Robust neural control of virtual locomotion enabled by a novel decoding strategy

Karen E Schroeder* [1,2], Sean M Perkins* [2,3], Qi Wang[3], & 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

Abstract

Brain-machine interfaces (BMIs) for reaching have enjoyed continued performance improvements. Yet there remains significant need for locomotor BMIs (e.g., for wheelchair control), which could potentially benefit a much larger patient population. Fewer studies have addressed this need, and the most effective approach remains undetermined. Here, we develop a locomotor BMI based on cortical activity as monkeys cycle a hand-held pedal to progress along a virtual track. Unlike most reach-based BMIs, we did not directly map neural states to commanded velocity or position. Instead, we leveraged features of the neural population response that were robust during rhythmic cycling. These included an overall shift in neural state when moving, and rotational trajectories with direction-specific paths. We used nonlinear means to infer kinematics from these features. Online BMI-control success rates approached those during manual control. Our results illustrate that different use-cases can require very different approaches to guiding a prosthetic via neural activity.

Introduction

Brain-machine interfaces (BMIs) interpret neural activity and provide control signals to external devices such as computers and prosthetic limbs. Intracortical BMIs for reach-like tasks have proved successful in primates and human clinical trials[1–8]. More widespread use appears imminent. Yet at the same time, there exist non-reach-like movements whose restoration is valuable to patients. For example, many patients could benefit from BMIs that control locomotion through an environment (e.g., movement of a wheelchair). Recent work has demonstrated that this is feasible[9,10]. While locomotor BMIs can be guided by reach-inspired decoding approaches, other viable strategies exist and remain unexplored. For example, it may be more natural or robust to decode locomotor commands from continuous rhythmic neural activity. The specific case of locomotor BMIs highlights a broader issue. A focus on reaching and grasping has produced state-of-the-art decode algorithms that may not generalize well to non-reaching applications, where the structure of neural activity may depart dramatically from that assumed when those decoders were designed. Locomotor BMIs provide a natural place to start exploring what other strategies might be pursued.

The early success of reach-based BMIs[11–16] was built upon a 'biomimetic' decode strategy that sought to invert the cortical encoding of kinematic variables such as hand velocity / direction. Multiple lines of evidence now argue against the hypothesis that motor cortex activity literally encodes kinematic variables[17–22]. Yet within the confines of a given task, robust correlations between neural activity and reach direction afford excellent BMI performance[11–16]. Indeed, as decoder sophistication has increased, improvements have derived largely from better estimating the neural state, resulting in improved decoding of intended reach velocity[1,23–26]. Indeed, this strategy is employed even when the hypothesis of literal kinematic encoding is rejected[23]. Thus, a decoding approach originally inspired by biomimetic considerations has endured for a practical reason: during reaching, aspects of the neural response covary usefully with the direction and magnitude of hand velocity.

Recent findings support this 'opportunistic' decoding strategy but indicate that it may need to be deployed differently for different movement classes. Based upon data collected while monkeys perform a cycling task, we argued that the dominant signals in motor cortex (those with the greatest influence on neural firing rates) exist to ensure noise-robust dynamics and do not encode any external quantity[18]. In contrast, neural signals that relate directly to outgoing commands (e.g., downstream muscle activity) appear relatively small. Because small signals cannot be reliably decoded on single trials, a reasonable strategy is to abandon attempts at biomimeticism, and simply leverage the most statistically robust relationships between neural responses and the parameters one wishes to control. Such relationships may be strongly movement-type dependent.

Here, we demonstrate the effectiveness of this strategy by developing a BMI system to decode virtual motion during cycling. The task, in which monkeys rotate a hand-held pedal to move along a virtual track, does not emulate natural locomotion. Instead, cycling provides a view of cortical activity during a rhythmic voluntary movement. The resulting rhythmic neural

responses represent a class of activity with which a patient might wish to control prosthetic locomotion. The cycling task also provides a usefully different view of neural activity outside the confines of a standard reaching task.

In agreement with recent results, we found that neural population activity during cycling differs from that during reaching in fundamental ways. It was not simply that activity was sustained rather than transient. More critically, 'directional' signals related to kinematics (hand velocity or position) were of small magnitude – i.e., they were only weakly reflected in neural firing rates. Signals related to muscle activity were of similarly small magnitude.

Thus, rather than directly map neural activity to decoded limb kinematics, we sought to leverage features of the neural response that had robust (but not necessarily linear) relationships with the variable we most wished to decode: self-motion through the environment. We identified three such features. The first was a translation of the neural state that rapidly and robustly indicated whether the monkey was moving. The second was an elliptical neural-trajectory while cycling. That trajectory did not reverse with cycling direction (as would a representation of hand velocity) but instead occupied direction-dependent dimensions. As a result, cycling direction / velocity could be estimated by comparing the angular momentum of the neural state between pairs of dimensions. This feature was robust during ongoing cycling but not at movement initiation. We thus leveraged a third feature: at movement initiation, neural activity corresponding to forward and backward cycling was briefly linearly separable.

A decoder that leveraged these features provided excellent online control of virtual locomotion. Success rates and acquisition times were very close to those achieved under manual control. Almost no training or adaptation time was needed; the low-latency and accuracy of the decoder were such that monkeys appeared to barely notice transitions from manual control to BMI control. These results demonstrate the feasibility of BMI locomotion based on rhythmic neural activity. More broadly, they establish a strategy for approaching BMI control in situations beyond traditional reach-based BMIs.

Results

Behavior

We trained two monkeys (G and E) to rotate a hand-held pedal to move through a virtual environment (**Fig. 1**). The magnitude of juice reward depended on their accuracy in stopping on those targets. All motion was along a linear track – no steering was necessary. Consistent with this, a single pedal was cycled with the right arm only. Our goal when decoding was to reconstruct the virtual motion produced by that single pedal.

On each trial, a target appeared in the distance. To acquire that target, monkeys produced virtual velocity in proportion to the rotational velocity of the pedal. The color of the environment (lush and green versus desert-like and tan) instructed cycling direction. When the environment was green (**Fig. 1a**, *left*) forward virtual motion was produced by cycling 'forward' (i.e., with the hand moving away from the body at the top of the cycle). When the environment

was tan (**Fig. 1a**, *right*) forward virtual motion was produced by cycling 'backward' (the hand moving toward the body at the top of the cycle). Cycling in the wrong direction produced motion away from the target. Trials were presented in blocks of forward or backward trials. Within each block, targets were separated by a randomized distance of 2, 4 or 7 cycles. Acquisition of a target was achieved by stopping and remaining stationary 'on top' of the virtual target for a specified time. Reward was then given and the next target appeared.

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

BMI control was introduced after monkeys were adept at performing the task under manual control. Task structure and the parameters for success were unchanged under BMI control, and no cue was given regarding the change from manual to BMI control. The switch to BMI control was made at the beginning of the session, after completion of a block of manual-control trials (25 forward and 25 backward 7-cycle trials). These manual-control trials were used to train the decoder. The switch was then made to BMI control for the remainder of the session. For monkey G, we occasionally included blocks of manual-control trials later in the session to allow comparison between BMI and manual performance. Monkey E performed fewer trials per session, and we thus used separate (interleaved) sessions to assess manual-control performance.

During both BMI control and manual control, the monkey's ipsilateral (non-cycling) arm was restrained. The contralateral (cycling) arm was never restrained. We intentionally did not dissuade the monkey from continuing to physically cycle during BMI control. Indeed, our goal was that the transition to BMI control would be sufficiently seamless to be unnoticed by the monkey, such that he would still believe that he was in manual control. An advantage of this strategy is that we are decoding neural activity when the subject attempts to actually move, as a patient presumably would. Had we insisted the arm remain stationary, monkeys would have needed to actively avoid patterns of neural activity that drive movement – something a patient would not have to do. Allowing the monkey to continue to move normally allowed us to extensively quantify the performance of our decoder by comparing decoded with intended (i.e., actual) movement. This is often not possible when using other designs. For example, in Rajangam et. al.[9], performance could only be assessed via indirect measures (such as time to target) because what the monkey was actually intending to do at each moment was unclear. We considered these advantages to outweigh a potential concern: a decoder could potentially 'cheat' by primarily leveraging activity due to proprioceptive feedback (which would not be present in a paralyzed patient). This is unlikely to be a large concern. Recordings were made from motor cortex, where robust neural responses precede movement onset. Furthermore, we

have documented that motor cortex population activity during cycling is quite different from that within the proprioceptive region of primary somatosensory cortex[18]. Thus, while proprioceptive activity is certainly present in motor cortex[27–30] (especially during perturbations[31]) the dominant features of M1 activity, described below, are unlikely to be primarily proprioceptive.

Neural activity and decoding strategy

We recorded motor cortical activity using 96-channel Utah arrays. For monkey G, one array was implanted in primary motor cortex (M1) and a second in dorsal premotor cortex (PMd). For monkey E, a single array was implanted in M1. For each channel we recorded times when the voltage passed a threshold. Threshold crossings typically reflected the spikes of a handful of neurons (a neural 'unit'). Spikes from individual neurons could be clearly seen on many channels, but no attempt was made to spike-sort, as the benefit of doing so is typically modest when controlling a prosthetic device[32]. Unit activity was strongly modulated during cycling (**Fig. 1d**) and the phase, magnitude, and temporal pattern of activity depended on whether cycling was forward (*green* traces) or backward (*red* traces). A key question is how these unit-level features translate into population-level features that might be leveraged to estimate intended motion through the virtual environment.

In traditional decoding approaches (**Fig. 2a**, *top*) neural activity is hypothesized (usefully if not literally) to encode kinematic signals, which can be decoded by inverting the encoding scheme. Although nonlinear methods (including variations of Kalman filtering) are often used to estimate the neural state, the final conversion to a kinematic command is typically linear or roughly so. To explore kinematic encoding in the present task, we used linear regression to identify neural dimensions where activity correlated well with kinematics (including hand velocity and position). Regression was performed using single trials. Use of single trials provides a large quantity of training data and is implicitly regularizing: regression must find signals that are robust in the face of single-trial spiking variability. We computed the variance captured by the dimensions found via regression (variance captured was computed using trial-averaged data so that results were not diluted by noise). Surprisingly, the neural dimensions that best captured kinematic signals captured little population response variance (**Fig. 2b**, green bars). This was also true of neural dimensions that captured muscle activity (**Fig. 2b**, yellow bar). This was initially surprising, as single-neuron responses were robustly sinusoidally modulated, as were many kinematic variables. Yet those sinusoidal response features were often superimposed upon other response features (e.g., overall shifts in rate when moving versus not moving). Sinusoidal features also did not display phase relationships, across forward and backward cycling, that were consistent with kinematic encoding[18].

Low-variance signals are a poor candidate for decoding intended action, as they are likely to be non-robust with respect to multiple challenges. Some of these challenges (e.g., spiking noise) can be anticipated and estimated in advance, but others cannot. This includes recording instabilities, changes in strategy or behavior with time, and (outside the laboratory setting) external sources of noise or variability. Given these challenges, it is worth stressing that there

are two practical reasons why reach-focused BMIs have leveraged signals that correlate with hand velocity. First, such signals are high-variance during reaching – so much so that M1 responses have often been summarized in terms of a preferred direction[33,34]. Second, movement direction / velocity are the variables one wishes to decode during reaching. In the present case neither motivation holds. Signals related to hand kinematics are low-variance, and we most wish to decode self-motion through the virtual environment. This suggests an alternative strategy (**Fig. 2a**, *bottom*): identifying neural response features that are both robust (high-variance) and relate reliably to the presence and direction of self-motion.

To pursue this strategy, we considered three sets of high-variance dimensions. The first set included four 'rotational dimensions' (two each for forward and backward cycling) which captured elliptical trajectories present during steady-state cycling[18]. The second set included a single 'moving-sensitive' dimension, in which the neural state distinguished whether the monkey was stopped or moving regardless of movement direction[35]. The third set was a triplet of 'initial-direction' dimensions. In these dimensions, cycling direction could be transiently but readily distinguished in the moments after cycling began.

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

Performance

Monkeys performed the task very well under closed-loop BMI control (**Fig. 3** and **Supp. Movie 1**). Monkeys continued to cycle as normal, presumably not realizing that the pedal had been disconnected from the control system. The illusion that the pedal still controlled the task was supported by a high similarity between decoded virtual velocity and intended virtual velocity (i.e., what would have been produced by the pedal were it still controlling the task). The cross-correlation between these peaked at 0.93 ± .02 and 0.81 ± .03 (monkey G and E, mean ± SD) at a short lag: 76 ± 4 ms and 102 ± 7 ms (**Fig. 3a**). That illusion was also aided by a low rate of false starts; it was exceedingly rare for decoded motion to be non-zero when the monkey was attempting to remain stationary on top of a target. False starts occurred on 0.29% and 0.09% of trials (monkeys G and E), yielding an average of 1.9 and 0.12 occurrences per day. This is notable because combatting unintended movement has been a persistent challenge for BMI decoding, leading some studies to incorporate discrete moving and stopped states into otherwise continuous decoders[2,36,37]. These features of our decode – high correlation with intended movement, low latency, and few false starts – led to near-normal performance under BMI control (**Fig. 3b,c**). Success rate under BMI control (**Fig. 3d**, *magenta symbols*) was almost as

high as that under manual control (*open symbols*), and the time to move from target to target was only slightly longer under BMI control (**Fig. 3e**).

The only respect in which BMI control suffered noticeably was accuracy in stopping on the middle of the target. Under manual control, monkeys stopped very close to the middle of the target (**Fig. 3f,** *gray histogram*), which always corresponded to the 'pedal-straight-down' position. Stopping was considerably less accurate under BMI control (*magenta histogram*). Errors were due to two sources, only one of which relates to sub-optimal decoding. The first source is that, during rapid cycling, even small errors in decoded stopping time become relevant. Monkeys cycled at a virtual speed equivalent to what would be produced by physical cycling at ~2 Hz. At this speed, a 100 ms error in the decoded stopping time produces an error equivalent to that produced by 0.2 cycles of physical motion. Our decode had ~ ±100 ms of variability in the decoded stopping time relative to the actual stopping time (standard deviations of 133 and 99 ms for monkey G and E, averaged across sessions). Trial-to-trial variability, unlike a bias, cannot be compensated for and thus necessarily limited performance. Increased stopping error in BMI-control trials was also due to an incidental advantage of manual control: the target center was always aligned with the pedal-straight-down position, a fact which monkeys leveraged to stop very accurately in that position. This strategy was not available during BMI control because the correct time to stop rarely aligned perfectly with the pedal-straight-down position (this occurred only if decoded and intended virtual velocity matched perfectly when averaged across the cycling bout).

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

State machine

The performance documented above was achieved by a state-dependent decode (**Fig. 4**). State transitions were governed by activity in the moving-sensitive dimension, which was translated into a probability of moving, p_{move}. If p_{move} was low, the STOP state was active and decoded virtual velocity was zero. When p_{move} became high, the INIT state was entered but decoded velocity remained zero. After 175 ms, the EARLY state was entered and velocity was decoded using the initial-direction dimensions. After an additional 200 ms, the STEADY state was entered and virtual velocity depended only on the neural state in the rotational dimensions.

Regardless of present state, $p_{move} < 0.1$ always produced a transition back to STOP. Decoded velocity was filtered to smooth fluctuations during STEADY. Below we describe how virtual velocity was estimated while in STEADY, how p_{move} was derived, and how we decoded the early direction of movement.

Direction of steady-state movement inferred from rotational structure

The dominant feature of the neural response during steady-state cycling was a repeating elliptical trajectory[18]. Our decoder leveraged the fact that forward-cycling and backward-cycling trajectories occurred in non-identical dimensions. We employed an optimization procedure to find a two-dimensional 'forward plane' that maximized the size of the forward trajectory relative to the backward trajectory. We similarly found an analogous 'backward plane'. Forward and backward planes were identified based on trial-averaged responses from the 50 trials of training data collected under manual control (**Fig. 5a**). With the aid of filtering (*Methods*), these planes continued to capture rotational features on individual trials (**Fig. 5b**). Although forward and backward trajectories were not orthogonal to one another (subspace angles between planes found for forward and backward trajectories separately were <90°) the above procedure was still able to find planes where strongly elliptical trajectories were present of one cycling direction but not the other.

Note that the rotating neural state did not have a linear relationship with the quantity we wished to decode: the direction of intended virtual motion. For example, forward motion was associated with neural states to both the left and right of zero in the forward plane. Furthermore, this relationship was not static. Whether a neural state was consistent with forward motion depended on the conjunction of its location and its derivative. For example, a state to the right of zero in the forward plane indicates forward motion only if the derivative points upwards. We thus did not attempt to directly map the location of the neural state onto decoded velocity. Instead, we assessed the angular momentum (the cross product of the state vector with its derivative) in each plane. When moving backward (first three cycling bouts in **Fig. 5c**) angular momentum was sizeable in the backward plane (*dark blue*) but not the forward plane (*bright blue*). The opposite was true when moving forward (subsequent three bouts).

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

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

Inferring the probability of moving

Decoders that directly translate neural state to cursor velocity have historically had difficulty remaining stationary when there is no intended movement. The ability to do so is of even more importance for a locomotor prosthetic. To meet this challenge, we adopted the strategy of a state machine with distinct stopped and moving states[36–38]. Transitions between these states were governed by a probability of moving, p_{move}, derived from the neural state in the moving-sensitive dimension.

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

Using the training data, we fit a Gaussian to the distributions of neural states when stopped (**Fig. 6c**, *orange*), and another to that when moving (*blue*). During BMI control, p_{move} was estimated by considering the likelihood of observing the current state (and all previous states) given that the monkey is currently moving or stopped, combined with the baseline (prior) probability of moving. During BMI control, p_{move} (**Fig. 6d**, *blue*) was near unity during intended movement (i.e., when the monkey was actually cycling, *black*) and near zero otherwise.

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

327 ***Inferring initial movement direction***

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

342 During BMI control, upon transition from INIT to EARLY, we computed the likelihood of the
343 neural state under each distribution. A simple winner-take-all computation determined the
344 direction of virtual velocity during the EARLY state. The inference of movement direction
345 during EARLY was correct on 94% and 82% of trials (monkeys G and E). After 200 ms, the
346 STEADY state was entered and virtual velocity was controlled thereafter by activity in the
347 rotational dimensions. **Figure 7d** illustrates the moments where the above strategy was used to
348 decode motion direction for three backward and three forward trials (under BMI control, with
349 physical pedal velocity shown for reference). These moments were brief, and had a very modest
350 effect on the overall time to reach the target. However, we still employed this strategy because
351 our goal was to build a BMI decode that closely tracked intended movement and felt
352 responsive.

353 ***Speed control***

354 The excellent performance of the decoder was aided by the relative simplicity of behavior: when
355 monkeys moved, they did so at a stereotyped speed. This allowed us to concentrate on building
356 a decode algorithm that decoded intended direction with accurate timing, and remained
357 stationary if movement was not intended. However, that decode provided only limited control
358 of movement speed. An obvious extension is to allow finer-grained speed control. This would
359 presumably be desired by users of a locomotor prosthetic. Furthermore, speed control provides
360 one possible way of steering: e.g., by decoding the relative intensity of intended movement on
361 the two sides of the body. While we do not attempt that here, we still considered it important to
362 determine whether the neural features we identified could support speed control.

363 That assessment required a task where speed control is necessary for success. We thus trained
364 one monkey to track various speed profiles as he progressed through the virtual environment.
365 Two floating targets were rendered in the foreground as the monkey cycled and the distance
366 between them reflected his error in matching the instructed speed. Obtaining juice required

aligning the two floating targets while progressing towards a final target, on which he stopped to obtain additional juice reward. The task was divided into trials, each of which required moving a distance equivalent to twenty cycles under manual control. We used eight total trial-types, four each for forward and backward cycling. Two of these employed a constant target speed (equivalent to 1 or 2 Hz cycling) and two involved a ramping speed (from 1 Hz to 2 Hz or vice versa). As above, the decoder was trained based on a small number of manual-control trials performed at the beginning of each session. Blocks of manual-control trials were also included for comparisons between manual and BMI-based performance.

Our decode strategy was largely preserved from that described above. However, we used a modified state machine (**Supp. Fig. 2**) and a slightly different algorithm for transforming rotations of the neural state into decoded virtual velocity. Direction was determined based on which distribution (forward or backward) produced the higher likelihood of observing the measured angular momentums (as in **Fig. 5d**). Once that choice was made, speed was determined by the angular velocity of the neural state in that plane. Thus, faster rotational trajectories led to faster decoded virtual velocity. We chose a scaling factor so that a given neural angular velocity produced the speed that would have been produced by physical cycling at that angular velocity. Neural angular velocity was exponentially filtered with a time constant of 500 ms. The filter memory was erased on entry into a movement state (EARLY or STEADY) from a stopped state (INIT or EXIT) to allow brisk movement onsets (see *Methods).*

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

Under BMI control, monkeys were able to match decoded virtual speed to the instructed speed. They did so both across trials with different constant speeds, and within trials where speed modulated with time (**Fig. 8a,b**). Monkeys were able to match instructed speed nearly as accurately under BMI control as under manual control (**Fig. 8b** shows a typical example trials). This was true judged both by time within the rewarded speed window (**Fig. 8c**) and by the error between virtual and instructed velocity (**Fig. 8d**).

Discussion

We have argued that the largest signals in motor cortex are not 'representational' – they do not encode variables but are instead essential for noise-robust dynamics. Those dynamics produce outgoing commands that *are* representational (they covary with the variables they control) but are low-variance. This perspective argues that decoders should not attempt to invert encoding unless a great many neurons can be recorded. Instead, decoders should opportunistically

leverage whichever high-variance response features have a robust relationship with the variables one wishes to decode. In retrospect, traditional reach-based prosthetics can be seen as taking exactly this approach; during reaching there exist high-variance neural signals that correlate (linearly) with reach direction. Similarly, decoding of muscle activity for prosthetic control[6] likely leverages signals that coincidentally but usefully correlate with muscle force during the task of interest.

During cycling, the neural signals that correlate strongly with hand position and velocity are low variance. Yet there exist different high-variance response features that possess robust relationships with intended movement. Those relationships may be nonlinear, but can nevertheless be decoded using straightforward means (e.g., computing angular momentum in state space). Doing so yielded BMI control that was sufficiently natural that monkeys appeared not to notice that the task was no longer under manual control. By most measures (success rate, time to target) performance under BMI control was remarkably close to that under manual control. The main limitation of BMI control was stopping accuracy. Although our algorithm typically detected stopping within a few tens of milliseconds, these small discrepancies could lead to the target being over or undershot by a noticeable amount. Perhaps most importantly, our BMI decode almost never produced movement when it was not intended. A truly zero velocity was decoded whenever the monkey was intending to remain stopped on the target. We consider this a particularly important attribute of any locomotor-BMI algorithm, due to the potentially large consequences of unintended movement of the whole body.

The present approach relates to recent studies that modeled neural dynamics to improve online[23] or offline[24,39] decodes of movement kinematics. A key insight of those studies is that signals that do not correlate directly with kinematics can be used to infer those that do. For example, when a given variable is linearly decoded, it depends only upon the neural state in one dimension: the dimension defined by the regression weights. Nevertheless, inferring the neural state in that dimension may benefit from a dynamical model that spans multiple dimensions. Much like the present approach, this allows the decode to leverage features that are robust, even if they do not directly correlate with the kinematic parameters of interest. The present approach extends this idea to situations where there may be no high-variance dimensions that can be linearly decoded, and/or where the most prominent features are not well-described by linear dynamics.

We had three motivations for developing a BMI algorithm for virtual locomotion. First, our recently developed cycling task naturally lends itself to this application. Second, BMI-controlled locomotion is likely to be desired by a large patient population (potentially much larger than the population that desires BMI-controlled cursors or robot arms). Third, prior work has demonstrated that BMI control of locomotion is viable[9,10]. In particular, Rajangam et al.[9] demonstrated BMI control of a physical wheelchair based on neural activity recorded from within motor and somatosensory cortex. Our work supports their conclusion that BMI-controlled locomotion is possible, and demonstrates the feasibility of an alternative decode strategy. Rajangam et al. employed a traditional decode strategy: linear filters transformed neural activity into the key variables: translational and angular velocity. That strategy allowed

448 monkeys to navigate ~2 meters to a target (which had to be approached with an accuracy of +/-
449 ~0.2 meters, or 10% of the distance traveled) in an average of 27-49 seconds (depending on the
450 monkey and degree of practice). In our task, monkeys traveled up to seven cycles and had to
451 stop with an accuracy of +/- 0.5 cycles (7% of the distance traveled). Under BMI control,
452 monkeys traveled those seven cycles in 4.3 and 3.7 seconds on average. While this is a
453 considerable reduction, we stress that movement durations are not directly comparable
454 between our task and theirs. Success requirements differed in multiple ways. For example,
455 Rajangam et al. did not require that monkeys stop on a target location (which made the task
456 easier) but did require that they turn *en route* (which adds considerable challenge). Still, an
457 approximately ten-fold improvement in time-to-target argues that non-traditional decode
458 strategies can be effective and should be explored further.

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

471 For convenience, monkeys were trained to control the pedal with their forelimb (allowing them
472 to be seated in a traditional primate chair) and we thus recorded from the forelimb region of
473 motor cortex. Recordings during natural locomotion in monkeys reveal broadly similar signals
474 in both the forelimb[40] and hindlimb[41] regions of motor cortex. These signals are dominated, as in
475 our task, by rotations of the neural state during ongoing locomotion. It is thus likely that
476 prosthetic locomotion could be driven by signals derived from either region.

477 Like many proof-of-concept prosthetic systems developed in primates[1,16,42–44], decoder training
478 depended on observations of neural activity under manual control. Whether this approach
479 translates depends upon the assumption that useful patterns of neural activity will emerge
480 when a paralyzed patient tries to move but can't. A number of existing studies indicate that
481 motor cortex is active, in reasonably normal ways, when paralyzed patients attempt to move.
482 Yet it remains possible that the dominant neural features will be subtly or even dramatically
483 different in paralyzed patients, which would require adapting the decode strategy. We stress
484 that our larger point is not the specifics of the decoder used here, but rather the broader
485 approach of identifying and leveraging population-level response features that relate robustly
486 to the variables one wishes to control. Such features are presumably present in the motor
487 cortices of paralyzed patients (and if they are not, then motor cortex would not be a suitable
488 area for guiding a prosthetic regardless of the decode strategy).

The present results indicate that a nonlinear, yet relatively simple, decode strategy can afford excellent one-dimensional control of locomotion. Although other approaches remain possible[45], our findings support the idea that cortical control of prosthetic locomotion is viable and should be explored further. More broadly, the present results argue that many of the decode strategies that proved effective for reach-based prostheses are unlikely to generalize across tasks. An alternative approach is to identify, for each task, the dominant features and determine how they might be usefully translated into decoded movement. While this approach abandons the elegant idea of inverting a literal encoding of kinematics, it opens up possibilities for improved prosthetic control across a variety of contexts.

Methods

Subjects and Tasks

All procedures were approved by the Columbia University Institutional Animal Care and Use Committee. Subjects G and E were two adult male macaque monkeys (*Macaca mulatta*). Monkeys sat in a primate chair facing an LCD monitor (refresh rate: 144 Hz) that displayed a virtual environment, generated by the Unity engine (Unity Technologies, San Francisco, CA). The head was restrained via a titanium surgical implant. While the monkey's left arm was comfortably restrained to keep it out of the way, the right arm grasped a hand pedal to perform the task. Tape was used to ensure consistent placement of the hand on the pedal. The pedal connected via a shaft to a motor (Applied Motion Products, Watsonville, CA), which measured the position of the pedal with a precision of 1/10,000 of the cycle. The motor was also used to apply forces to the pedal, to counteract gravity, and to give it virtual mass and viscosity, making the cycling motion feel more natural.

Manual-control sessions for the primary cycling task required that the monkey turn the hand pedal in the instructed direction to move through the virtual environment, and stop on top of a lighted target to collect juice reward. The color of the landscape indicated whether cycling must be 'forward' (green landscape, the hand moved away from the body at the top of the cycle) or 'backward' (tan landscape, the hand moved toward the body at the top of the cycle). There were 6 total conditions, defined by cycling direction (forward or backward) and target distance (2, 4, or 7 cycles). Distance conditions were randomized within same-direction blocks (3 trials of each distance per block), and directional blocks were randomized over the course of each experiment. A trial would be failed without reward if the monkey began moving before the reaction time limit of 170 ms, if the monkey moved past the target without stopping, or if the monkey stopped on the final target but then moved while waiting for the reward. A variable inter-trial interval ensured that the monkey could not anticipate the appearance of the target. Monkeys performed until they received enough liquid reward that they chose to desist. As their motivation waned, they would at times take short breaks. For decoder performance analysis of both manual-control and BMI-control sessions, we discarded any trials that monkeys failed to initiate and did not count them as 'failed'. These trials occurred 2 ± 2 times per session (Monkey G, maximum 10) and 3 ± 3 times per session (Monkey E, maximum 11).

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

(average of 231 trials/session over 5 sessions) were interleaved with BMI-control sessions on different days.

An additional three manual-control sessions (189 trials, 407 trials, 394 trials), collected in monkey G, contained simultaneous EMG recordings (5-7 muscles per session, total of 19 muscles). The specific muscles recorded varied across sessions, but included: all three heads of the *deltoid*, the lateral and long heads of *triceps brachii*, the long head of *biceps brachii* (along with additional *biceps brachii* recordings of uncertain identity), *trapezius*, and *latissimus dorsi*. These muscles were selected due to their clear activations during the cycling task. These manual-control sessions included two extra conditions (1-cycle forward and 1-cycle backward) and were only analyzed in the variance captured analysis (**Fig. 2b,c**).

In the speed-tracking task, the monkey was required to match his avatar's speed to an instructed speed. A target was rendered a short distance ahead of the monkey's avatar and remained there (at a fixed offset to the avatar) throughout the trial. A secondary target drifted ahead or behind the fixed target to indicate the error between the avatar's speed and the instructed speed. Thus, successful speed-tracking in this task consisted of aligning these targets on top of one another. This design ensured that the monkey would have real-time feedback about his error throughout speed-tracking and disallow accrual of error throughout the trial that would alter his intention to match the instructed speed. The same colored-background directional cue was given as in the standard task, and speed target profiles were randomized within directional blocks. There were four conditions per direction: constant speed targets of 1 or 2 cycles/s, and ramps that increased or decreased between 1 to 2 cycles/s. For ramp conditions, the trial started at a constant initial speed (first 3 seconds of the trial), followed by ramping (next 8 seconds of the trial), followed by the constant final speed for the remainder of the trial. Each trial was 20 cycles total and began with a brief period during which the monkey got up to speed, followed by a rewarded speed-tracking epoch that lasted until the monkey was within one cycle of the final target, on which he stopped to receive additional juice reward. During the speed-tracking epoch, the monkey received frequent dispensation of juice droplets while matching instructed speed within a tolerance of 0.2 Hz. This task was only performed by monkey G, who completed an average of 166 trials/session over 2 sessions in manual control and an average of 116 trials/session over 3 sessions in BMI control.

The speed-tracking task was very difficult to perform in manual control due to the monkey's high intra-cycle variability in cycling speed. Given that a patient using a locomotive prosthetic would likely prefer to modulate their speed slowly, we chose to smooth the monkey's cycling speed with an exponential moving average ($\tau = 1$ s) prior to computing the speed error he received visual feedback on. This was effective; however, the large time constant made it difficult for the monkey to quickly get up to the instructed speed in the early portion of each trial. This was addressed by creating a speed window (0.07 Hz to 0.2 Hz) in which speed through the virtual environment was unfiltered. If speed dropped below the lower threshold or exceeded the upper threshold, the exponential moving average was set to the current speed and was then smoothed over time from that point onward. This allowed the monkey to quickly match the instructed speed and then slowly vary his mean speed to match the instructed speed

profile throughout each trial. This smoothing was not directly applied to decoded velocity in BMI control because similar filtering was performed within the decoder. In manual-control trials, we also required that his smoothed velocity could only deviate so far from the instructed velocity (± 0.5 Hz), or else he would fail the trial. This was done to discourage the monkey from ignoring the speed command and simply pedaling as rapidly as possible to the final target. BMI control for this task was implemented with all the same task parameters as manual control, with the exception of the aforementioned failure mode. We chose not to implement this failure-mode in BMI control because we wanted to quantify how well the monkey could match speed via the decode and did not want to effectively down-select to speed-matched trials by terminating trials with large errors in speed. Accordingly, we did not compare success rates between BMI control and manual control for this task.

Surgery and Neural/Muscle Recordings

Neural activity was recorded using chronic 96-channel Utah arrays (Blackrock Microsystems, Salt Lake City, UT) that were implanted using standard surgical techniques. Arrays were placed in left hemisphere arm area of primary motor cortex (M1) of both monkeys based on MRI scans and on anatomical landmarks. In monkey G, a second array was placed in dorsal premotor cortex (PMd), just anterior to the first array. Data were collected 1-8 months and 3-4 months after surgical implantation for monkeys G and E, respectively.

Electrode voltages were analog filtered (band-pass 0.3 Hz – 7.5 kHz) and digitized at 30 kHz using Digital Headstages, Digital Hubs, and Cerebus Neural Signal Processors from Blackrock Microsystems. Digitized voltages were then high-pass filtered (250 Hz) and thresholded at a value between -4.5 and -3 times the RMS voltage on each channel, depending on the array quality on a given day, to detect unit spike times.

Intra-muscular EMG recordings were made using pairs of hook-wire electrodes inserted with 30 mm x 27 gauge needles (Natus Neurology, Middleton, WI). The raw voltages were amplified and analog filtered (band-pass 10 Hz – 10 kHz) with ISO-DAM 8A modules (World Precision Instruments, Sarasota, FL), and digitized at 30 kHz with the Cerebus Neural Signal Processors. EMG was then digitally band-pass filtered (50 Hz – 5 kHz) prior to saving.

A real-time target computer (Speedgoat, Bern, CH) running Simulink Real-Time environment (MathWorks, Natick, MA) processed behavioral and neural data and controlled the decoder output in online experiments with millisecond precision. It also streamed variables of interest to another computer that saved these variables for offline analysis.

Preprocessing

Spike trains were causally converted to firing rates by convolving each spike with a beta kernel. The beta kernel was defined by temporally scaling a beta distribution (shape parameters: $\alpha = 3$ and $\beta = 5$) to be defined over the interval [0, 275] ms and normalizing the kernel such that the firing rates would be in units of spikes/second. The firing rates were then mean-centered and soft-normalized (as discussed in prior work, see[18]). Offline, EMG recordings were rectified, low-pass filtered by convolving with a Gaussian (standard deviation: 25 ms), downsampled to 1

kHz, and then fully normalized such that the maximum value achieved on each EMG channel was 1.

Trial Alignment

Aligning trials within conditions was necessary for creating trial averages and computing certain decoder parameters (e.g. Kalman filter measurement noise). Due to the nature of the task, cycling trials could be quite long (up to 20 cycles in the speed-tracking task) and of variable length, rendering the approach of aligning to movement onset insufficient for preserving the temporal structure of the data through alignment. Indeed, stretching/compressing the data in time was necessary to align the data across trials. This problem was solved in two ways. Method A is a simplified alignment method that was used to align trials within conditions during the parameter fitting procedure for online BMI control. This method only aligns the movement data. Method B is a more sophisticated alignment procedure that was utilized for the variance captured analysis and visualization of the neural and kinematic data (**Figs. 1,2,5-7**). This method aligns the entire trial, including pre- and post-movement data. For visualization, conditions with the same target distance (e.g., 7 cycles), but different directions, were also aligned to the same time base. Critically, any data processing that relied on temporal structure was completed in the original, unstretched time base prior to alignment.

Method A*:* The world position for each trial resembles a ramp between movement onset and offset (**Fig. 1a**). First, we identify the portion of each trial starting ¼ cycle into the movement and ending ¼ cycle before the end of the movement. We fit a line to the world position in this period and then extend that line until it intercepts the starting and ending positions. The data between these two intercepts is considered the movement data for each trial and is extracted. This movement data is then uniformly stretched in time to match the average trial length for each trial's associated condition. This approach compresses slower than average movements and stretches faster than average movements within a condition, such that they can be averaged while still preserving many of the cycle-specific features of the data.

Method B: This method consists of a mild, non-uniform stretching of time in order to match each trial to a condition-specific template. For complete details, see Russo et al. 2018.[18]

Variance Captured Analysis

For the variance captured analysis (**Fig. 2**), successful manual-control trials were analyzed from three sessions containing simultaneous EMG and neural recordings. The feature dimensions corresponding to our decoder (**Fig. 2c**) were found using the fitting procedure described for our decode algorithm (see subsequent sections). The traditional regressor dimensions (**Fig. 2b**) were found by linearly regressing single-trial neural activity against kinematic variables or EMG activity one at a time. Although this linear regression allowed for a constant offset, the regressor dimensions were defined by the weights of the regression excluding the offset term to ensure the dimension was a linear projection of the data. Each dimension was normalized to unit length after fitting. Additionally, we computed subspaces to evaluate variance captured by

combining subsets of dimensions and orthonormalizing the space spanned by those dimensions.

We then proceeded to generate trial-averaged activity for each condition using trial-alignment (Method B) and concatenated them into a matrix $\overline{Y} \in \mathbb{R}^{N \times \sum_c T_c}$ where T_c is the number of time points in the c-th condition. Given that $\overline{Y}$ is already mean-centered, percent variance captured for a dimension (of unit length) with weights organized in a column vector w can be computed as

$$\%\, Variance\ Captured = \frac{w^{\top} \overline{Y}\, \overline{Y}^{\top} w}{\mathrm{tr}\left(\overline{Y}\, \overline{Y}^{\top}\right)} \times 100$$

Variance captured for a subspace is simply the sum of the variance captured for each of the component orthonormal dimensions that span the space.

Decoder State Machine

During the primary experiments, BMI motion at each millisecond was determined by a state machine with four states: STOP, INIT, EARLY, and STEADY (**Fig. 4**), corresponding to the different stages of a typical trial. Decoder output in each state and the transitions between states were governed by the features described in the main text, as well as some additional timing controls. Details on each of these features are given in the subsequent sections.

An estimate of decoded velocity through the virtual environment, v_{dec}, was computed at each millisecond in a state-dependent manner. During the EARLY and STEADY states, this estimate was smoothed with a 500 ms trailing average. However, the span of the trailing average was shortened during the initial portion of a decoded movement to ensure it never averaged over values of v_{dec} generated in the STOP or INIT states (where $v_{dec} = 0$), i.e., the history of the trailing average grew in length from 0 to 500 ms over the first 500 ms of decoded movement. This approach to smoothing throughout the trial ensured brisk movement onsets and offsets and a smoothly varying estimate of decoded velocity during movement. The post-smoothing estimate of velocity was integrated every millisecond to create an estimate of position in the virtual environment.

During the speed-tracking experiment, a different state machine was used, which included a fifth state: EXIT (**Supp. Fig. 1**). The additional state provided some flexibility to reenter the EARLY or STEADY states without having to pass through the INIT state in circumstances where the decoder only briefly believed that movement had ceased. This allowed the monkey to continue moving with minimal interruption in the event of a momentarily incorrect stop decode. This state machine also introduced new criteria for certain state transitions, e.g., angular momentums were used in transitioning to the EXIT state and the transition from INIT to EARLY required that the estimate of initial direction meet some threshold confidence. More details on these criteria are provided in subsequent sections. No smoothing was performed directly on v_{dec} prior to integration into position. However, speed was smoothed with an

692 exponential moving average ($\tau = 500$ ms). On entry or reentry into EARLY or STEADY from
693 INIT or EXIT, the output of this exponential moving average was reset to an initial value of 1.5
694 Hz, which was the average starting speed across conditions.

695 *Moving-sensitive Dimension*

696 To control the decoder's movement onset and offset behavior, we used activity in one
697 dimension that captured an overall translation of the population activity that occurs and then
698 reverses at movement onset and offset, respectively. Linear discriminant analysis was used to
699 find a dimension that best separated labeled 'moving' (speed > 1 Hz) from 'stopped' (speed <
700 .05 Hz) training data. We computed binned spikes (10 ms for monkey G, 20 ms for monkey E)
701 and applied a square-root transform to these counts as this has been shown to improve the
702 Gaussian fit for Poisson data with small counts[46]. After projecting these square-rooted counts
703 onto the learned dimension, we fit Gaussian distributions for 'moving' and 'stopped' data
704 separately. (**see Fig. 6c**). An HMM was used to track the probability of being in a 'moving' or
705 'stopped' state[37]. The probability (p_{move}) of being in the 'moving' state, given the entire
706 sequence of current and previously observed square-rooted spike counts, was computed
707 efficiently with a recursive algorithm that uses the state transition matrix

708 $$\Phi = \begin{bmatrix} p_{move|move} & p_{move|stop} \\ p_{stop|move} & p_{stop|stop} \end{bmatrix}$$

709 and knowledge of the Gaussian distributions learned for each state. Φ encodes prior
710 assumptions about the probability of transitioning from one state to the next at any given bin.
711 We used a benchmark set of manual-control training data from each monkey to determine a
712 reasonable value for Φ, which was then used in all experiments. For monkey G, we set
713 $p_{move|stop} = .0001$ and $p_{stop|move} = .002$; for monkey E, we set $p_{move|stop} = .0002$ and
714 $p_{stop|move} = .004$. The value p_{move} was used throughout the decoder state machine to control
715 transitions between various states, effectively dictating the movement onset and offset behavior
716 of the decoder (**Fig. 6d**).

717 In the speed-tracking experiment, we used a different set of parameters, largely due to changes
718 in recording quality in the intervening timing between data collection for the primary
719 experiment and data collection for the speed-tracking experiment. The bin size was increased to
720 100 ms and the following state transition values were used: $p_{move|stop} = .0005$ and $p_{stop|move} =$
721 $.0005$. In addition, we observed that the square-root transform seemed to be having a negligible
722 impact on the quality of the decoder at these parameter values, so we removed it for this task.

723 *Rotational Dimensions*

724 To identify dimensions that captured the rotational dynamics in the data, we first computed
725 trial-averaged movement-period neural activity for the forward and backward 7-cycle
726 conditions in the training set (25 forward trials, 25 backward trials). We computed these neural
727 responses by first preprocessing the spikes to compute mean-centered and normalized firing
728 rates (as described in 'Preprocessing' section), then high-pass filtering (1 Hz) these firing rates

(2nd order Butterworth, cutoff frequency: 1 Hz) in order to remove slowly varying signals in the data (i.e. condition-invariant translation) that could partially align with the rotational activity we were seeking to isolate. We then used trial-alignment (Method A) to generate the $N \times T_c$ matrices $\overline{Y}_f$ and $\overline{Y}_b$, which contained the trial-averaged movement-period responses for the two conditions.

Having identified that rotational activity was observed for forward and backward trials in the top principal components of the data, we sought a projection of the high-pass filtered firing rates that would maximally capture these rotations while segregating forward and backward data into different planes in order to more readily decode pedaling direction on single trials. Whereas the standard PCA cost function finds dimensions that maximize variance captured, we opted instead for a cost function that would maximize the difference in variance captured between the two conditions:

$$J(W) = \mathrm{tr}(W^\top \Sigma_f W) - \mathrm{tr}(W^\top \Sigma_b W)$$

where $\Sigma_f = \mathrm{cov}(\overline{Y}_f)$, $\Sigma_b = \mathrm{cov}(\overline{Y}_b)$, W is constrained to be orthonormal, and the $\mathrm{cov}(\cdot)$ operator computes a covariance matrix by treating rows as dimensions and columns as observations. Note that this cost function will be maximized when the projection of the data captures a great deal of variance for forward trials and very little variance for backward trials. Conversely, this cost function will be minimized when the projection favors large variances for backward trials and small variances for forward trials. We thus chose to define our forward rotational plane by the 2D matrix W_f that maximizes $J(W)$ and our backward rotational plane by the 2D matrix W_b that minimizes $J(W)$. An iterative optimization procedure was used to find W_f and W_b; full details of this in[47].

Projecting single-trial, high-pass filtered firing rates onto these planes yielded large rotational activity that differed considerably between forward and backward conditions; however, single-trial noise still contributed to jitter in the observed neural trajectories. We addressed this by denoising the neural state in these planes with a linear Kalman filter of the form

$$x_t = Ax_{t-1} + q_t$$
$$y_t = Cx_t + r_t$$

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

$$A = \overline{X}_2 \overline{X}_1^\top \left(\overline{X}_1 X_1^\top\right)^{-1}$$

$$C = \begin{bmatrix} W_f^\top \\ W_b^\top \end{bmatrix}^\dagger$$

$$Q = \mathrm{cov}(\overline{X}_2 - A\overline{X}_1)$$

$$R = \mathrm{cov}(Y - C\overline{X})$$

where

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

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

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

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

with Y_i denoting the neural activity (high-pass filtered firing rates) for the i-th trial in the training set, $\overline{Y}_i$ denoting the trial-averaged activity for the condition that the i-th trial is an instantiation of, and † denoting the Moore-Penrose pseudoinverse, and the colon symbol designating how to index matrices (e.g., $M(:,a{:}b)$ refers to the submatrix of M including all rows of M, but only the columns a through b). Lastly, the initial state parameter x_0 was computed by taking the average value of the trial-averaged projections over all times and conditions. Online inference of the underlying neural state, which yields an estimate $\hat{x}_t$ at each millisecond t, was computed recursively using the steady-state form of the Kalman filter[48].

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

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

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

In general, one can compute which of these three distributions is most likely by choosing the maximizing likelihood and assess confidence in that choice by comparing the relative values of

the three likelihoods. However, we wanted the decoder to err on the side of withholding movement. We therefore set a conservative threshold on f_{stop} corresponding to the point at which $L(t)$ would have a Mahalanobis distance of 3 to the stopped distribution of angular momentums. If f_{stop} ever exceeded this threshold, we set $speed_{steady}(t) = 0$. If this condition was not met, we decoded direction and speed as follows:

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

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

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 the likelihoods (yielding a slower velocity if the direction decode is uncertain) and β is a direction-specific constant learned from the training data whose purpose is simply to scale up the result to match steady-state cycling speed. In practice, $speed_{steady}(t)$ was frequently very close to the monkeys' steady-state cycling speeds (**Fig. 5e**).

In the speed-tracking experiment, direction was computed similarly, but with a minor change in filtering and added safeguards against making inferences based on ambiguous evidence. First, the cutoff frequency for the high-pass filter was dropped to 0.75 Hz to allow a greater dynamic range of speeds. Second, the proximity of the observed angular momentums to the distribution for 'stopped' data partially determined transitions into the EXIT state. The condition $L(t) \in$ Stopped, referred to in the decoder state machine (**Supp. Fig. 1**), was considered met when the observed angular momentums had a Mahalanobis distance of less than 4 to the 'stopped' distribution. Third, when in the STEADY state, a new direction was only computed if the following conditions were met: 1) the observed angular momentums had a Mahalanobis distance of less than 4 to the distribution corresponding to the decoded direction, 2) the observed angular momentums had a Mahalanobis distance of greater than 6 to the distribution corresponding to the opposite direction. These criteria ensured that a new steady-state direction was only decoded when the angular momentums were highly consistent with a particular direction. When these criteria were not met, the decoder continued to decode the same direction from the previous time step.

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

$$speed(t) = \begin{cases} \frac{d\theta_f}{dt}, & direction(t) = +1 \\ \frac{d\theta_b}{dt}, & direction(t) = -1 \end{cases}$$

where $\theta_f(t)$ and $\theta_b(t)$ are the phase of the two planes in the neural state estimate $\hat{x}_t$ and the derivative is computed in units of Hz. Additional saturation limits were set such that decoded

speed never dropped below 0.5 Hz or exceeded 3.5 Hz, so as to remain in the range typically seen during pedaling. See the section on 'Decoder State Machine' for details on how the speed estimate was smoothed over time.

Initial-direction Dimensions

To determine the direction of virtual motion at the very beginning of intended movement, we used activity in another set of dimensions. Using the training data, we applied PCA to the neural activity from the period of time around movement onset (0 to 200 ms from decoded movement onset, as detected in the moving dimension) and extracted the top three dimensions to form an initial-direction subspace. We then projected activity at 175 ms after decoded movement onset (the moment when the EARLY state is entered) into the space. We fit 3D Gaussian distributions to the resulting sets of forward and backward neural states.

Online, when the EARLY state was entered, we projected activity into the initial-direction subspace and computed likelihoods $g_{forward}$ and $g_{backward}$ with respect to the learned distributions. If the observed neural state in the initial-direction subspace was not an outlier (>10 Mahalanobis distance units) with respect to both distributions, then the initial direction and speed were computed as follows:

$$direction_{initial} = \mathrm{sgn}(g_{forward} - g_{backward})$$

$$speed_{initial} = \left| 2 \cdot \frac{g_{forward}}{g_{forward} + g_{backward}} - 1 \right| \beta$$

Once the initial direction and speed were computed, these values persisted throughout the remainder of the EARLY state. If the observed neural state was an outlier, initial direction and speed were computed in the same manner as is done in the STEADY state.

In the speed-tracking experiment, there was an additional refinement. Instead of computing the initial direction and speed at a fixed lag (175 ms) relative to decoded movement onset, the transition into the EARLY state was set to wait until a "confident initial direction" was obtained (**Supp. Fig. 1**). This condition was met when the Mahalanobis distance from the neural state in the initial-direction subspace to either the forward or backward distributions dropped below 4. If, after 500 ms, the decoder was still confident that movement was occurring, but the "confident initial direction" condition had not been met, the state machine skipped directly to the STEADY state. Speed during the EARLY state was computed identically to how it was computed during the STEADY state for this experiment.

Acknowledgements. We thank Y. Pavlova for expert animal care, and E. Oby for surgical expertise and assistance. This work was supported by NINDS 1DP2NS083037, NIH CRCNS R01NS100066, NINDS 1U19NS104649, the Simons Foundation (SCGB#325233 and SCGB#542957), the Grossman Center for the Statistics of Mind, the McKnight Foundation, P30 EY019007, a Klingenstein-Simons Fellowship, and the Searle Scholars Program.

References

1. Gilja, V. *et al.* A high-performance neural prosthesis enabled by control algorithm design. *Nat. Neurosci.* **15**, 1752–1757 (2012).
2. Gilja, V. *et al.* Clinical translation of a high-performance neural prosthesis. *Nat. Med.* **21**, 1142–1145 (2015).
3. Shanechi, M. M. *et al.* Rapid control and feedback rates enhance neuroprosthetic control. *Nat. Commun.* **8**, 13825 (2017).
4. Shenoy, K. V. & Carmena, J. M. Combining Decoder Design and Neural Adaptation in Brain-Machine Interfaces. *Neuron* **84**, 665–680 (2014).
5. Ajiboye, A. B. *et al.* Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration. *The Lancet* **0**, (2017).
6. Ethier, C., Oby, E. R., Bauman, M. J. & Miller, L. E. Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature* **485**, 368–371 (2012).
7. Collinger, J. L. *et al.* 7 degree-of-freedom neuroprosthetic control by an individual with tetraplegia. *Lancet* **381**, 557–564 (2013).
8. Wodlinger, B. *et al.* Ten-dimensional anthropomorphic arm control in a human brain−machine interface: difficulties, solutions, and limitations. *J. Neural Eng.* **12**, 016011 (2015).
9. Rajangam, S. *et al.* Wireless Cortical Brain-Machine Interface for Whole-Body Navigation in Primates. *Sci. Rep.* **6**, (2016).
10. Libedinsky, C. *et al.* Independent Mobility Achieved through a Wireless Brain-Machine Interface. *PLOS ONE* **11**, e0165773 (2016).
11. Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S. & Schwartz, A. B. Cortical control of a prosthetic arm for self-feeding. *Nature* **453**, 1098–1101 (2008).
12. Chapin, J. K., Moxon, K. A., Markowitz, R. S. & Nicolelis, M. A. L. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat. Neurosci.* **2**, 664–670 (1999).
13. Wessberg, J. *et al.* Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* **408**, 361–365 (2000).
14. Carmena, J. M. *et al.* Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates. *PLoS Biol* **1**, e42 (2003).
15. Taylor, D. M., Tillery, S. I. H. & Schwartz, A. B. Direct Cortical Control of 3D Neuroprosthetic Devices. *Science* **296**, 1829–1832 (2002).
16. Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R. & Donoghue, J. P. Brain-machine interface: Instant neural control of a movement signal. *Nature* **416**, 141–142 (2002).
17. Scott, S. H., Gribble, P. L., Graham, K. M. & Cabel, D. W. Dissociation between hand motion and population vectors from neural activity in motor cortex. *Nature* **413**, 161–165 (2001).
18. Russo, A. A. *et al.* Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron* **97**, 953-966.e8 (2018).
19. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).

20. Michaels, J. A., Dann, B. & Scherberger, H. Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLOS Comput. Biol.* **12**, e1005175 (2016).
21. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: A dynamical systems perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
22. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* **18**, 1025–1033 (2015).
23. Kao, J. C. *et al.* Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat. Commun.* **6**, 1–12 (2015).
24. Aghagolzadeh, M. & Truccolo, W. Inference and Decoding of Motor Cortex Low-Dimensional Dynamics via Latent State-Space Models. *IEEE Trans. Neural Syst. Rehabil. Eng.* **24**, 272–282 (2016).
25. Makin, J. G., O'Doherty, J. E., Cardoso, M. M. B. & Sabes, P. N. Superior arm-movement decoding from cortex with a new, unsupervised-learning algorithm. *J. Neural Eng.* **15**, 026010 (2018).
26. Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P. & Black, M. J. Bayesian Population Decoding of Motor Cortical Activity Using a Kalman Filter. *Neural Comput.* **18**, 80–118 (2006).
27. Fetz, E. E., Finocchio, D. V., Baker, M. A. & Soso, M. J. Sensory and motor responses of precentral cortex cells during comparable passive and active joint movements. *J. Neurophysiol.* **43**, 1070–1089 (1980).
28. Lemon, R. N., Hanby, J. A. & Porter, R. Relationship between the activity of precentral neurones during active and passive movements in conscious monkeys. *Proc. R. Soc. Lond. B Biol. Sci.* **194**, 341–373 (1976).
29. Suminski, A. J., Tkach, D. C. & Hatsopoulos, N. G. Exploiting multiple sensory modalities in brain-machine interfaces. *Neural Netw.* **22**, 1224–1234 (2009).
30. Schroeder, K. E. *et al.* Robust tactile sensory responses in finger area of primate motor cortex relevant to prosthetic control. *J. Neural Eng.* **14**, 046016 (2017).
31. Pruszynski, J. A. *et al.* Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature* **478**, 387–390 (2011).
32. Christie, B. P. *et al.* Comparison of spike sorting and thresholding of voltage waveforms for intracortical brain–machine interface performance. *J. Neural Eng.* **12**, 016009 (2014).
33. Moran, D. W. & Schwartz, A. B. Motor Cortical Representation of Speed and Direction During Reaching. *J. Neurophysiol.* **82**, 2676–2692 (1999).
34. Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. Neuronal population coding of movement direction. *Science* **233**, 1416–1419 (1986).
35. Kaufman, M. T. *et al.* The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eNeuro* **3**, (2016).
36. Kemere, C. *et al.* Detecting Neural-State Transitions Using Hidden Markov Models for Motor Cortical Prostheses. *J. Neurophysiol.* **100**, 2441–2452 (2008).
37. Kao, J. C., Nuyujukian, P., Ryu, S. I. & Shenoy, K. V. A High-Performance Neural Prosthesis Incorporating Discrete State Selection With Hidden Markov Models. *IEEE Trans. Biomed. Eng.* **64**, 935–945 (2017).
38. Ethier, C., Sachs, N. A. & Miller, L. E. Continuous state-dependent decoders for brain machine interfaces. in *2011 5th International IEEE/EMBS Conference on Neural Engineering* 473–477 (2011). doi:10.1109/NER.2011.5910589.
39. Kao, J. C., Ryu, S. I. & Shenoy, K. V. Leveraging neural dynamics to extend functional lifetime of brain-machine interfaces. *Sci. Rep.* **7**, (2017).
40. Foster, J. D. *et al.* A freely-moving monkey treadmill model. *J. Neural Eng.* **11**, 046020 (2014).
41. Xing, D., Aghagolzadeh, M., Truccolo, W. & Borton, D. Low-Dimensional Motor Cortex Dynamics Preserve Kinematics Information During Unconstrained Locomotion in Nonhuman Primates. *Front. Neurosci.* **13**, (2019).

42. Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A. & Shenoy, K. V. A high-performance brain-computer interface. *Nature* **442**, 195–198 (2006).
43. Sadtler, P. T., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M. & Batista, A. P. Brain–computer interface control along instructed paths. *J. Neural Eng.* **12**, 016015 (2015).
44. Nuyujukian, P., Kao, J. C., Ryu, S. I. & Shenoy, K. V. A Nonhuman Primate Brain–Computer Typing Interface. *Proc. IEEE* **105**, 66–72 (2017).
45. Fernández-Rodríguez, Á., Velasco-Álvarez, F. & Ron-Angevin, R. Review of real brain-controlled wheelchairs. *J. Neural Eng.* **13**, 061001 (2016).
46. Thacker, N. A. & Bromiley, P. A. *The Effects of a Square Root Transform on a Poisson Distributed Quantity.* Technical Report http://www.tina-vision.net/tina-knoppix/tina-memo/2001-010.pdf (2001).
47. Cunningham, J. P. & Ghahramani, Z. Linear Dimensionality Reduction: Survey, Insights, and Generalizations. *J Mach Learn Res* **16**, 2859–2900 (2015).
48. Malik, W. Q., Truccolo, W., Brown, E. N. & Hochberg, L. R. Efficient Decoding With Steady-State Kalman Filter in Neural Interface Systems. *IEEE Trans. Neural Syst. Rehabil. Eng.* **19**, 25–34 (2011).

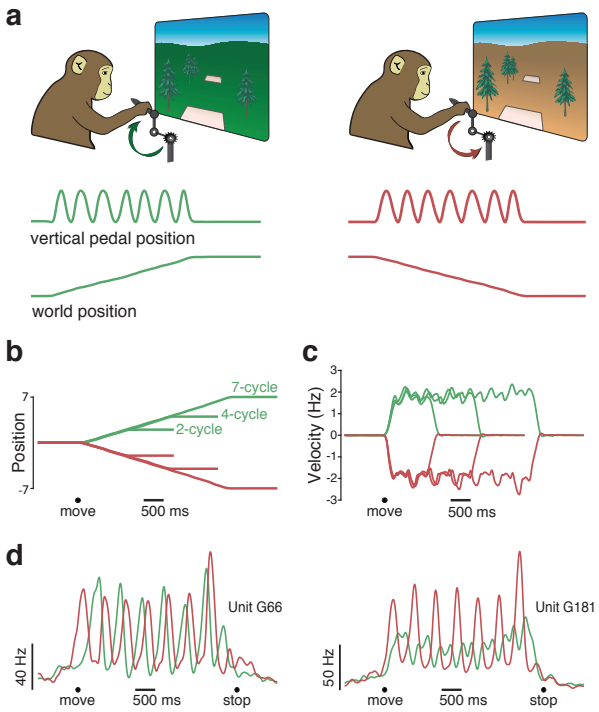

Figure 1. A cycling task that elicits rhythmic movements. (a) Monkeys rotated a hand-held pedal forward (*left*, cued by a green background) or backward (*right*, cued by a tan background) to progress through a virtual environment. Traces at bottom plot pedal kinematics (vertical position) and the resulting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

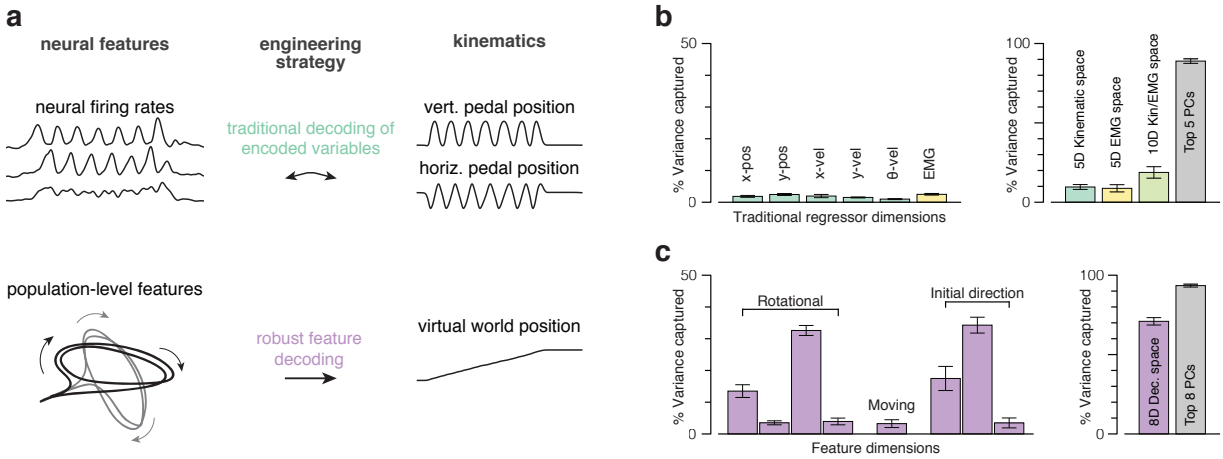

Figure 2. Different decode strategies leverage neural signals with different magnitudes. (a) Two possible decoding strategies. In the first (*top*) neural firing rates are assumed to predominantly encode the key variables. The encoding model is usually assumed to be roughly linear when variables are expressed appropriately. For example, cosine tuning for reach velocity is equivalent to a linear dependence on horizontal and vertical velocity. The goal of decoding is to invert encoding. Thus, decoding dimensions should capture the dominant signals in the neural data (because those are what is encoded). The second strategy (*bottom*) can be applied even if the dominant signals do not have the goal of encoding. This strategy seeks to find neural response features that have a robust relationship with the variable one wishes to decode. That relationship may be complex or even incidental, but is useful if it involves high-variance response features. **(b)** Variance of the neural population response captured by dimensions used to decode kinematic parameters (*green bars*) and muscle activity (*yellow bar*). Data are from three manual-control sessions where units (192 channels per day) and muscles (5-7 channels per day) were recorded simultaneously. Each bar plots the average and standard error across sessions (unless otherwise specified). *Left subpanel*: variance captured for kinematic variables (individual variables shown separately) and muscles (average across 19 recordings, standard error computed across recordings). *Right subpanel*: total variance captured by subspaces spanned by kinematic-decoding dimensions, muscle-decoding dimensions, or both. (These are not the sum of the individual variances as dimensions were not always orthogonal). We had different numbers of EMG recordings per day and thus always selected a subset of five. Variance captured by the top five principal components is shown for comparison. **(c)** Similar plot but for the dimensions upon which our decoder was built. *Left subpanel*: variance captured for each of these eight dimensions. *Right subpanel*: variance captured by the eight-dimensional subspace spanned by those dimensions. Variance captured by the top eight principal components is shown for comparison.

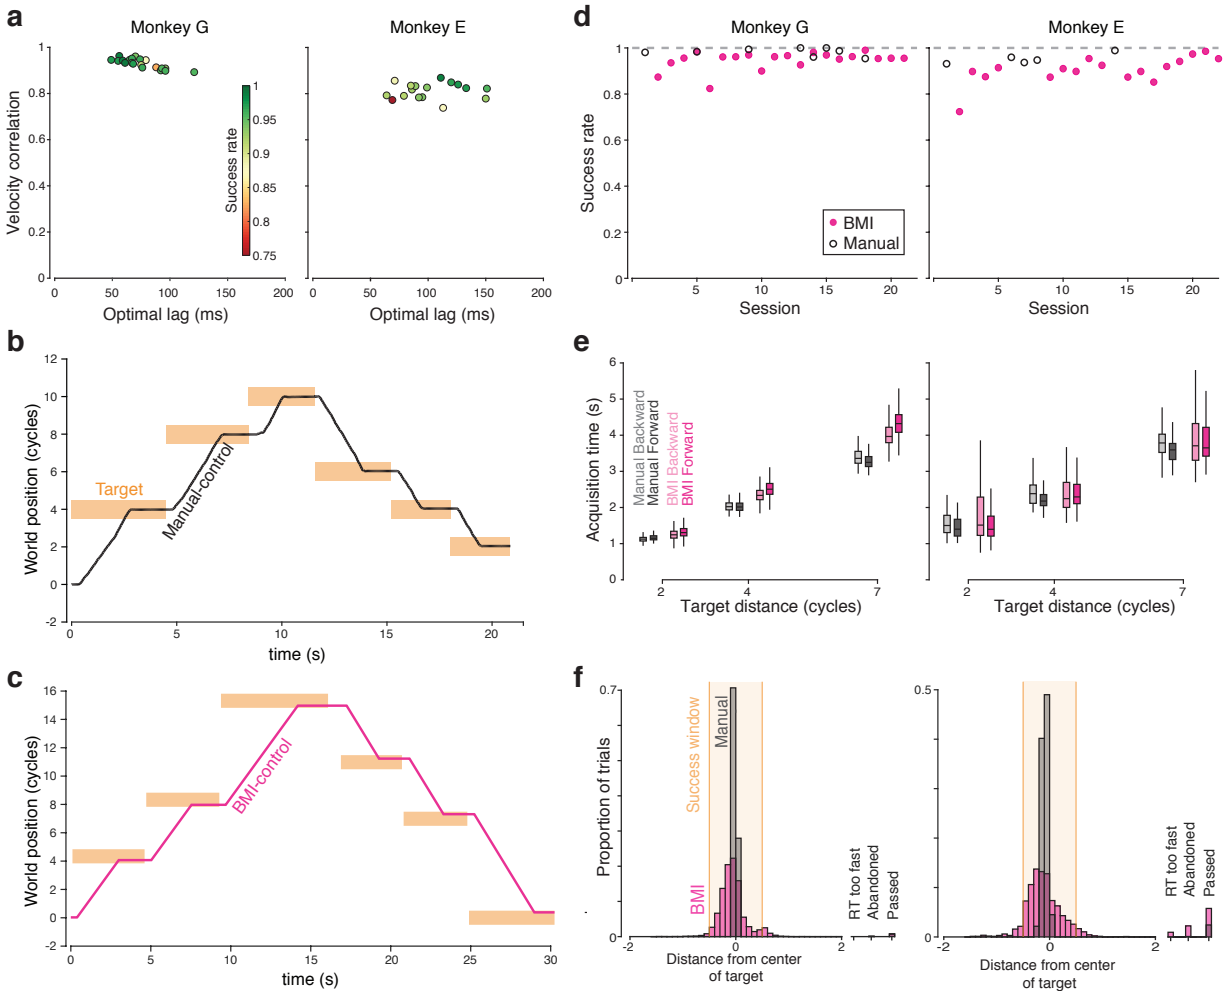

Figure 3. Decoder performance. (a) Summary of the cross-correlation between decoded virtual velocity under BMI control, and the virtual velocity that would have been produced by the pedal (which monkeys continued to manipulate normally). Each symbol corresponds to one BMI-control session, and plots the peak of the cross-correlation versus the lag where that peak occurred. Colors indicate success rate during that session. **(b)** Example manual-control performance for six consecutive trials, 3 forward and 3 backward. World position is expressed in terms of the number of cycles of the pedal needed to move that distance. For plotting purposes, the position at the beginning of this stretch of behavior was set to zero. Bars indicate the time that targets turned on and off (horizontal span) and the size of the acceptance window (vertical span). **(c)** Similar plot during BMI control. For ease of comparison, world position is still expressed in terms of the number of physical cycles that would be needed to travel that far, although physical cycling no longer had any impact on virtual velocity. **(d)** Success rate for both monkeys. Each symbol plots, for one session, the proportion of trials where the monkey successfully moved from the initial target to the final target, stopped within it, and remained stationary until reward delivery. Dashed line at 1 for reference. **(e)** Target acquisition times for successful trials. Center lines indicate median, the box edges indicate the first and third quartiles, and the whiskers include all non-outlier points (points less than 1.5 times the interquartile range from the box edges). Data are shown separately for the three target distances. **(f)** Histograms of stopping location from both monkeys. Analysis considers both

successful and failed trials. The bar at far right indicates the proportion of trials where the monkey failed for reasons other than stopping accuracy per se. This included trials where monkeys disrespected the reaction time limits, abandoned the trial before approaching the target, or passed through the target without stopping.

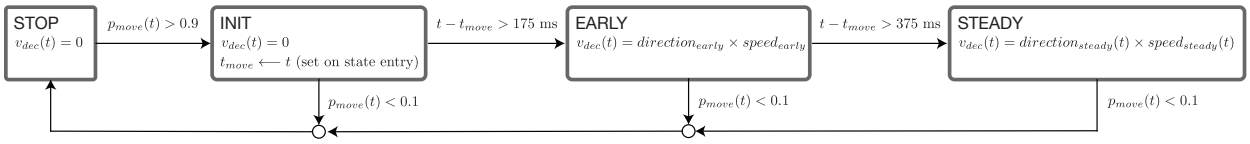

Figure 4. State machine diagram. BMI motion was determined by a state machine with four states: STOP, INIT, EARLY, and STEADY, corresponding to the different stages of a typical trial. The output of the state machine at every millisecond was an estimate of decoded velocity through the virtual environment, v_{dec}, which was then smoothed and integrated to compute virtual position. 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.

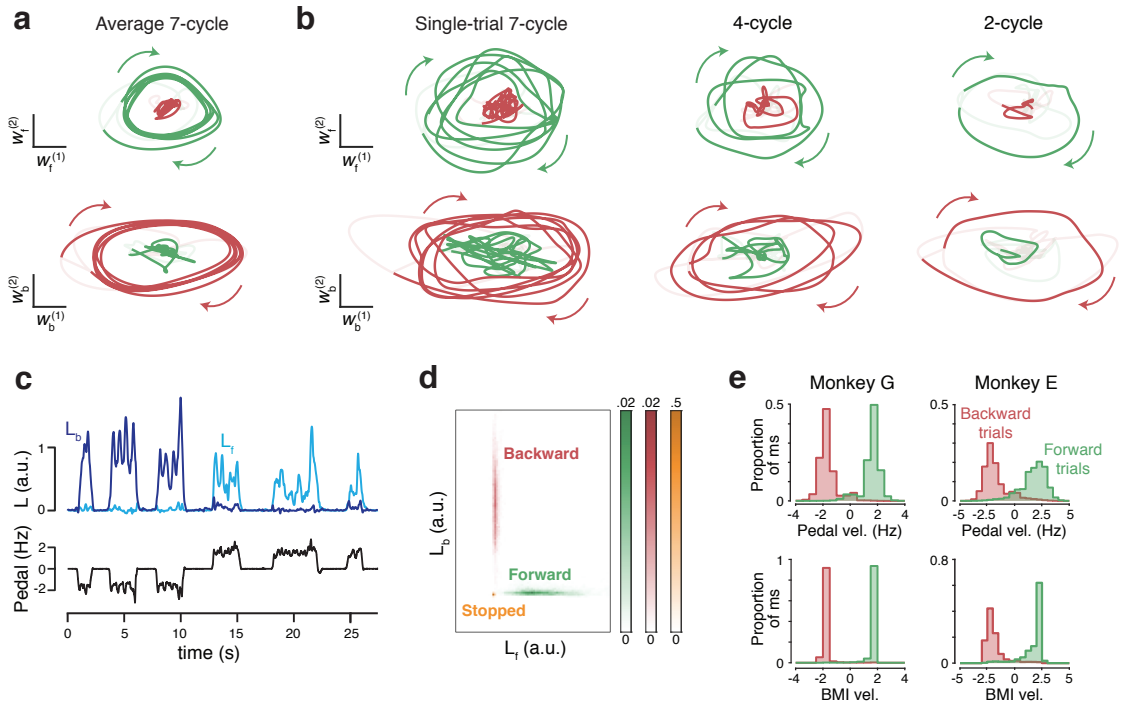

Figure 5. Leveraging rotational trajectories to decode velocity. (a) Trial-averaged population activity, during a manual-control block, projected onto the forward (*top*) and backward (*bottom*) rotational planes. Data are from seven-cycle forward (*green*) and backward (*red*) conditions. By design,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

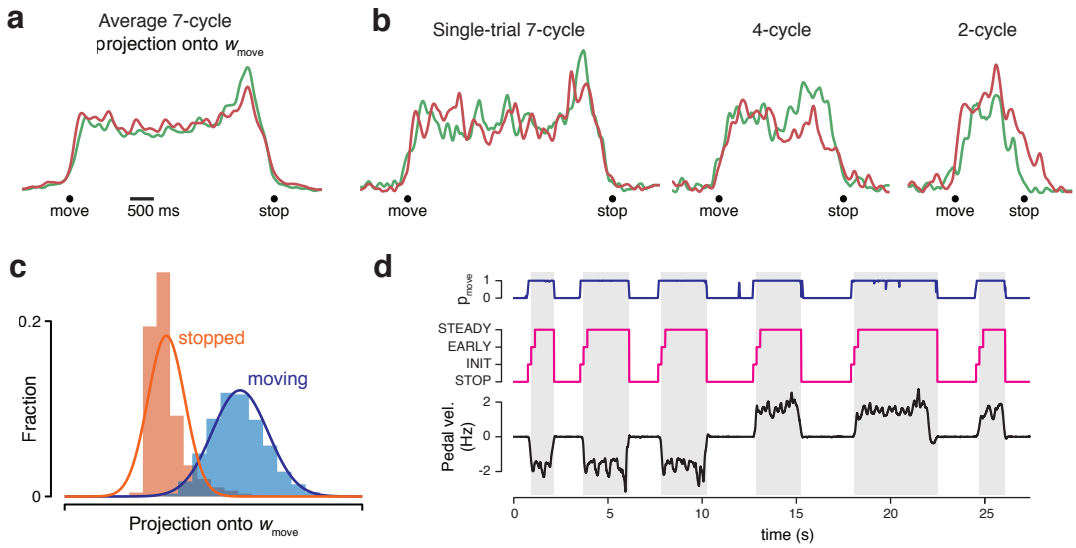

Figure 6. Leveraging the moving-sensitive dimension to infer probability of moving. (a) Trial-averaged population activity, during a manual-control block, projected onto the moving-sensitive dimension (same session and trials as Figure 5a). **(b)** As in panel (a), but for three example single trials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

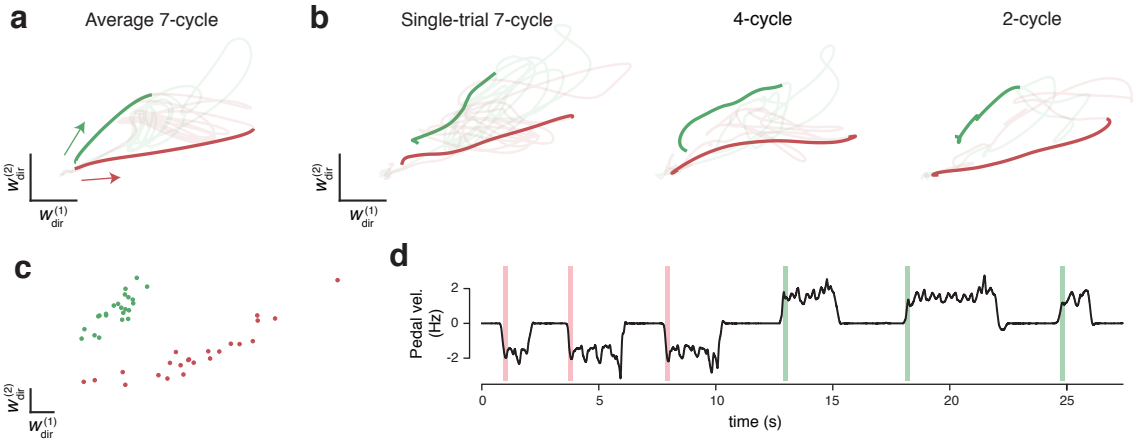

Figure 7. Leveraging initial-direction dimensions to allow low-latency decoding. (a) Trial-averaged population activity, during a manual-control block, projected onto two (of three) initial-direction dimensions (same session and trials as Figure 5a and 6a). Boldly colored portions of traces highlight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

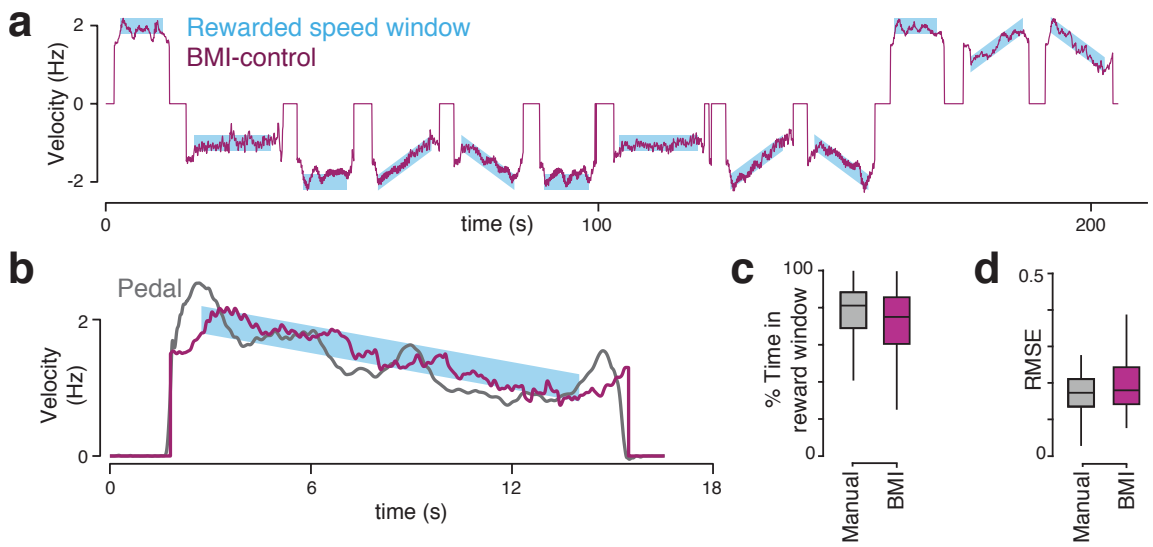

Figure 8. Performance in the modified task requiring speed tracking. (a) Instructed velocity and BMI-decoded virtual velocity during 12 contiguous trials of BMI-control. **(b)** Expanded view of one example trial (the last trial from panel a). The virtual velocity that would have been produced by the pedal is shown in gray for comparison. **(c)** Percentage of time spent in rewarded velocity window while 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)** Root-mean-squared error between instructed velocity and virtual velocity for both manual control and BMI control sessions. Same format as (c).

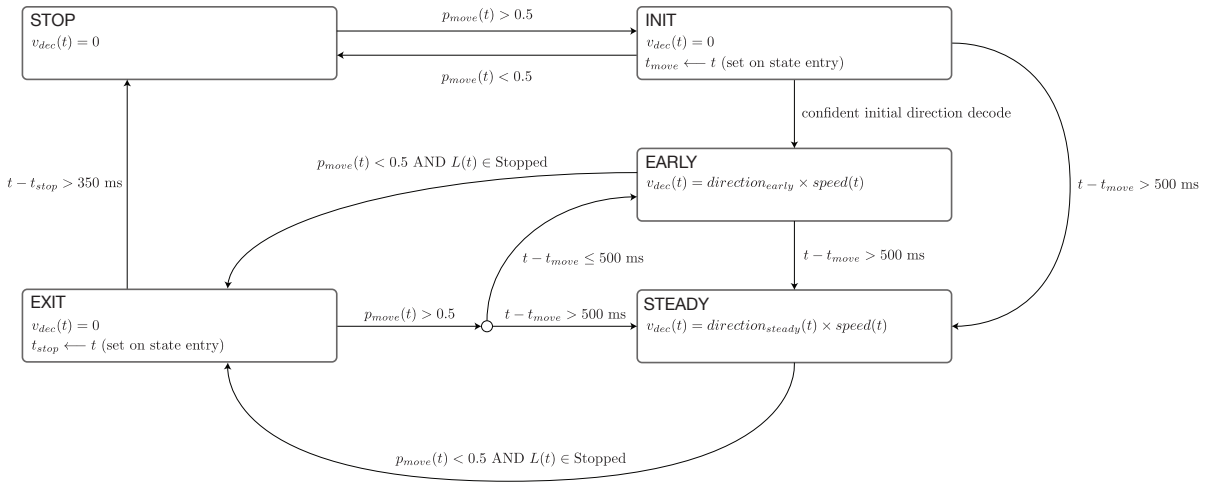

Supplementary Figure 1. State machine diagram for speed-control experiments. For these experiments, the state machine included a fifth state: EXIT. The additional state provided some flexibility to reenter the EARLY and STEADY states without having to pass through the INIT state. This could occur in circumstances where the decoder only briefly believed that movement had ceased. This allowed the monkey to continue moving with minimal interruption in the event of a momentarily incorrect stop decode. This state machine also introduced new criteria for certain state transitions. Angular momentums were used in transitioning to the EXIT state. Furthermore, the transition from INIT to EARLY required that the estimate of initial direction meet a threshold level of confidence. These modifications were non-essential but were deemed useful based on our experience with the decoder used in the primary experiment.

Supplementary Movie 1. BMI and manual control of cycling 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 Fig. 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 one cycle of the final green stopping target (total of 20 cycles from starting position).