

UCSF

UC San Francisco Previously Published Works

Title

Timescales of Local and Cross-Area Interactions during Neuroprosthetic Learning.

Permalink

<https://escholarship.org/uc/item/71k598kb>

Journal

Journal of Neuroscience, 41(49)

ISSN

0270-6474

Authors

Derosier, Katherine
Veuthey, Tess L
Ganguly, Karunesh

Publication Date

2021-12-08

DOI

10.1523/jneurosci.1397-21.2021

Peer reviewed

Timescales of Local and Cross-Area Interactions during Neuroprosthetic Learning

Katherine Derosier,^{1,2,3*} Tess L. Veuthey,^{1,2,3,4*} and  Karunesh Ganguly^{2,3}

¹Neuroscience Graduate Program, University of California San Francisco, San Francisco, California 94158, ²Neurology and Rehabilitation Service, San Francisco Veterans Affairs Medical Center, San Francisco, California 94158, ³Department of Neurology, University of California San Francisco, San Francisco, California 94158, and ⁴Medical Scientist Training Program, University of California San Francisco, San Francisco, California 94158

How does the brain integrate signals with different timescales to drive purposeful actions? Brain-machine interfaces (BMIs) offer a powerful tool to causally test how distributed neural networks achieve specific neural patterns. During neuroprosthetic learning, actuator movements are causally linked to primary motor cortex (M1) neurons, i.e., “direct” neurons that project to the decoder and whose firing is required to successfully perform the task. However, it is unknown how such direct M1 neurons interact with both “indirect” local (in M1 but not part of the decoder) and across area neural populations (e.g., in premotor cortex/M2), all of which are embedded in complex biological recurrent networks. Here, we trained male rats to perform a M1-BMI task and simultaneously recorded the activity of indirect neurons in both M2 and M1. We found that both M2 and M1 indirect neuron populations could be used to predict the activity of the direct neurons (i.e., “BMI-potent activity”). Interestingly, compared with M1 indirect activity, M2 neural activity was correlated with BMI-potent activity across a longer set of time lags, and the timescale of population activity patterns evolved more slowly. M2 units also predicted the activity of both M1 direct and indirect neural populations, suggesting that M2 population dynamics provide a continuous modulatory influence on M1 activity as a whole, rather than a moment-by-moment influence solely on neurons most relevant to a task. Together, our results indicate that longer timescale M2 activity provides modulatory influence over extended time lags on shorter-timescale control signals in M1.

Key words: BCI; cross-area; learning; motor cortex; premotor; timescale

Significance Statement

A central question in the study of motor control is whether primary motor cortex (M1) and premotor cortex (M2) interact through task-specific subpopulations of neurons, or whether tasks engage broader correlated networks. Brain-machine interfaces (BMIs) are powerful tools to study cross-area interactions. Here, we performed simultaneous recordings of M1 and M2 in a BMI task using a subpopulation of M1 neurons (direct neurons). We found that activity outside of direct neurons in M1 and M2 was predictive of M1-BMI task activity, and that M2 activity evolved at slower timescales than M1. These findings suggest that M2 provides a continuous modulatory influence on M1 as a whole, supporting a model of interactions through broad correlated networks rather than task-specific neural subpopulations.

Received July 7, 2021; revised Sep. 29, 2021; accepted Oct. 11, 2021.

Author contributions: K.D., T.L.V., and K.G. designed research; K.D. and T.L.V. performed research; K.D. analyzed data; K.D. and T.L.V. wrote the first draft of the paper; K.D. and K.G. edited the paper.

This work was supported in part by the Department of Veterans Affairs, Veterans Health Administration, Rehabilitation Research and Development Service Award 101RX001640-01A1 (to K.G.) and San Francisco VA Medical Center; the National Institutes of Health (NIH) National Institute of Neurological Disorders and Stroke Award K02NS093014; the University of California San Francisco (UCSF) Neuroscience Graduate Program (K.D. and T.L.V.); the National Defense Science and Engineering Graduate (NDSEG) fellowship awards; the UCSF Discovery Fellows Program; the Markowski–Leach Fellowship (T.L.V.); the UCSF Medical Scientist Training Program (T.L.V.); the NIH National Institute of Mental Health Grant 5R01MH111871 (to K.G.); and start-up funds from the UCSF Department of Neurology (K.G.). K.G. holds a Career Award for Medical Scientists from the Burroughs Wellcome Fund (1009855) and an Independent Scientist Award (1K02NS093014) from the National Institute of Neurological Disorders and Stroke. We thank R. Kulkarni and N. Hoglen for providing comments on this manuscript.

*K.D. and T.L.V. contributed equally to this work.

The authors declare no competing financial interests.

Correspondence should be addressed to Karunesh Ganguly at karunesh.ganguly@ucsf.edu.

<https://doi.org/10.1523/JNEUROSCI.1397-21.2021>

Copyright © 2021 the authors

Introduction

Mammalian cortex is characterized by dense local and sparse cross-area connections, suggesting a vital role for dynamic interactions between local and cross-area processes (Sporns et al., 2004). Take, for example, premotor (M2) and motor cortex (primary motor cortex; M1): they have connectivity both within and across areas, and both are known to play an important role in motor learning (Hikosaka et al., 2002; Ganguly and Carmena, 2009; Ganguly et al., 2011; Peters et al., 2014; Cao et al., 2015; Kawai et al., 2015; Athalye et al., 2017; Makino et al., 2017). While there are known changes in activity patterns localized to both M2 (Cao et al., 2015; Makino et al., 2017) and M1 (Peters et al., 2014; Ramanathan et al., 2015; Lemke et al., 2019) with learning, there is also greater task-related cross-area coordination with learning (Veuthey et al., 2020). Together, this supports a

hierarchical “top-down” influence of M2 on M1 during learning and adaptation (Hikosaka et al., 2002; Siniscalchi et al., 2016; Omlor et al., 2019). However, it is also well known that both M2 and M1 have projections to subcortical and spinal regions (Nudo and Masterton, 1990; Dum and Strick, 2002; Wang et al., 2017). Thus, with learning, it is difficult to tease apart whether changes in interactions between M2 and M1 are because of alterations in cross-area communication, or instead because of their roles coordinating the upper limb (Whishaw and Pellis, 1990).

How then might we isolate M1 and M2’s roles in movement control from M2’s top-down influence on M1? Brain-machine interfaces (BMIs) offer a powerful tool to causally test how networks change with learning of neuroprosthetic skills (Taylor et al., 2002; Ganguly et al., 2011; Gulati et al., 2014, 2017; Sadtler et al., 2014; Kim et al., 2019). During “brain control,” actuator movements are causally linked to an ensemble of neurons in M1, “direct neurons” (Taylor et al., 2002). Such direct neurons change their firing properties during neuroprosthetic learning (Ganguly et al., 2011; Arduin et al., 2013; Gulati et al., 2014, 2017; Athalye et al., 2017, 2018); in contrast, the vast majority of neurons in the larger M1 network, “indirect neurons,” generally become less task coupled (Fetz, 2007; Ganguly et al., 2011; Koralek et al., 2012, 2013; Arduin et al., 2013; Clancy et al., 2014; Gulati et al., 2017; Silversmith et al., 2021; but see Liu and Schieber, 2020). It remains unknown how neural activity in M2, i.e., indirect neurons which shape M1, plays a role in neuroprosthetic learning. Here, we hypothesized that M2 neurons will demonstrate task-related firing during neuroprosthetic learning. We also predicted that analyzing interactions between M2 and M1 can reveal timescales and modes of interactions within and across structures during learning.

Here, we use a BMI paradigm in which a small selection of M1-direct neurons controlled a simple 1-D actuator (Koralek et al., 2012; Gulati et al., 2014, 2017; Kim et al., 2019). We additionally record neural activity from indirect neurons in both M1 and M2. First, we found that M2 neurons are modulated by M1-BMI learning, reinforcing M2’s indirect involvement in downstream M1-activity. Intriguingly, M2 population activity predicted M1-direct neural activity at many time lags, suggesting continuous rather than intermittent influence on M1. This diversity of time lags was also present when predicting M1-indirect activity from M2, indicating that M2 interacts with M1 as a whole for BMI tasks, rather than specifically modulating direct neurons as a task-relevant group. However, this diversity of time lags was not the case for M1-indirect to M1-direct predictions, which instead were optimal at short time lags, indicating a narrower temporal relationship within M1 than between M1 and M2. This difference in time lags may be inherent to each population, as M2 activity patterns evolved slower than M1-indirect patterns, neither of which directly controlled the actuator. Together, our results support a framework for population-based cross-area interactions in which M2 activity provides a modulatory influence that evolves at longer timescales.

Materials and Methods

Animal care

All procedures were in accordance with protocols approved by the Institutional Animal Care and Use Committee at the San Francisco Veterans Affairs Medical Center. Adult male Long-Evans rats ($n = 4$, 250–400 g; Charles River Laboratories) were housed in a 12/12 h light/dark cycle. All experiments were done during the light cycles. Rats were housed in groups of two animals before surgery and individually after surgery.

Surgeries

All surgical procedures were performed using a sterile technique under 2–4% isoflurane. Surgery involved cleaning and exposure of the skull, preparation of the skull surface (using cyanoacrylate) and then implantation of the skull screws for overall headstage stability. Reference screws were implanted posterior to λ and ipsilateral to the neural recordings. Craniotomy and durotomy were performed, followed by implantation of the neural probes. Postoperative recovery regimen included the administration of 0.02 mg kg⁻¹ buprenorphine for 2 d, and 0.2 mg kg⁻¹ meloxicam, 0.5 mg kg⁻¹ dexamethasone and 15 mg kg⁻¹ trimethoprim sulfadiazine for 5 d. All animals were allowed to recover for one week before further behavioral training.

Electrode array and cannula implants

Long-Evans hooded rats were implanted with two 32-channel tungsten wire probes (TDT or Innovative Neurophysiology), one each in the forelimb regions for M1 (Neafsey et al., 1986; Ramanathan et al., 2015, 2018; Lemke et al., 2019; +0.5 AP, +3.5 ML, -1.5 DV) and M2 (Neafsey et al., 1986; Murakami et al., 2014; +4.0–4.5 AP, +1.5 ML, -1.5 DV). Subject-specific implantation sites were not based on motor mapping. However, a subset of rats also performed reaching tasks and had robust activation during reaching (Veuthey et al., 2020). One rat (T391) had a cannula attached to the lateral side of the M2 electrode array before surgery for use in a different experiment (Veuthey et al., 2020).

General BMI paradigm

Rats were trained using an automated behavior box, with components controlled by MATLAB R2015a and an Arduino running the Adafruit Motor Library V1. Within the box, rats were unrestrained. Neural data were recorded and sorted online using software from Tucker Davis Technologies, either OpenEx or Synapse. Spike counts from online sorting were imported into MATLAB and used to control the feedback stimuli (for details, below, BMI control). Trials started with an auditory cue and the opening of the plastic gate covering a slot in the back of the behavior box. When rats achieved the neural firing rate target, success was indicated with an auditory cue, and water reward was delivered via a metal spout through the slot. If rats failed to achieve the target in the set time, failure was indicated with an auditory cue and the closing of the gate, followed by a timeout period. The maximum trial length and the timeout period following failures were both manipulated over the course of the experiments to encourage learning, and ranged from 10 to 20 and 5 to 10 s, respectively.

BMI control

The spout BMI paradigm was used to train $n = 4$ rat. In this paradigm, feedback about progress to the firing rate target was given via the movement of the water spout used for reward. A total of 4–8 direct units were chosen based on having good signal-to-noise and neither unusually high nor unusually low firing rates. Of the direct units, 2–4 units were arbitrarily assigned to the “positive pool,” and 2–4 units were arbitrarily assigned to the “negative pool.” The same channels were used for all sessions, but we did not directly test for unit similarity across days. At the beginning of each session, a 5–30 min baseline recording was taken, and used to fit mean firing rates for each unit. During the task, for every 100-ms bin, direct unit firing rates were computed, mean subtracted, and summed within pools. For one rat, the “neural state” was computed as $s = g \times (p - n)$, where p is the firing rate of the positive pool, n is the firing rate of the negative pool, and g is an experimenter-controlled gain parameter. For three rats, γ distributions were fitted to the histogram of firing rate differences using the MATLAB function *fitdist*. During the task, for every 100-ms bin, firing rates were summed within the positive and negative pools, and the difference between the two pools was computed. This difference was fed into the cumulative distribution function of the baseline distribution to obtain the neural state. The neural state was smoothed by averaging it with its previous value, and then used to control the position of the water spout, such that increasing the difference between positive and negative pools moved the spout toward the rat. Once the neural state crossed a threshold value, the trial was considered a success.

Three rats received visual feedback on their performance. In this paradigm, feedback about progress to the firing rate target was given via the movement of both a visual cue on a computer monitor placed outside the behavior box and of the water spout used for reward. First, a computer monitor outside the behavior box displayed a circular “cursor” that moved along a line toward a stationary “target” circle. The position of the cursor along the line was a direct readout of neural state, moving from the top left to the bottom right of the screen (i.e., closer to the rat) as neural state increased.

Electrophysiology data collection

We recorded extracellular neural activity using tungsten microwire electrode arrays (MEAs; $n=4$ rats, TDT or Innovative Neurophysiology). We recorded spike and LFP activity using a 128-channel TDT-RZ2 system (TDT). Spike data were sampled at 24,414 Hz. Analog headstages with a unity gain and high impedance (~ 1 G Ω) were used. Snippets of data that crossed a high signal-to-noise threshold (four standard deviations away from the mean) were time-stamped as events, and waveforms for each event were peak aligned. For BMI control, spikes were sorted online using either TDT OpenEx ($n=1$) or TDT Synapse ($n=3$) software. MEA recordings were sorted offline using a density-based clustering algorithm (Chung et al., 2017). Clusters interpreted to be noise were discarded, but multiunits were kept for analysis. Trial-related timestamps (i.e., trial onset, trial completion) were sent to the RZ2 analog input channel using an Arduino digital board and synchronized to neural data.

The pool of M1 direct neurons ranged from four to eight neurons. The M1 indirect pool contained 23 units ± 10 (SD), and the M2 pool contained 29 units ± 9 (SD).

Circular shuffle test for single unit modulation

A circular shuffle test was used to determine whether single units were significantly task-modulated. First, the unit’s true perievent time histogram (PETH) was calculated by aligning data to the end of successful trials and binning at 15 ms. Then, the PETH was given a modulation score by calculating its difference from a flat mean using the MATLAB function *immse*. Then, for 10^5 repetitions, trials were independently randomly circularly shuffled, and a shuffled PETH and modulation score were calculated. A unit was considered significantly task-modulated if its true modulation score exceeded the 95th percentile of shuffled modulation scores.

M1 BMI-potent space

M1 BMI-potent activity was reconstructed from offline-sorted data by binning the data according to online timestamps and identifying the units on M1 direct channels that were most correlated with the recorded online firing rates. Once the offline-sorted M1 direct units had been identified, neural data could be binned at any binwidth, and M1 BMI-potent activity could be computed as the difference between summed M1d⁺ activity and summed M1d⁻ activity and used as the response variable in the generalized linear models (GLMs).

GLMs

Regression was performed using GLMs fitted with the MATLAB function *fitglm*. Predictors were binned firing rates of either all M2 units or all M1 indirect units, and the response variable was either M1 BMI-potent space activity (described above), or surrogate M1 indirect activity (for GLM-2i models, see below). In all cases, each predictor unit was represented multiple times, with data shifted to include all time lags from -0.2 to $+0.2$ s relative to the response variable. For the response variable, data used was the final 2 s of all trials of at least 2-s duration; for the predictors, the incorporation of time lags meant that data between 2.2 s before the end of the trial to 0.2 s after the end of the trial was used. Models were fit to data binned at 15, 25, 35, 50, and 100 ms. Crossvalidated R^2 values were computed for all binwidth models, and weight structures were qualitatively similar across binwidths. Statistics were performed on 25-ms models. In the case of GLM-2i models, a “surrogate BMI-potent space” was created from M1 indirect activity by randomly selecting matched numbers of indirect units to stand in for the

true direct units. The randomly chosen indirect units were summed within pools and the difference between the pools was used as the response variable. This process was repeated for 50 choices of indirect units per dataset.

For each model, a crossvalidated R^2 was computed and used as a metric of how well the model generalized to held out data. The full dataset was randomly partitioned into 10 equal folds (ignoring trial structure, i.e., timepoints from the same trial could be assigned to different folds). Then, ten different times, 1-fold was assigned to be the test data and the other nine to be the training data. GLM models were fit to the training dataset. The test data were then projected onto the training model, and R^2 values were computed between the true response variable and the model output. The R^2 values reported are the average across all 10 combinations of testing/training data.

Population correlations

Population correlations were computed for either all M2 or all M1 indirect units. Neural data from the given population was binned at 50 ms, and the set of firing rates at a given time bin was considered a population vector. For all successful trials of at least 2-s duration, pairwise correlations were computed between population vectors for all time bins from -2 to 0 s (relative to success) using the MATLAB function *corrcoef*. The results were averaged across all correct trials of at least 2-s duration and all datasets.

Statistical analysis

Statistical tests were done using bootstrapping and permutation tests. Unless otherwise specified, simulations had 10^3 repetitions, and the significance threshold was set at $\alpha = 0.05$, meaning that tests were significant when the true value was greater than the 95th percentile of the reference distribution.

To calculate success rate across animals, the success rate from the last 1/3 of the session (late) was subtracted from the success rate from the first 1/3 of the session (early) for each dataset and averaged across datasets. This difference between success rates was compared with a bootstrap distribution of differences with the early and late success rates shuffled between categories.

To determine whether GLMs had significant predictive power, their crossvalidated R^2 (described above) was compared with a reference distribution of crossvalidated R^2 values for models fitted to trial-shuffled data. Shuffling the trials in the response variable before fitting the models created a control dataset with equivalent firing rates and the same characteristic pattern of neural state increasing at the end of successful trials, but broke true moment-by-moment correlations. Trial-shuffling was performed for 25-ms models only. A total of 10^3 shuffles were performed. Models were considered to have significant predictive power if the true R^2 exceeded the 95th percentile of the reference distribution, which was the case for all seven datasets.

To determine whether the distribution of timelags that had the largest magnitude GLM weight (τ_{\max}) were significantly nonuniform, the τ_{\max} histograms were compared with surrogates drawn from a uniform distribution as follows: first, the probability-normalized τ_{\max} histogram was considered to be the empirical distribution, and a theoretical discrete uniform distribution was calculated as the mean value of the empirical distribution. The difference between the empirical and theoretical distributions was computed using the MATLAB function *immse*. Then, for each bootstrap repetition, a surrogate set of τ_{\max} values was drawn from a uniform distribution using the MATLAB function *datasample*, and the difference from the theoretical distribution was calculated in the same manner to create the reference distribution.

To determine whether the GLM-2d models had significantly different R^2 values from the GLM-2i models, permutation tests were performed separately on each dataset. First, for each dataset, the true difference between the GLM-2d R^2 value and the mean GLM-2i R^2 value was calculated. Then the true GLM-2d R^2 value was pooled with the 50 GLM-2i R^2 values (from different choices of indirect units, see above, GLMs), resulting in a set of 51 R^2 values. Then, for each permutation, one of the R^2 values was chosen to stand in for the GLM-2d R^2 value, and the difference between that and the mean of the remaining 50 R^2

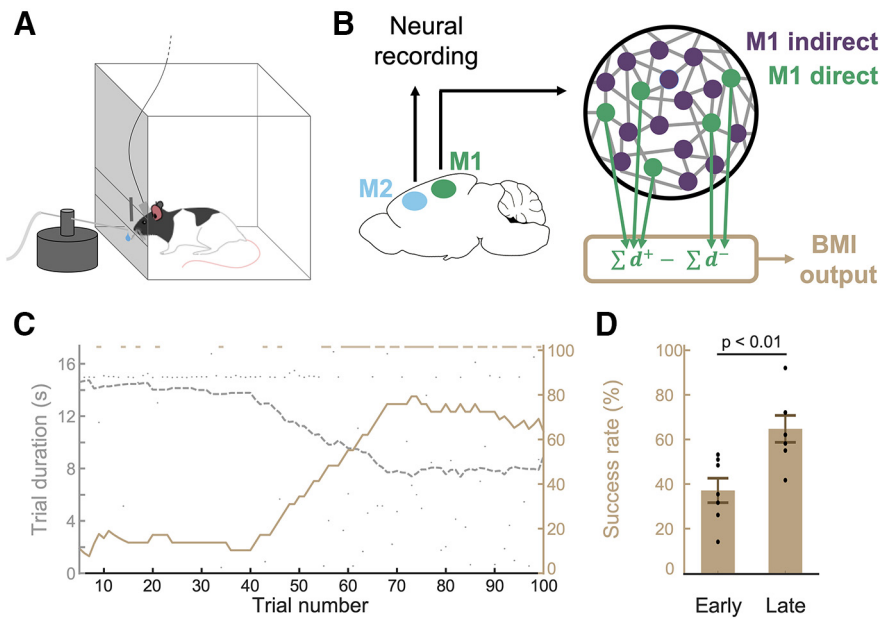


Figure 1. Learning the BMI task. **A**, Experimental setup (not to scale). Rats were allowed to move freely inside a clear plastic behavior box during the task. Online feedback about neural state was provided via the angular position of the waterspout (all $n = 4$ rats, see Materials and Methods). When the target neural state was reached, the waterspout moved within reach and a water reward was given. Auditory cues marked trial start, reward, or task timeout. **B**, left, Only M1 neurons were used in the BMI decoder. Neural activity from M2 was recorded but was not directly decoded. Right, Subsets of M1 neurons were arbitrarily chosen as direct (green), used to drive BMI decoder, while rest were indirect units (purple). One to four direct units were assigned to the positive pool (d^+) and one to four direct units were assigned to the negative pool (d^-). **C**, Example learning curve. Gray points indicate actual trial duration. Dashed gray line shows running average trial duration (smoothed over 30 trials). Light brown boxes at the top of the plot indicate successful trials. Brown line shows running average success rate (smoothed over 30 trials). **D**, Trials early in the session were significantly more successful than trials late (early: first 1/3 trials, $37.30 \pm 5\%$ SEM; late: last 1/3 trials, $64.87 \pm 6\%$ SEM success; $p < 0.01$). Error bars represent SE ($n = 7$).

values was calculated to create the reference distribution. By this metric, the GLM-2d R^2 values were not significantly different from the GLM-id R^2 values for six of the seven datasets; for the remaining dataset the GLM-2d R^2 value was significantly higher than the mean GLM-id R^2 value.

To determine whether the GLM-2d models had significantly lower R^2 values from the GLM-id models, a permutation test was used. First, the mean GLM-2d and GLM-id R^2 values were calculated across datasets, and the true difference between them was computed. Then, for each permutation, the GLM-2d and GLM-id R^2 values were randomly swapped (or not) with 50% probability, and the difference in means between the two groups was calculated to create the reference distribution.

To determine whether M2 single units had a significantly wider autocorrelation function than M1 indirect units, a permutation test was used. First, the autocorrelation of each M2 and M1 indirect unit was computed using the MATLAB function *xcorr*, the width-at-half-max was calculated for each autocorrelation, and the true difference in mean values between M2 and M1 indirect was calculated. Then, for each permutation, units were randomly assigned to two groups, matched in size to the true M2 and M1 indirect populations, and the difference in mean autocorrelation width between the two groups was calculated to create the reference distribution.

To determine whether M2 population activity was more correlated across time than M1 indirect population activity, a permutation test was used. First, M2 and M1 indirect population correlations were computed as described above. Then, the mean M2 and M1 indirect population correlations was taken across all time bin pairs, and the true difference between the two means was calculated. Then, for each permutation, correlation values were randomly swapped between M2 and M1 and the difference between the mean values was calculated to create the reference distribution.

Data sharing

The datasets generated and analyzed in the current study are available from the corresponding author on reasonable request.

Results

M2 units are modulated by a simple M1 BMI task

We recorded simultaneously from both M1 and M2 in rats learning a simple linear BMI task (Fig. 1A,B; Koralek et al., 2012; Gulati et al., 2014, 2017; Athalye et al., 2018; Kim et al., 2019). In each recording session, a subset of M1 neurons were chosen to drive the decoder (M1 direct, or M1d, units). M1 neurons that were excluded from the decoder were considered indirect (M1i) units. The M1 direct units were split into two groups: the positive pool (M1d⁺) and the negative pool (M1d⁻), each consisting of one to four arbitrarily selected units (Fig. 1B). At each timestep during the experiments, online firing rates were computed for all of the M1 direct units. Firing rates were summed within the two pools, and then the summed negative pool firing rate was subtracted from the positive pool to obtain the neural state used to drive the BMI. Neural state values were compared with a baseline distribution, and a trial was successful when the neural state exceeded a target threshold (see Materials and Methods). Therefore, rats were rewarded when they increase the difference in firing rate between the positive and negative pools.

Feedback was given via changes in the angular position of a waterspout ($n = 4$ rats; Fig. 1A). Water reward was delivered on successful task completion; trials that were not completed within 15 s ($n = 1$ rat) or 10 s ($n = 3$ rats) were considered failures. Rats were generally able to learn the task over the course of 100–200 trials, as indicated by increased success rate (Fig. 1C,D).

First, we asked whether M2 units and M1 indirect units were task modulated. Although these neurons do not directly drive the BMI decoder (Fig. 2A), it is possible that they directly affect the M1 direct units. We computed PETHs for all offline-sorted units relative to the end of successful trials, and used a circular shuffle modulation test to determine which units were significantly modulated. We found that 31.68% of M2 units and 36.02% of M1 indirect units were significantly modulated, compared with 70.83% of offline-sorted units from channels that had contained an online-sorted direct unit. Significantly task modulated units are shown in Figure 2B,C. This indicates that during this relatively early stage of BMI training, a large portion of recorded M1 and M2 neurons were involved in rats' attempts to perform the BMI task, although only a few neurons were directly causal to the BMI.

M2 population activity predicts M1 BMI-potent activity at diverse time lags

Although we found that many M2 neurons were significantly task-modulated, it is not necessarily the case that their activity pattern is related to BMI control: an M2 neuron could have

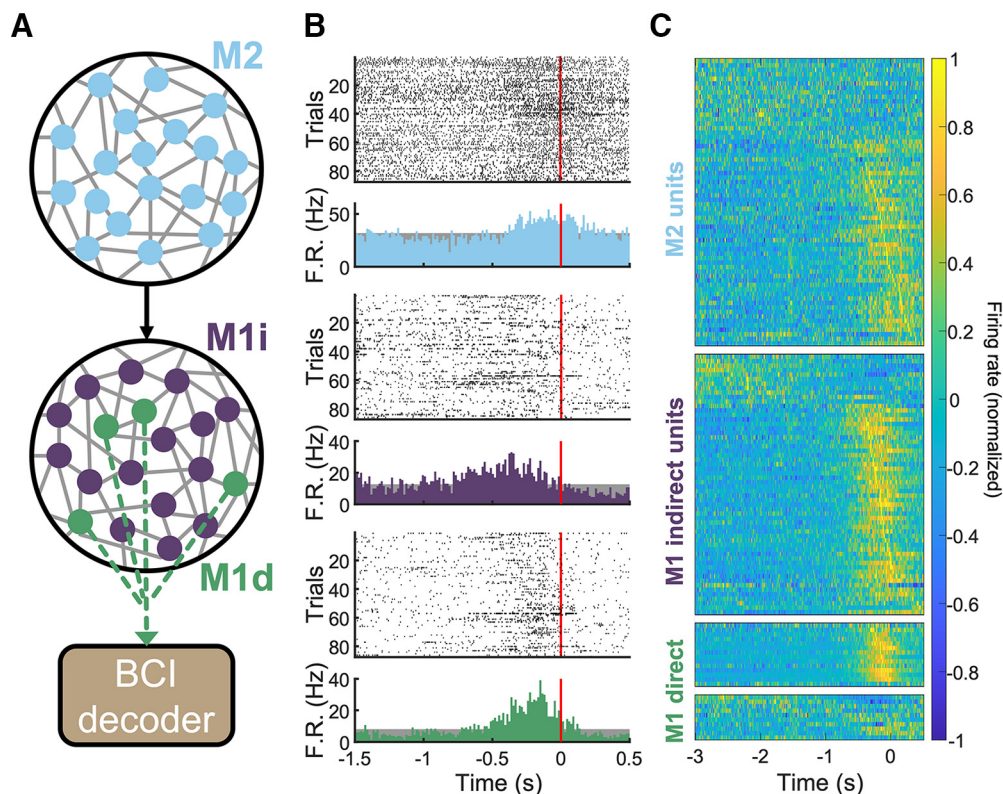


Figure 2. Indirect neurons in both M2 and M1 are modulated during BMI control. **A**, Illustration of the relationship between recorded neurons and the BMI decoder. **B**, Rasters and PETHs for example M2 (top), M1i (middle), and M1d (bottom) units. Data are aligned to the end of successful trials (red bar). Gray shaded region represents a circular shuffle test. **C**, PETH showing average activity patterns from significantly modulated units in all rats. Units are grouped as: M2, M1i, M1d⁺, M1d⁻. Within each group, units are ordered by the time of their peak firing rate. Time is relative to the end of successful trials. Color indicates normalized firing rate ($n = 387$ units from 7 BMI sessions in 4 animals).

a stereotyped activity pattern without being correlated with moment-by-moment BMI activity. To address this question, we used GLMs to perform regression predicting M1 BMI-potent activity from M2 population activity (GLM-2d models; Fig. 3A). As the goal of this analysis was to investigate patterns of simultaneously-recorded population activity in M2 and M1, we included single units both with and without significant trial-averaged PETH task modulation. First, M1 BMI-potent activity was reconstructed, re-binned at 25 ms, and used as the response variable in the GLMs. The predictors were binned M2 firing rates, where each neuron appeared more than once with variable time lags ranging from -0.2 to $+0.2$ s relative to the M1 BMI-potent activity. Model significance was computed for 25-ms models by comparison to a reference distribution of models fit to trial-shuffled data. By this metric, M2 population activity had significant predictive power for moment-by-moment M1 BMI-potent activity in all datasets (Fig. 4B; see Materials and Methods).

Next, we examined the temporal structure of the regression weights. Across the population, different M2 units had their largest magnitude weight at a wide range of time lags (Fig. 3C). Additionally, many single M2 units had large regression weights at multiple time lags, often including both positive and negative weights (Fig. 3D). Across all datasets, the distribution of time lags at which any neuron had its largest magnitude regression weight was not significantly different from a uniform distribution (Fig. 3E; bootstrap, $p = 0.2435$; for details, see Materials and Methods). These results indicate that M2 and M1 interactions occur at a broad range of time lags.

M1 direct units do not have a privileged functional relationship with M2

Next, we wondered whether M2 activity was particularly predictive of M1 direct unit activity, or whether it had similar predictive power for any arbitrarily chosen subset of M1 units. To address this, we randomly chose M1 indirect units to use as “surrogate direct” units, matched in number to the M1 direct units from each dataset, and built GLMs to predict their activity as before (GLM-2i models). For each dataset, we repeated this process with 50 different sets of M1 indirect units. As with the GLM-2d models, we found that many M2 units had large regression weights at multiple time lags, with both positive and negative weights (Fig. 4C). The distribution of time lags for maximally predictive weights was also consistent with a uniform distribution (Fig. 4D,E; bootstrap, $p = 0.491$). For the 25-ms models, we also compared crossvalidated R^2 values, and found that the R^2 values for the GLM-2d models were within the range of R^2 values for the GLM-2i models (Fig. 4B), and were not significantly different from the mean GLM-2i R^2 models for six of seven datasets (permutation tests, $p = 0.1164$, $p = 0.1568$, $p = 0.8609$, $p = 0.9609$, $p = 0.3701$, $p = 0.6454$; remaining dataset, $p = 0.0192$). This suggests that, at least at this early stage of learning, the relationship between M2 and M1 is quite similar for both M1 direct and indirect neural populations.

M2 units are distinct from M1 indirect units

Do M2 and M1 indirect units have similar roles? M2 and M1 are heavily interconnected, both include task-modulated indirect units, and are jointly involved in natural movements. To test this idea, we built GLMs to predict M1 BMI-potent activity from M1

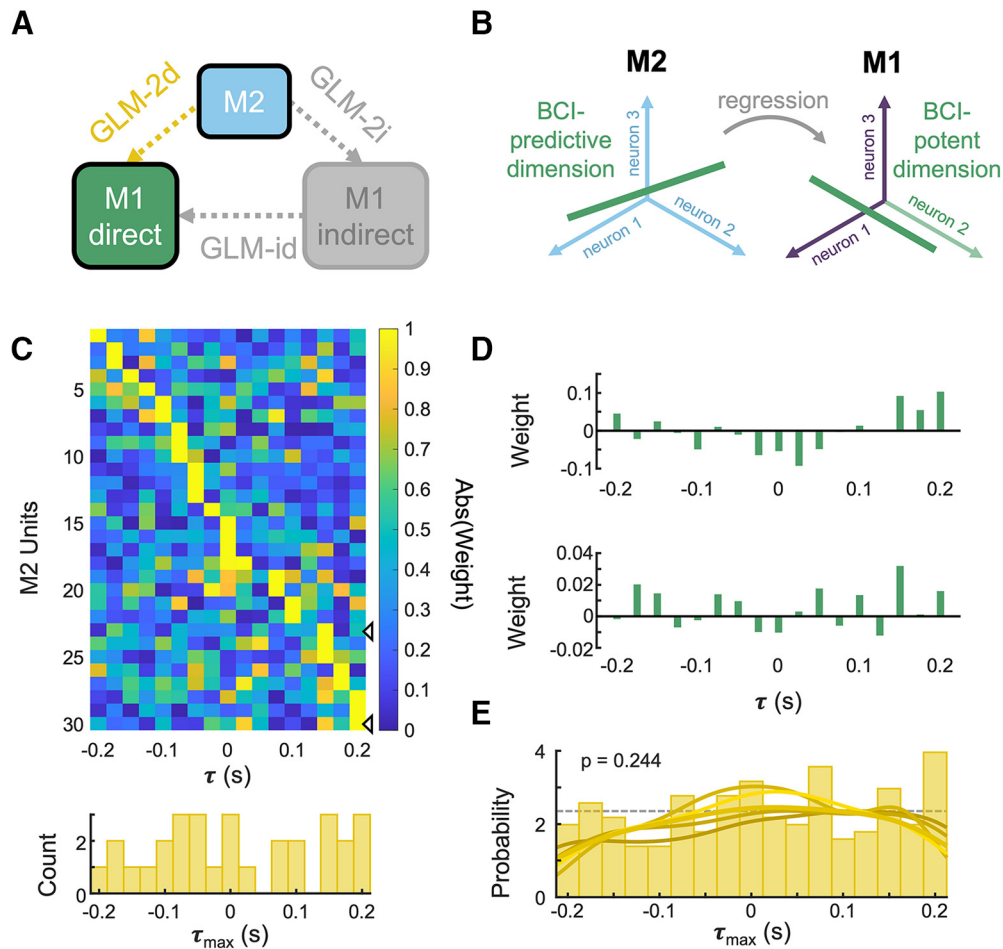


Figure 3. M2 neural activity predicts M1 BMI-potent neural activity at multiple time lags. **A**, GLM-2d indicates models which predict M1 direct activity from M2 activity. In color are regions and interactions analyzed in this figure. In gray are regions and interactions analyzed in subsequent figures. **B**, Regression was used to identify a M2 neural population space that predicted BMI-potent M1 activity. GLMs were fit to predict the M1 BMI neural state from M2 population activity; multiple time lagged copies of each M2 unit were used as predictors. **C**, Distribution of regression weight magnitude in one example session. Top, For each M2 unit, regression weights were assigned for a variety of time lags. To emphasize the time of the maximum absolute weight of each neuron, values here are normalized to each neuron's maximum value, and the absolute value of those weights is indicated by color. Units are sorted according to the time of the largest magnitude weight. Tick marks on the right edge indicate the units shown in **D**. Bottom, Histogram of the τ values with the largest magnitude weight for this dataset. **D**, Example non-normalized weights for two M2 neurons from one example session (neural data binned at 25 ms). Height of green lines indicates weights for example neurons at different time lags (τ) relative to the M1 BMI-potent activity, with negative τ values meaning that M2 leads. **E**, Histogram of time lags that had the largest magnitude regression weight for all M2 units in all animals, for models fitted to neural data binned at 25 ms. Yellow lines show kernel density fit to equivalent histograms for models fit to neural data binned at 15 (lightest), 25, 35, 50, and 100 ms (darkest). Dashed gray line shows theoretical uniform distribution, empirical distribution is not significantly different.

indirect population activity (GLM-id models). Unlike the GLM-2d and GLM-2i models, predictor units in GLM-id models tended to have large weights for small time lags (Fig. 4F), and simultaneous data ($\tau = 0$) was most commonly given the largest magnitude weight (Fig. 4G,H). In contrast to the GLM-2d and GLM-2i models, the distribution of time lags with the largest magnitude weights was significantly nonuniform (bootstrap, $p < 0.00001$). Additionally, the GLM-id models had higher crossvalidated R^2 values than the GLM-2d models (Fig. 4B; permutation test, $p = 0.0079$). Together, this shows that there is a meaningful difference between M2 and M1 indirect units' respective functional relationships with M1 direct units. Specifically, M1 indirect units tend to be more correlated with M1 BMI-potent activity than M2 units, and at shorter time lags; reflecting their belonging to the same local population.

M2 population activity evolves more slowly than M1 indirect activity

Finally, we examined the timescales at which M2 and M1 indirect populations activities evolved immediately before successful

trial completion. We observed that M2 population activity included broad modulation near the end of successful trials (Fig. 2B,C), and hypothesized that M2 population activity was changing at a slower timescale than M1 indirect population activity. We tested this hypothesis in two ways. First, we computed single unit autocorrelation functions, and found that the average width-at-half-max was significantly wider for M2 units than for M1 indirect units (permutation test, $p < 0.00001$). Second, we binned neural data at 50 ms and correlated M2 or M1i population vectors with their respective population vectors during the final 2 s of correct trials. We found that the mean correlation coefficient was significantly higher for M2 than M1i (Fig. 5A–D; permutation test, $p < 0.00001$), supporting a model in which M2 population activity patterns evolve at a longer timescale than M1 population activity patterns.

Discussion

In this study, we found differential relationships between M1 direct units and M1 indirect versus M2 units during M1 BMI

control. While M2 population activity changed slowly and was predictive of BMI performance at a broad range of time lags, M1 indirect population activity changed quickly and was predictive of performance at short time lags. These findings suggest a model in which M2 plays a “hierarchical role” by providing a longer timescale influence on the short-timescale M1 activity, which is closer to neuroprosthetic control. Additionally, M2-based models predicting M1 direct or M1 indirect activity had similar correlation values, suggesting M2 interacts with M1 neurons as a population, rather than through privileged relationships with effector-specific neurons. In contrast, models predicting M1 direct unit activity from M1 indirect unit activity had higher correlation values, indicating closer within-area interactions. Together, these findings suggest a model of M2-M1 communication in which M2 provides a slow modulatory input. Thus, moment-to-moment BMI control might be largely internal to M1, while presumed contextual or other top-down input might be provided by M2 (Siniscalchi et al., 2016; Omlor et al., 2019). Understanding these differential roles provides insight into natural movement control and can be harnessed to improve BMI functionality. For example, BMI paradigms could incorporate longer-timescale contextual signals for situation-specific controllers.

Multiarea BMIs to study motor control

BMI experiments allow selection of target neural patterns in one region and concurrent examination of correlates of task performance in connected regions. We thus aimed to isolate the top-down influences from M2 to M1. Notably, both M1 and M2 have direct connections to the spinal cord (Nudo and Masterton, 1990; Dum and Strick, 2002), they are implicated in normal movement control (Harrison et al., 2012; Wang et al., 2017), and can compensate after damage to either area (Whishaw et al., 1991; Ramanathan et al., 2018). In our BMI study, our arbitrarily selected target neural pattern, enforced by the decoder, is unlikely to be correlated with processes in other brain regions. This allows us to examine how cross-area communication during BMI control facilitates control. Specifically, natural movement studies suggest that activity in M2 may be context specific (Murakami et al., 2014; Saiki et al., 2014; Siniscalchi et al., 2016; Omlor et al., 2019) and appears to be organized in a top-down manner. Thus, while it

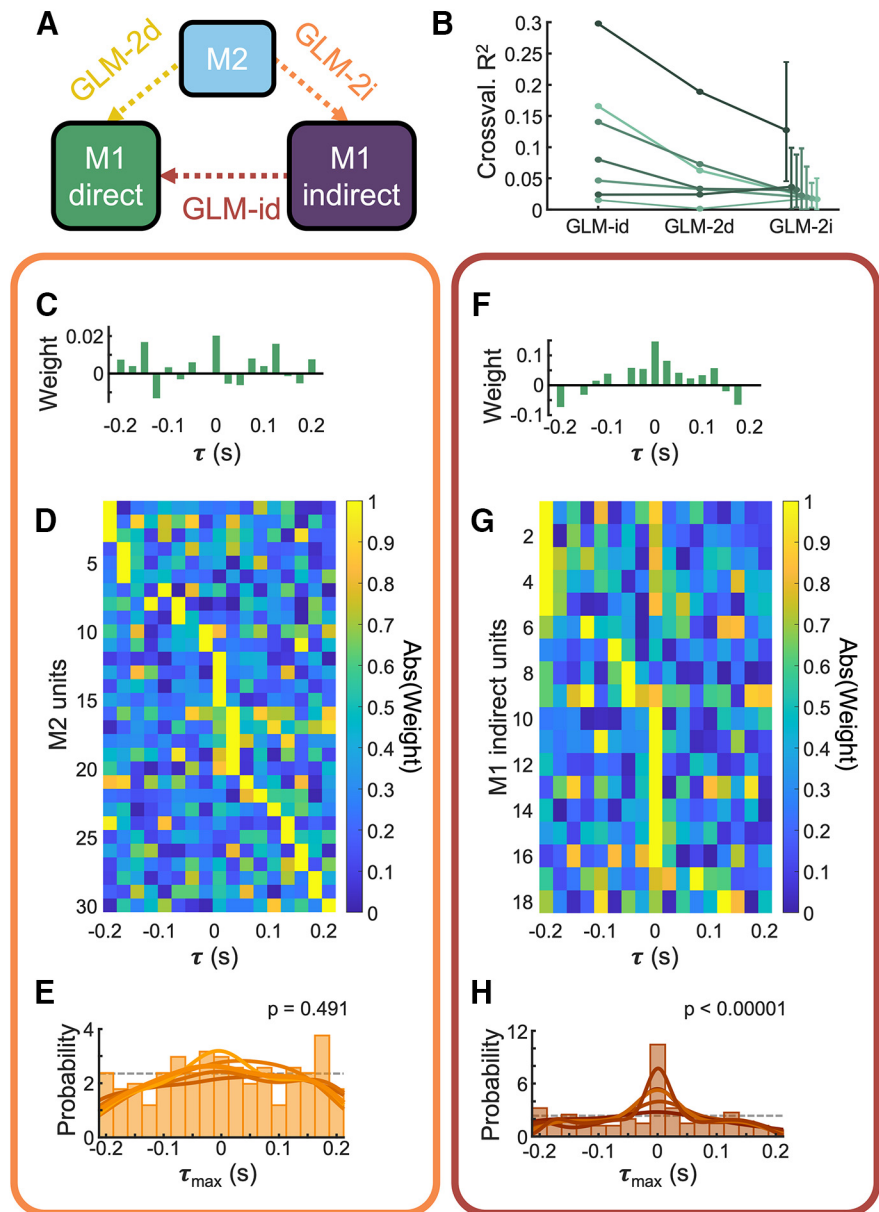


Figure 4. Comparison of time lags of M2 and M1-indirect units' predictions of M1-direct units. **A**, Models which predict M1 direct activity from M2 activity (GLM-2d; see Fig. 3) were compared with models which predict M1 direct activity from M1 indirect activity (GLM-id) and to models which predict a subset of M1 indirect neurons from M2 activity (GLM-2i). **B**, Crossvalidated R^2 values for GLM-id, GLM-2d, and GLM-2i models. Each line is one dataset. For the GLM-2i models, 50 different models were fitted to different choices of indirect units. Circle shows mean, error bars show full range. **C**, Example regression weights for an M2 neuron from one example learning session, for a model fitted to neural data binned at 25 ms. Height of the green lines indicates regression weights for the example neurons at different time lags (τ), with negative values meaning that M2 leads. **D**, Distribution of regression weight magnitudes in one example learning session. For each unit, regression weights were assigned for a variety of time lags (as in **C**), and the absolute value of those weights is indicated by color. Units are sorted according to the time of the largest magnitude weight. **E**, Histogram of time lags that had the largest magnitude regression weight for all M2 units in all animals, for models fitted to neural data binned at 25 ms. Solid lines show kernel density fit to equivalent histograms for models fit to neural data binned at 15 (lightest), 25, 35, 50, and 100 ms (darkest). Dashed gray line shows theoretical uniform distribution. **F**, As in **C**, but for GLM-id models. **G**, As in **D**, but for GLM-id models. **H**, As in **E**, but for GLM-id models.

would be surprising if M2 modulated neurons were not functionally coupled to M1 neurons (Makino et al., 2017; Veuthey et al., 2020), it is unknown how M2 signals might interact with movement-specific versus nontask pools of M1 neurons.

For example, we find that the correlation between M2 and M1 direct neurons is weaker and at longer time lags than the correlation between M1 indirect and direct neurons. This suggests

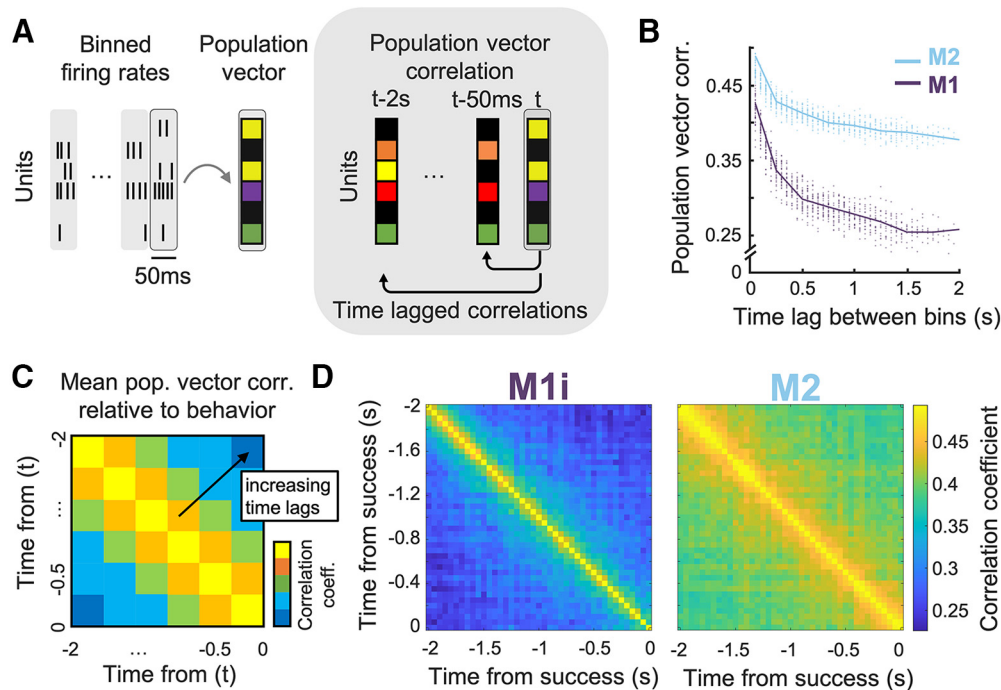


Figure 5. M2 population activity evolves more slowly across trials than M1 indirect population activity. **A**, left, Patterns of binned firing rates across a neural population were represented as population vectors, where the firing rate for each unit is an entry. Right, Population vectors are correlated across different time lags to quantify temporal stability. **B**, Population vector correlations for different time lags averaged across populations of M2 versus M1 indirect units. **C**, Explanatory graphic representing correlation values between population vectors from different time bins relative to behavior. **D**, Correlation of population firing vectors across time for neural data binned at 50 ms for (left) M1i indirect and (right) M2 neural populations. Color indicates the average correlation coefficient between population vectors at the given time for all successful trials of at least 2-s duration, for all animals. Mean M2 correlations coefficients were significantly higher than for M1i (M2: 0.41, M1: 0.31; permutation test, $p < 0.00001$, $n = 820$ values for each region).

that M2 provides inputs to M1 that are task-specific, albeit temporally imprecise; in contrast, internal M1 population dynamics are more robust but are temporally limited. Because we locked activity to the end of a successful trial, one interpretation is that M2 inputs are better positioned to steer to the target state by providing a steady modulatory input. This is analogous to the finding in sensory systems that task information in higher level regions is present at greater latencies and maintained for longer durations (Murray et al., 2014; Runyan et al., 2017).

Timescales of multiplexed local versus cross-area interactions

Little is known about how signals of different timescales interact during motor control (Cavanagh et al., 2020). Specifically, it is unclear how long-duration signals, perhaps representing abstract goals, are integrated with short-duration signals controlling movements. Studies with simultaneous recordings are best suited to address these questions of interactions (Perich et al., 2018; Lemke et al., 2019; Smedo et al., 2019; Stringer et al., 2019; Perich and Rajan, 2020). However, most studies with simultaneous recordings in motor regions use well-trained animals performing natural movement tasks (Kaufman et al., 2014; Perich et al., 2018; Musall et al., 2019; Stringer et al., 2019). When it comes to understanding M1-M2 interactions, this raises two potential complications: that M2 signals may be directly controlling movement rather than modulating M1, and that well-trained animals, who have often completed thousands of training trials, have transitioned to automatic movements with compressed or reduced top-down goals signals (Kawai et al., 2015). These confounds make it difficult to interpret the timescale of M1-M2 signal interactions in these experiments (Kaufman et al., 2014; Perich et al., 2018).

Here, we show that, in a M1 BMI task, M2 contains M1-predictive population signals at a broad range of time lags, in contrast to the short time lag predictive signals within M1. This supports two important concepts in motor control. First, that the activity of M1 effector neurons contains multiplexed signals coordinated with both local and cross-area activity. Second, that these multiplexed signals have timescales corresponding to their hierarchical functions (Murray et al., 2014; Runyan et al., 2017; Cavanagh et al., 2020). In turn, these findings reinforce the idea that variances in M1 task activity previously perceived as noise may instead be important activity related to coordination in the larger motor network (Smedo et al., 2019; Stringer et al., 2019; Veuthey et al., 2020). While the function of such large-network activity is still unknown, it has been hypothesized to represent internal attentional states (Ni et al., 2018), to provide modulatory motivational signals (Decot et al., 2017), and to coordinate movement with sensory feedback (Stringer et al., 2019).

Dynamic versus static neural activity goals

BMIs are often framed as a substitute for natural movement, but there are important differences between BMI and natural movement tasks. For one, while natural movements involve body parts which must move continuously through space, BMIs do not have this restriction. Studies also suggest that, within motor cortex, neural activity during natural movement progresses through continuous states (Churchland et al., 2012; Shenoy et al., 2013). However, high-dimensional neural activity is not inherently constrained by continuous paths through activity space. Instead, neural activity can jump between states. This has been exploited in nonhuman primate and human experiments in which participants use a BMI to type on a digital keyboard (Stavisky et al., 2015) rather than trace paths through virtual space, as is more

common with artificial limb research. In the simple BMI used in our experiments, BMI control occurred over discrete time bins and output was only smoothed over short timescales. Neural activity was therefore less constrained to be continuous than for BMIs mimicking movements (Liu and Schieber, 2020). Furthermore, the target neural activity pattern was a static threshold, and the approach to it was unconstrained within the BMI-potent space. Because of this, it is particularly interesting that population activity in M2 evolved relatively slowly, since the task did not require a smooth evolution even in downstream M1. However, to better understand natural movement dynamics, experimenters may want to use BMIs with dynamic rather than static neural activity targets (Collinger et al., 2013; Athalye et al., 2017).

This does not mean threshold-based BMI tasks are irrelevant for understanding natural movements. Instead, threshold-based BMI tasks may best be compared with reaction time and waiting time tasks. In waiting time tasks, animals must hold a position until an arbitrary internal threshold is reached (Murakami et al., 2014), similar to common simple BMI paradigms in which the neural activity must be maintained above a target threshold for a period of time. Waiting experiments have demonstrated encoding of hold time within medial prefrontal cortex (mPFC; Xu et al., 2014) and M2 (Murakami et al., 2014) single neurons, with single neurons displaying a variety of activity patterns, including a slow ramp to threshold. This may contribute to the slow evolution of M2 population activity seen here.

Activity of local and cross-area indirect neurons over BMI learning

Natural motor learning is considered to occur in several phases, from early exploration to consolidated control (Hikosaka et al., 2002; Kleim et al., 2004; Floyer-Lea and Matthews, 2005; Lutz and Buitrago, 2005). We focus on early exploratory BMI learning by using a single-session paradigm. BMI studies which allow sleep consolidation (Gulati et al., 2014, 2017; Kim et al., 2019) or use multiday learning paradigms (Ganguly et al., 2011; Silversmith et al., 2021) have found that M1 indirect units become less modulated. It is possible that at later stages, M2 task-modulation is similarly decreased; however, it is also possible that consolidation has a different effect on cross-area indirect neurons. Only one other study has analyzed the activity of cross-area indirect neurons in cortex; they found robust task-related modulation in several cortical areas, including premotor cortex, after multiday learning (Liu and Schieber, 2020). This suggests that M2 activity patterns may be maintained. This parallels evidence from studies of natural learning that overall correlated activity between M2-M1 does not change with learning (Veuthey et al., 2020).

In conclusion, this project leverages a multiarea BMI paradigm to probe cross-area interactions. We show that task-related local interactions are more robust and take place at shorter time lags than cross-area interactions. While our paradigm imposes an artificial target activity pattern within M1, its relatively small neural footprint allows us to examine internal motor network dynamics without the constraints of movement performance. Thus, multiarea BMIs may allow for a more natural inside-out view of network dynamics.

References

- Arduin PJ, Frégnac Y, Shulz DE, Ego-Stengel V (2013) “Master” neurons induced by operant conditioning in rat motor cortex during a brain-machine interface task. *J Neurosci* 33:8308–8320.
- Athalye VR, Ganguly K, Costa RM, Carmena JM (2017) Emergence of coordinated neural dynamics underlies neuroprosthetic learning and skillful control. *Neuron* 93:955–970.e5.
- Athalye VR, Santos FJ, Carmena JM, Costa RM (2018) Evidence for a neural law of effect. *Science* 359:1024–1029.
- Cao VY, Ye Y, Mastwal S, Ren M, Coon M, Liu Q, Costa RM, Wang KH (2015) Motor learning consolidates Arc-expressing neuronal ensembles in secondary motor cortex. *Neuron* 86:1385–1392.
- Cavanagh SE, Hunt LT, Kennerley SW (2020) A diversity of intrinsic timescales underlie neural computations. *Front Neural Circuits* 14:615626.
- Chung JE, Magland JF, Barnett AH, Tolosa VM, Tooker AC, Lee KY, Shah KG, Felix SH, Frank LM, Greengard LF (2017) A fully automated approach to spike sorting. *Neuron* 95:1381–1394.e6.
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. *Nature* 487:51–56.
- Clancy KB, Koralek AC, Costa RM, Feldman DE, Carmena JM (2014) Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat Neurosci* 17:807–809.
- Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJ, Velliste M, Boninger ML, Schwartz AB (2013) High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet* 381:557–564.
- Decot HK, Namboodiri VMK, Gao W, McHenry JA, Jennings JH, Lee S-H, Kantak PA, Jill Kao YC, Das M, Witten IB, Deisseroth K, Shih YYI, Stuber GD (2017) Coordination of brain-wide activity dynamics by dopaminergic neurons. *Neuropsychopharmacology* 42:615–627.
- Dum RP, Strick PL (2002) Motor areas in the frontal lobe of the primate. *Physiol Behav* 77:677–682.
- Fetz EE (2007) Volitional control of neural activity: implications for brain-computer interfaces. *J Physiol* 579:571–579.
- Floyer-Lea A, Matthews PM (2005) Distinguishable brain activation networks for short- and long-term motor skill learning. *J Neurophysiol* 94:512–518.
- Ganguly K, Carmena JM (2009) Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol* 7:e1000153.
- Ganguly K, Dimitrov DF, Wallis JD, Carmena JM (2011) Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat Neurosci* 14:662–667.
- Gulati T, Ramanathan DS, Wong CC, Ganguly K (2014) Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nat Neurosci* 17:1107–1113.
- Gulati T, Guo L, Ramanathan DS, Bodepudi A, Ganguly K (2017) Neural reactivations during sleep determine network credit assignment. *Nat Neurosci* 20:1277–1284.
- Harrison TC, Ayling OGS, Murphy TH (2012) Distinct cortical circuit mechanisms for complex forelimb movement and motor map topography. *Neuron* 74:397–409.
- Hikosaka O, Nakamura K, Sakai K, Nakahara H (2002) Central mechanisms of motor skill learning. *Curr Opin Neurobiol* 12:217–222.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci* 17:440–448.
- Kawai R, Markman T, Poddar R, Ko R, Fantana AL, Dhawale AK, Kampff AR, Ölveczky BP (2015) Motor cortex is required for learning but not for executing a motor skill. *Neuron* 86:800–812.
- Kim J, Gulati T, Ganguly K (2019) Competing roles of slow oscillations and delta waves in memory consolidation versus forgetting. *Cell* 179:514–526.e13.
- Kleim JA, Hogg TM, VandenBerg PM, Cooper NR, Bruneau R, Remple M (2004) Cortical synaptogenesis and motor map reorganization occur during late, but not early, phase of motor skill learning. *J Neurosci* 24:628–633.
- Koralek AC, Jin X, Long JD, Costa RM, Carmena JM (2012) Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* 483:331–335.
- Koralek AC, Costa RM, Carmena JM (2013) Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron* 79:865–872.
- Lemke SM, Ramanathan DS, Guo L, Won SJ, Ganguly K (2019) Emergent modular neural control drives coordinated motor actions. *Nat Neurosci* 22:1122–1131.

- Liu Z, Schieber MH (2020) Neuronal activity distributed in multiple cortical areas during voluntary control of the native arm or a brain-computer interface. *eNeuro* 7:ENEURO.0376-20.2020.
- Luft AR, Buitrago MM (2005) Stages of motor skill learning. *Mol Neurobiol* 32:205–216.
- Makino H, Ren C, Liu H, Kim AN, Kondapaneni N, Liu X, Kuzum D, Komiyama T (2017) Transformation of cortex-wide emergent properties during motor learning. *Neuron* 94:880–890.e8.
- Murakami M, Vicente MI, Costa GM, Mainen ZF (2014) Neural antecedents of self-initiated actions in secondary motor cortex. *Nat Neurosci* 17:1574–1582.
- Murray JD, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Padoa-Schioppa C, Pasternak T, Seo H, Lee D, Wang XJ (2014) A hierarchy of intrinsic timescales across primate cortex. *Nat Neurosci* 17:1661–1663.
- Musall S, Kaufman MT, Juvinett AL, Gluf S, Churchland AK (2019) Single-trial neural dynamics are dominated by richly varied movements. *Nat Neurosci* 22:1677–1686.
- Neafsey EJ, Bold EL, Haas G, Hurley-Gius KM, Quirk G, Sievert CF, Terrence RR (1986) The organization of the rat motor cortex: a microstimulation mapping study. *Brain Res Reviews* 11:77–96.
- Ni AM, Ruff DA, Alberts JJ, Symmonds J, Cohen MR (2018) Learning and attention reveal a general relationship between population activity and behavior. *Science* 359:463–465.
- Nudo RJ, Masterton RB (1990) Descending pathways to the spinal cord, III: sites of origin of the corticospinal tract. *J Comp Neurol* 296:559–583.
- Omlor W, Wahl A-S, Sipilä P, Lütcke H, Laurency B, Chen I-W, Sumanovski LT, van 't Hoff M, Bethge P, Voigt FF, Schwab ME, Helmchen F (2019) Context-dependent limb movement encoding in neuronal populations of motor cortex. *Nat Commun* 10:4812.
- Perich MG, Rajan K (2020) Rethinking brain-wide interactions through multi-region ‘network of networks’ models. *Curr Opin Neurobiol* 65:146–151.
- Perich MG, Gallego JA, Miller LE (2018) A neural population mechanism for rapid learning. *Neuron* 100:964–976.e7.
- Peters AJ, Chen SX, Komiyama T (2014) Emergence of reproducible spatio-temporal activity during motor learning. *Nature* 510:263–267.
- Ramanathan DS, Gulati T, Ganguly K (2015) Sleep-dependent reactivation of ensembles in motor cortex promotes skill consolidation. *PLoS Biol* 13:e1002263.
- Ramanathan DS, Guo L, Gulati T, Davidson G, Hishinuma AK, Won SJ, Knight RT, Chang EF, Swanson RA, Ganguly K (2018) Low-frequency cortical activity is a neuromodulatory target that tracks recovery after stroke. *Nat Med* 24:1257–1267.
- Runyan CA, Piasini E, Panzeri S, Harvey CD (2017) Distinct timescales of population coding across cortex. *Nature* 548:92–96.
- Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP (2014) Neural constraints on learning. *Nature* 512:423–426.
- Saiki A, Kimura R, Samura T, Fujiwara-Tsukamoto Y, Sakai Y, Isomura Y (2014) Different modulation of common motor information in rat primary and secondary motor cortices. *PLoS ONE* 9:e98662.
- Semedo JD, Zandvakili A, Machens CK, Yu BM, Kohn A (2019) Cortical areas interact through a communication subspace. *Neuron* 102:249–259.
- Shenoy KV, Sahani M, Churchland MM (2013) Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci* 36:337–359.
- Silversmith DB, Abiri R, Hardy NF, Natraj N, Tu-Chan A, Chang EF, Ganguly K (2021) Plug-and-play control of a brain-computer interface through neural map stabilization. *Nat Biotechnol* 39:326–310.
- Siniscalchi MJ, Phoumthippavong V, Ali F, Lozano M, Kwan AC (2016) Fast and slow transitions in frontal ensemble activity during flexible sensorimotor behavior. *Nat Neurosci* 19:1234–1242.
- Sporns O, Chialvo D, Kaiser M, Hilgetag C (2004) Organization, development and function of complex brain networks. *Trends Cogn Sci* 8:418–425.
- Stavisky SD, Kao JC, Sorokin JM, Ryu SI, Shenoy KV (2015) System identification of brain-machine interface control using a cursor jump perturbation. 2015 7th International IEEE/EMBS Conference on Neural Engineering (NER), pp 643–647. Available at <https://ieeexplore.ieee.org/abstract/document/7146705>.
- Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD (2019) Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364:eaav7893.
- Taylor DM, Tillery SIH, Schwartz AB (2002) Direct cortical control of 3D neuroprosthetic devices. *Science* 296:1829–1832.
- Veuthey TL, Derosier K, Kondapavulur S, Ganguly K (2020) Single-trial cross-area neural population dynamics during long-term skill learning. *Nat Commun* 11:4057.
- Wang X, Liu Y, Li X, Zhang Z, Yang H, Zhang Y, Williams PR, Alwahab NSA, Kapur K, Yu B, Zhang Y, Chen M, Ding H, Gerfen CR, Wang KH, He Z (2017) Deconstruction of corticospinal circuits for goal-directed motor skills. *Cell* 171:440–455.e14.
- Whishaw IQ, Pellis SM (1990) The structure of skilled forelimb reaching in the rat: a proximally driven movement with a single distal rotatory component. *Behav Brain Res* 41:49–59.
- Whishaw IQ, Pellis SM, Gorny BP, Pellis VC (1991) The impairments in reaching and the movements of compensation in rats with motor cortex lesions: an endpoint, videorecording, and movement notation analysis. *Behav Brain Res* 42:77–91.
- Xu M, Zhang S, Dan Y, Poo M (2014) Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. *Proc Natl Acad Sci USA* 111:480–485.