**:046016. doi:10.1088/1741-2552/aa7329

- 1073 Scott SH, Gribble PL, Graham KM, Cabel DW. 2001. Dissociation between hand motion and
1074 population vectors from neural activity in motor cortex. *Nature* **413**:161–165.
1075 doi:10.1038/35093102
- 1076 Sergio LE, Hamel-Pâquet C, Kalaska JF. 2005. Motor Cortex Neural Correlates of Output
1077 Kinematics and Kinetics During Isometric-Force and Arm-Reaching Tasks. *J*
1078 *Neurophysiol* **94**:2353–2378. doi:10.1152/jn.00989.2004
- 1079 Serruya MD, Hatsopoulos NG, Paninski L, Fellows MR, Donoghue JP. 2002. Brain-machine
1080 interface: Instant neural control of a movement signal. *Nature* **416**:141–142.
1081 doi:10.1038/416141a
- 1082 Shanechi MM, Orsborn AL, Moorman HG, Gowda S, Dang S, Carmena JM. 2017. Rapid
1083 control and feedback rates enhance neuroprosthetic control. *Nat Commun* **8**:13825.
1084 doi:10.1038/ncomms13825
- 1085 Shenoy KV, Carmena JM. 2014. Combining Decoder Design and Neural Adaptation in Brain-
1086 Machine Interfaces. *Neuron* **84**:665–680. doi:10.1016/j.neuron.2014.08.038
- 1087 Shenoy KV, Sahani M, Churchland MM. 2013. Cortical control of arm movements: A dynamical
1088 systems perspective. *Annu Rev Neurosci* **36**:337–359. doi:10.1146/annurev-neuro-
1089 062111-150509
- 1090 Suminski AJ, Tkach DC, Hatsopoulos NG. 2009. Exploiting multiple sensory modalities in brain-
1091 machine interfaces. *Neural Netw* **22**:1224–1234. doi:10.1016/j.neunet.2009.05.006
- 1092 Sussillo D, Churchland MM, Kaufman MT, Shenoy KV. 2015. A neural network that finds a
1093 naturalistic solution for the production of muscle activity. *Nat Neurosci* **18**:1025–1033.
1094 doi:10.1038/nn.4042
- 1095 Sussillo D, Nuyujukian P, Fan JM, Kao JC, Stavisky SD, Ryu S, Shenoy K. 2012. A recurrent
1096 neural network for closed-loop intracortical brain–machine interface decoders. *J Neural*
1097 *Eng* **9**:026027. doi:10.1088/1741-2560/9/2/026027
- 1098 Sussillo D, Stavisky SD, Kao JC, Ryu SI, Shenoy KV. 2016. Making brain–machine interfaces
1099 robust to future neural variability. *Nat Commun* **7**:13749. doi:10.1038/ncomms13749
- 1100 Taylor DM, Tillery SIH, Schwartz AB. 2002. Direct Cortical Control of 3D Neuroprosthetic
1101 Devices. *Science* **296**:1829–1832. doi:10.1126/science.1070291
- 1102 Thacker NA, Bromiley PA. 2001. The Effects of a Square Root Transform on a Poisson
1103 Distributed Quantity. (Technical Report). Manchester, UK: Univ. of Manchester.
- 1104 Trautmann EM, Stavisky SD, Lahiri S, Ames KC, Kaufman MT, O’Shea DJ, Vyas S, Sun X, Ryu
1105 SI, Ganguli S, Shenoy KV. 2019. Accurate Estimation of Neural Population Dynamics
1106 without Spike Sorting. *Neuron* **0**. doi:10.1016/j.neuron.2019.05.003
- 1107 Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. 2008. Cortical control of a
1108 prosthetic arm for self-feeding. *Nature* **453**:1098–1101. doi:10.1038/nature06996
- 1109 Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, Kim J, Biggs SJ,
1110 Srinivasan MA, Nicolelis MAL. 2000. Real-time prediction of hand trajectory by
1111 ensembles of cortical neurons in primates. *Nature* **408**:361–365. doi:10.1038/35042582
- 1112 Willett FR, Avansino DT, Hochberg LR, Henderson JM, Shenoy KV. 2020. High-performance
1113 brain-to-text communication via imagined handwriting. *bioRxiv* 2020.07.01.183384.
1114 doi:10.1101/2020.07.01.183384
- 1115 Wodlinger B, Downey JE, Tyler-Kabara EC, Schwartz AB, Boninger ML, Collinger JL. 2015.
1116 Ten-dimensional anthropomorphic arm control in a human brain–machine interface:
1117 difficulties, solutions, and limitations. *J Neural Eng* **12**:016011. doi:10.1088/1741-
1118 2560/12/1/016011