



OPEN

Distinct neural representations during a brain–machine interface and manual reaching task in motor cortex, prefrontal cortex, and striatum

Ellen L. Zippi^{1,4}, Gabrielle F. Shvartsman^{2,4}, Nuria Vendrell-Llopis^{1,2}, Joni D. Wallis^{1,3} & Jose M. Carmena^{1,2}✉

Although brain–machine interfaces (BMIs) are directly controlled by the modulation of a select local population of neurons, distributed networks consisting of cortical and subcortical areas have been implicated in learning and maintaining control. Previous work in rodents has demonstrated the involvement of the striatum in BMI learning. However, the prefrontal cortex has been largely ignored when studying motor BMI control despite its role in action planning, action selection, and learning abstract tasks. Here, we compare local field potentials simultaneously recorded from primary motor cortex (M1), dorsolateral prefrontal cortex (DLPFC), and the caudate nucleus of the striatum (Cd) while nonhuman primates perform a two-dimensional, self-initiated, center-out task under BMI control and manual control. Our results demonstrate the presence of distinct neural representations for BMI and manual control in M1, DLPFC, and Cd. We find that neural activity from DLPFC and M1 best distinguishes control types at the go cue and target acquisition, respectively, while M1 best predicts target-direction at both task events. We also find effective connectivity from DLPFC → M1 throughout both control types and Cd → M1 during BMI control. These results suggest distributed network activity between M1, DLPFC, and Cd during BMI control that is similar yet distinct from manual control.

The ability to volitionally modulate the activity of neurons has been demonstrated in both human subjects and animals¹. This skill is fundamental to the operation of brain–machine interfaces (BMIs), in which external devices, such as computer cursors or robotic arms, are controlled by the production of specific patterns of neural activity. Motor BMIs specifically rely on the modulation of motor cortical neural activity to control external devices^{2–4}, while also relying on higher-level cognitive processes to integrate sensory information, plan and initiate actions, and monitor visual feedback of the effector⁵. These systems hold great potential for rehabilitation and restoration of lost motor function^{4,6}, but the underlying neural processes involved in controlling a motor BMI are still unclear.

Controlling a BMI shares characteristics with both overt motor control and abstract cognitive tasks^{7–9}. BMI control requires users to produce specific patterns of neural activity to manipulate their physical environment without physical movement and often without proprioceptive feedback. Thus, BMI control may rely on networks underlying both overt motor control and abstract cognitive tasks. While both BMI control and overt motor control rely on the direct modulation of the motor cortex, the mapping between neural activity and output is distinct between these types of control because a specific subpopulation of neurons is selected for BMI control. The prefrontal cortex and dorsal striatum are important for both motor control^{10–13} and for learning abstract associations^{14–17}.

The primary motor cortex (M1), the dorsolateral prefrontal cortex (DLPFC) and the caudate nucleus of the striatum (Cd) are extensively interconnected with one another, as well as with other sensory and higher-level associational areas involved in goal-directed behavior^{18–22}. Cortico-striatal plasticity is necessary for motor

¹Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, CA, USA. ²Department of Electrical Engineering and Computer Science, University of California, Berkeley, Berkeley, CA, USA. ³Department of Psychology, University of California, Berkeley, Berkeley, CA, USA. ⁴These authors contributed equally: Ellen L. Zippi and Gabrielle F. Shvartsman. ✉email: jcarmena@berkeley.edu

control²³ and cortico-striatal and cortico-cortical interactions are involved in abstract skill learning^{18,24,25}, suggesting that these interactions may play an important role in BMI control.

Successful control of a BMI depends on multiple cognitive processes that are employed differentially across events within the task. For example, in a center-out BMI task, the go cue may be more dependent on action selection and initiation, while control at target acquisition may depend more on error correction and fine movement control. Therefore, it is possible that different regions of interest (ROIs), such as motor cortex, prefrontal cortex, or striatum, are more or less involved at different points within a trial. For instance, the prefrontal cortex may be more involved at the go cue when action planning and selection are critical^{13,18,26} whereas the primary motor cortex may be more involved at target acquisition when movement execution and fine effector control are more important²⁷. Thus, it is important to investigate neural activity in each ROI at multiple task events to better understand its role during manual and BMI control.

Previous studies that have used BMIs to dissect cognitive processes have primarily focused on neural activity in BMI-selected neurons, motor cortical neurons whose activity is directly input to the BMI²⁸. A few studies have investigated the role of neural activity from non-BMI neurons within the motor cortex^{29,30}, highlighting that BMI control involves large-scale networks that extend beyond BMI-selected neurons. However, it is becoming apparent that extending these investigations to brain areas beyond the motor cortex will be essential for gaining a more complete understanding of BMI control³¹. Previous studies have shown the involvement of cortico-striatal networks during a BMI task in rodents^{32–34} and of the prefrontal cortex in BMI skill acquisition in humans^{8,35,36}. However, due to experimental constraints, these studies relied on simplified, one-dimensional BMI tasks, making it difficult to gain insight into the distributed network activity that may be specific to more complex and cognitively demanding two-dimensional (2-D) BMI control. To overcome these limitations, we successfully trained non-human primates (NHPs) to perform a 2-D BMI center-out task with recordings from a semi-chronic microdrive. These microdrives have previously been used to study large-scale networks^{37,38}, but never in the context of BMI control. To our knowledge, this is the first time recordings from a semi-chronic microdrive have been used to drive a BMI in NHPs, allowing us to obtain simultaneous cortical and subcortical recordings to make direct comparisons of activity in these regions and investigate their interactions during BMI control.

Here, we use this unique experimental opportunity to investigate the modulation of and interactions between M1, DLPFC, and Cd activity during a motor BMI task and compare this activity to that during manual (overt motor) control. We perform these analyses at both the go cue and at target acquisition allowing us to compare the role of each ROI across two task events that involve distinct cognitive processes. We demonstrate that DLPFC holds the most information for distinguishing BMI and manual control at the go cue, while M1 is most informative at target acquisition. Additionally, we show that directed information flow is present from DLPFC → M1 in both BMI and manual control and from Cd → M1 during BMI control. These directed interactions were present at both the go cue and at target acquisition. Ultimately, our findings confirm the involvement of distributed cortical and subcortical networks during BMI control and identify similar but distinct network activity during manual control.

Results

To explore whether there are BMI-specific neural representations in M1, DLPFC, and Cd, we simultaneously recorded from these regions during a BMI control task, a manual control task, and a baseline rest period. Two rhesus macaques (Monkey H and Monkey Y) were implanted with a custom-fit large-scale semi-chronic microdrive array on the left hemisphere (Fig. 1a). Single- and multi-unit recordings from the motor cortex were used as input to a BMI decoder, while local-field potentials (LFP) were simultaneously recorded from three regions of interest (ROIs): M1, DLPFC, and Cd. Each day, the animals performed a two-dimensional, self-initiated, center-out task, in which they were instructed to move a cursor from a center target to one of eight pseudo-randomly instructed peripheral targets for a juice reward. A successful center-out trial required a brief hold at the center target, moving to the peripheral target within a specified time, and a brief hold at the target (Fig. 1b). First, they performed the task under manual control (Fig. 1c), followed by a four-minute baseline period (Fig. 1d). Then, they performed the same center-out task under BMI control (Fig. 1e,f).

Successful BMI control with a semi-chronic array

Both monkeys successfully learned to perform the eight-target, center-out task using a cursor under BMI control. Task performance improved across days, despite using different direct units and decoders each day (see “Methods” for details). Although only one monkey significantly increased the fraction of successfully completed trials (Fig. 1g; Linear regression, Monkey H: $R^2 = 0.003$, $p = 0.852$, Monkey Y: $R^2 = 0.593$, $p = 2.78e-5$), both monkeys learned to produce faster, straighter cursor trajectories, resulting in a decreased target acquisition time over days (Fig. 1h; Linear regression, Monkey H: $R^2 = 0.035$, $p = 3.20e-27$, Monkey Y: $R^2 = 0.093$, $p = 4.67e-141$). This improved performance demonstrates that the animals were able to learn to control a BMI using neural activity recorded from a semi-chronic microdrive.

Spectral power from M1, DLPFC, and Cd distinguishes BMI control, manual control, and baseline

To understand whether activity from M1, DLPFC, and Cd differs between BMI control, manual control, and baseline, we compared the spectral power in 5 distinct frequency bands across these three conditions (Supplementary Figs. 1, 2, see “Methods” for details). Using the LFP recorded from each ROI, we computed the mean power across electrode channels in each frequency band. For BMI and manual control, power was computed in the 500 ms window following the go cue and the 500 ms window preceding target acquisition on successful trials. The same calculations were performed in random non-overlapping 500 ms windows during the baseline

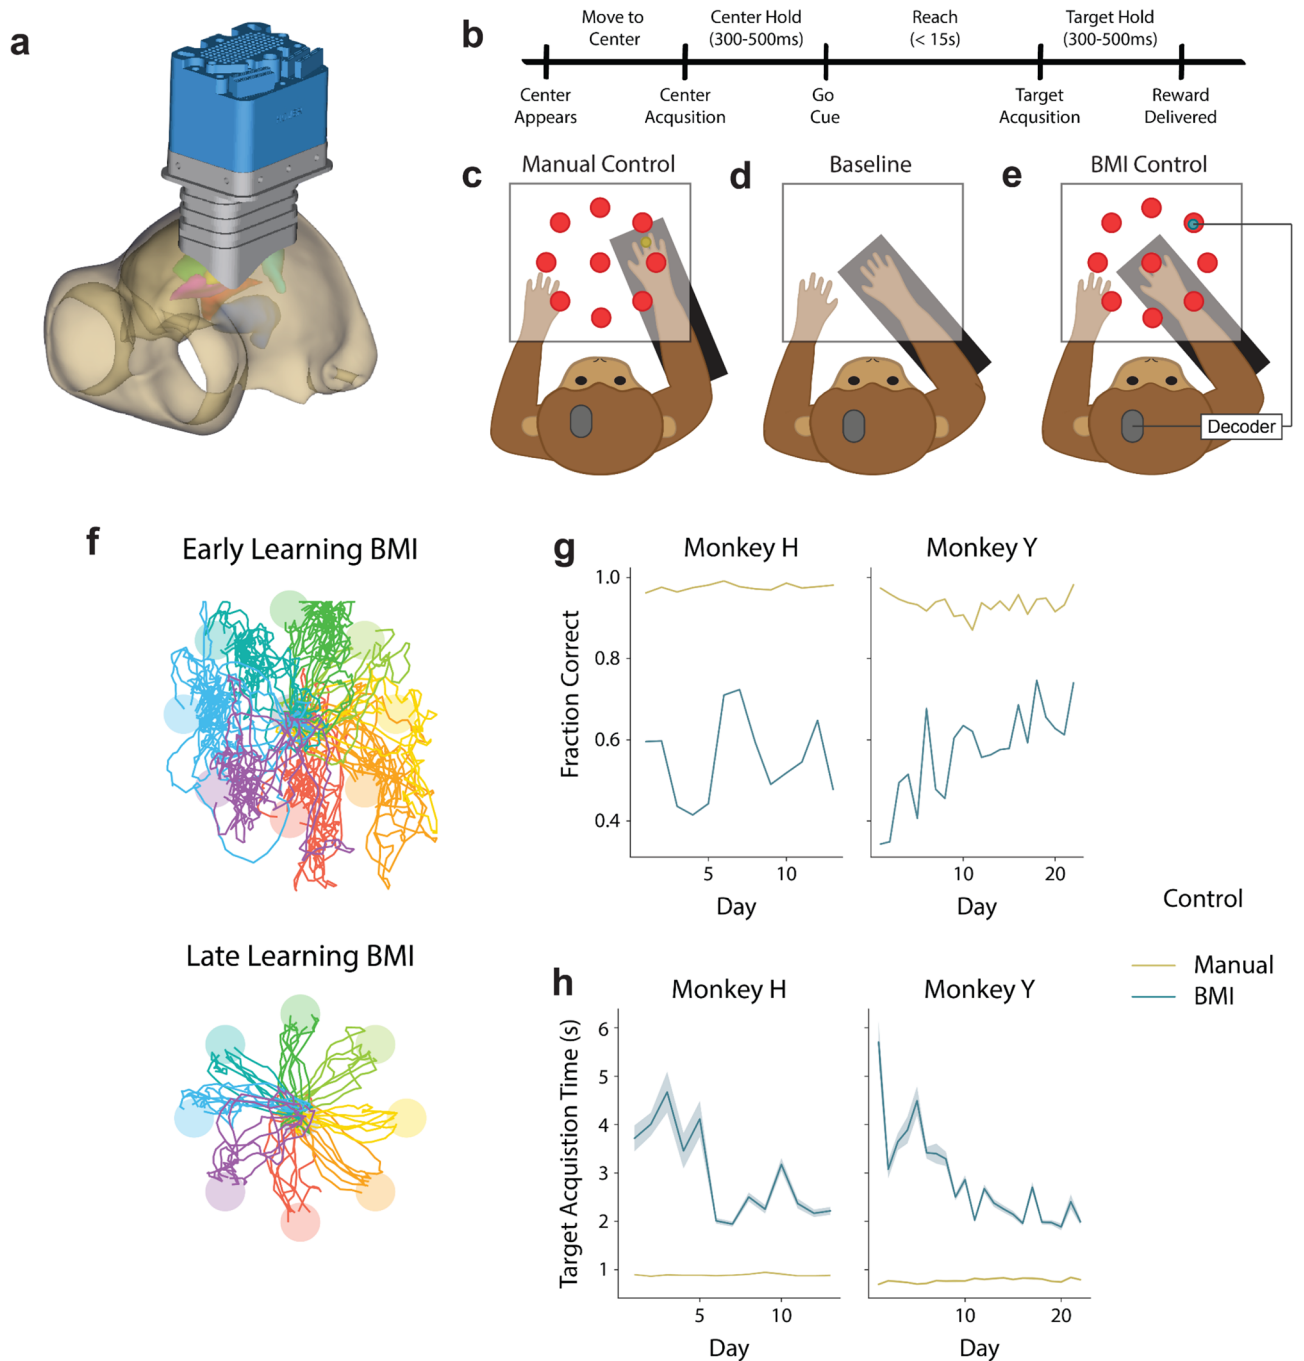


Figure 1. Experimental set-up and task performance. **(a)** A 3-D model of a custom-fit, large-scale, semi-chronic microdrive from Gray Matter Research used for simultaneous recordings of neural activity from ROIs at different depths. **(b)** Timeline of the center-out task (see “Methods” for details). **(c)** Two monkeys completed a 2-D, self-initiated, center-out movement task under manual control. **(d)** During a baseline period, neural activity was recorded during the absence of visual cues and the subject’s arm locked in a fixed position. **(e)** During BMI control, the subject’s arm was locked in a fixed position while they modulated neural activity to control a cursor in the same 2-D, self-initiated, center-out task. **(f)** Example BMI control trajectories from early learning (right) and late learning (left) to each target from Monkey Y. **(g)** Fraction of self-initiated trials that were successful in BMI control (teal) and manual control (yellow) for Monkey H (left) and Monkey Y (right). **(h)** Time from go cue to target acquisition in BMI control (yellow) and manual control (teal) for Monkey H (left) and Monkey Y (right). Shading represents the standard error of the mean across trials within a day.

period. The most notable differences occurred between baseline and the two control tasks in the theta and beta frequency bands (Fig. 2). At the go cue, theta power at baseline was significantly lower than during both manual

and BMI control across all ROIs, while beta power at baseline was significantly higher than during either of the two control tasks (Fig. 2a). At target acquisition, beta power at baseline remained significantly higher than during the two control tasks in M1 but was less consistent at other ROIs and frequency bands (Fig. 2b). While these differences provide insight into how spectral power differs between the baseline rest period and the two control tasks at different task events, differences in spectral power within individual frequency bands between the two types of control were less consistent across animals.

To determine whether combinations of spectral power from multiple frequency bands across our ROIs yield more distinct representations between task types, we trained quadratic discriminant analysis (QDA) classifiers to distinguish between these task types (Fig. 3a). Although only theta and beta power showed consistent significant differences across task types, incorporating power from all frequency bands as input features yielded a higher classification accuracy than including any individual feature for each individual ROI, so all power features were included for each ROI in subsequent analyses (Supplementary Fig. 3). We compared neural activity recorded during BMI control to activity recorded during both manual control and baseline, allowing us to identify task-related activity that is specific to BMI control (Fig. 3a). At both the go cue and target acquisition, classifiers using power features from all ROIs, M1 + DLPFC + Cd, either yielded the highest classification accuracy or were not significantly different from the highest accuracy (Fig. 3b,c, Supplementary Fig. 4a,b). At target acquisition, classifiers that incorporated M1 power features (M1 + DLPFC + Cd, M1 + DLPFC, M1 + Cd, and M1 only) yielded significantly higher classification accuracy than classifiers without M1 power features (DLPFC + Cd, DLPFC only, and Cd only) (Fig. 3c, Supplementary Fig. 4b). With BMI performance improving across days (Fig. 1g,h), we considered whether these results varied across different stages of BMI learning. We repeated these analyses within early (first third of recording days; Monkey H: $n = \text{first 4 days}$, Monkey Y: $n = \text{first 7 days}$) and late learning periods (last third of recording days; Monkey H: $n = \text{last 4 days}$, Monkey Y: $n = \text{last 7 days}$). Both the go cue and target acquisition results were similar across BMI learning stages (Supplementary Fig. 5). Overall, these results indicate that it is possible to distinguish between BMI control, manual control, and baseline using spectral power features from M1, DLPFC, and Cd. Furthermore, the increase in classification accuracy resulting from the inclusion of M1 activity indicates that there are distinct neural representations of the three task types in this region at target acquisition.

To further explore the differences in the neural representations of BMI control, manual control, and baseline in M1, DLPFC, and Cd, we computed confusion matrices for the classifiers trained using power features from all ROIs (Fig. 3d,e). Here, we assessed which type of misclassifications were most common. If the majority of misclassifications occur between BMI control and baseline, this would imply that these two classes are more similar to one another and distinct from manual control. Thus, we would infer that the differences in neural representations primarily correspond to the presence or absence of large, physical arm movements. On the other hand, a majority of misclassifications occurring between BMI control and manual control could indicate that differences in neural representations were largely task-related. Thus, quantifying the 3-class task-type QDA misclassifications allowed us to identify whether the model performed better at separating movement-related conditions or task-related conditions. At both the go cue and target acquisition, we found that the primary misclassifications occurred between BMI control and manual control, demonstrating that neural activity during the two tasks is more similar to one another than it is during baseline (Fig. 3d,e). This suggests that the primary

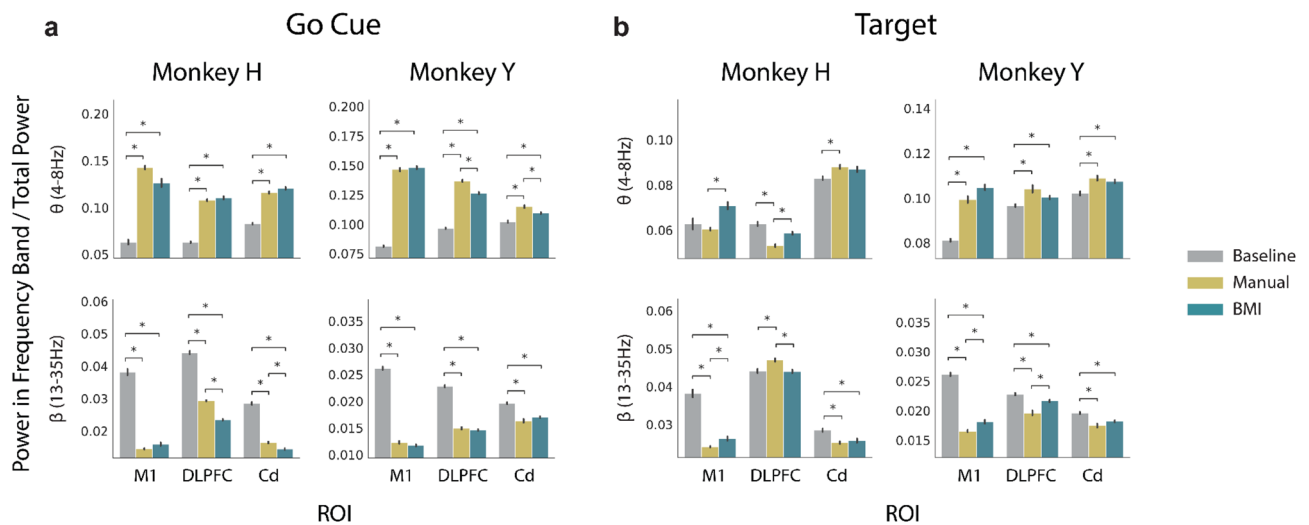


Figure 2. Theta and beta power in M1, DLPFC, and Cd is significantly different between baseline rest period and manual or BMI control. (a) Average theta (top) and beta (bottom) power normalized to total power in each ROI during baseline (gray), manual (yellow), and BMI (teal) at the go cue for Monkey H (left) and Monkey Y (right). Error bars represent standard error mean across days. (b) Same as (a), but for power features obtained at target acquisition rather than at the go cue. Normalized spectral power estimates significantly differing between task types after Bonferroni correction for multiple comparisons are indicated with an asterisk (Supplementary Tables 1, 2, Bonferroni corrected p -value < 0.05).

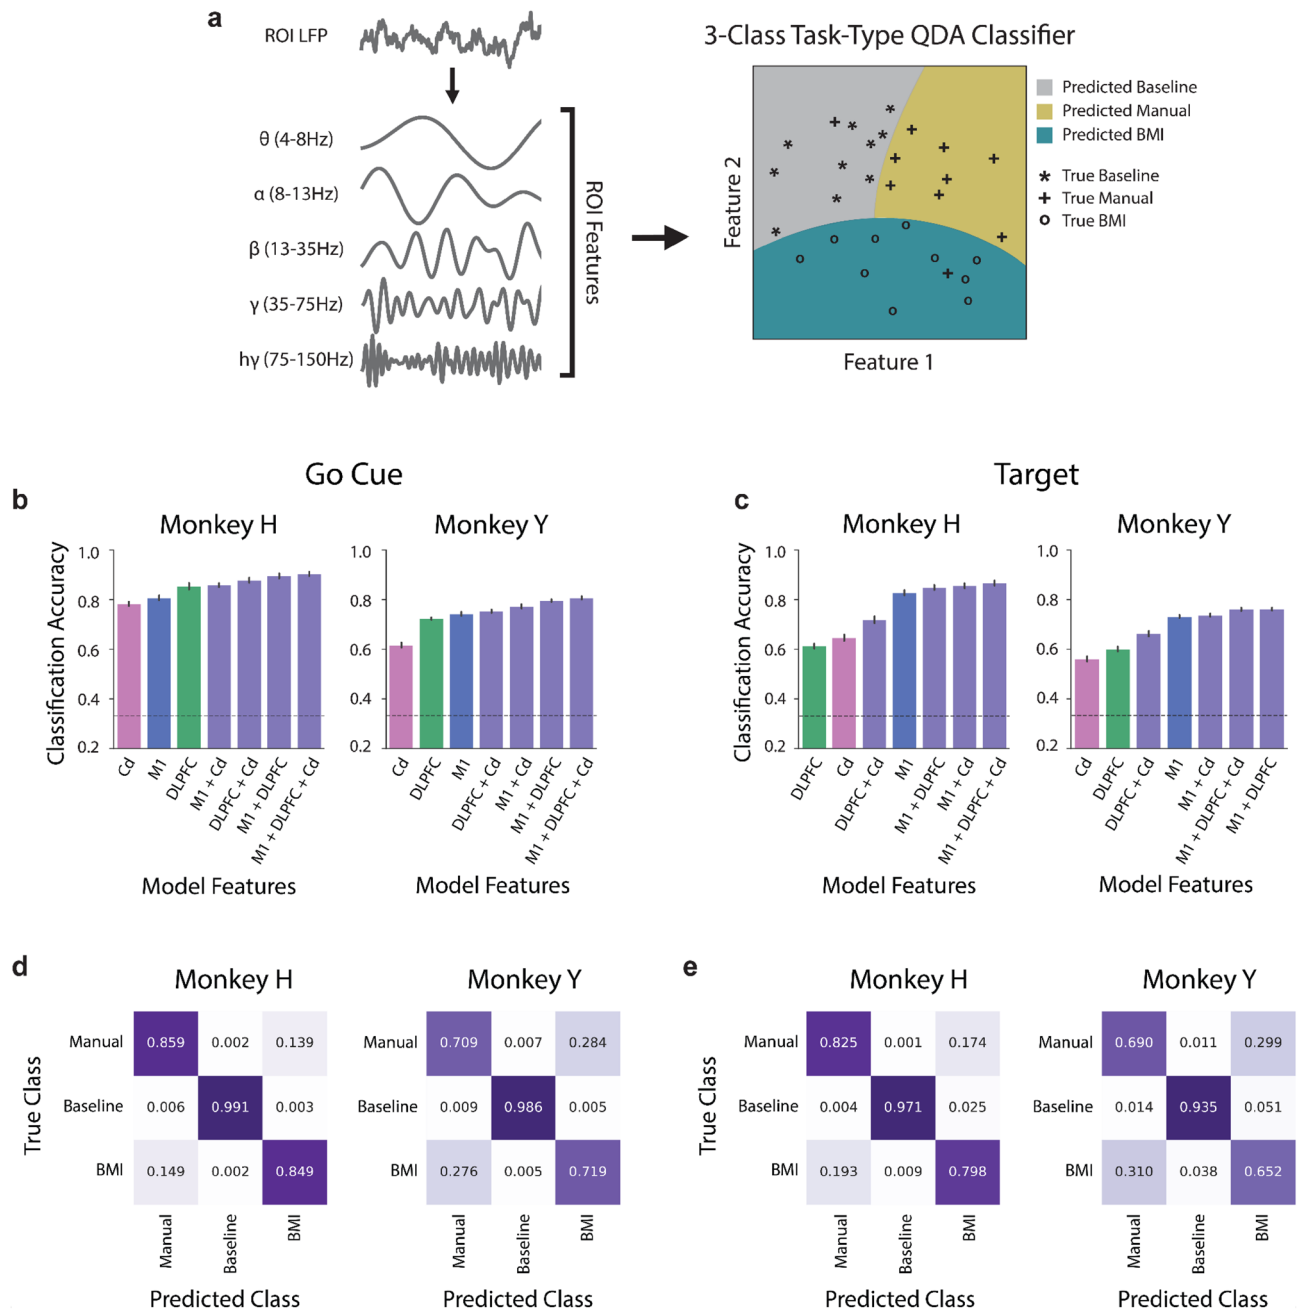


Figure 3. Spectral power in M1, DLPFC, and Cd distinguishes between BMI control, manual control, and baseline at the go cue and at target acquisition. **(a)** LFP was decomposed into 5 distinct frequency bands. The average power in each of these bands in each ROI was used as input to the classifiers. 3-class task-type QDA classifiers were trained to distinguish between BMI control, manual control, and baseline. **(b)** Mean tenfold cross-validated classification accuracy across days for a 3-class task-type QDA classifier using all frequency bands from individual ROIs (M1 in blue, DLPFC in green, Cd in pink) or combinations of ROIs (purple) at the go cue for Monkey H (left) and Monkey Y (right). Classifiers are presented in order of ascending accuracy. Chance accuracy shown as a dashed line. Error bars represent standard error mean across days. **(c)** Same as **(b)**, but for power features obtained at target acquisition rather than at the go cue. **(d)** Most misclassifications occurred between BMI control and manual control. Confusion matrix for the model including all frequency bands from all ROIs at the go cue. Numbers and shading correspond to the fraction of correct predictions within a class for Monkey H (left) and Monkey Y (right). **(e)** Same as **(d)**, but for power features obtained at target acquisition rather than at the go cue.

distinction in neural representations was due to the presence or absence of a task, rather than the presence or absence of movement.

DLPFC and M1 best distinguish between control-type at go cue and target acquisition, respectively

To gain further insight into the differences between neural representations of BMI control and manual control, we used a 2-class task-type QDA model to distinguish between BMI control and manual control only (Fig. 4a). At the go cue, using power features from DLPFC resulted in significantly higher classifier performance than using power features from M1 or Cd for Monkey H (Fig. 4b,c, Supplementary Fig. 6a). For Monkey Y, classifiers using power features from DLPFC had better performance than those using features from M1 or Cd, but there was only a significant improvement from Cd to DLPFC (Fig. 4b,c, Supplementary Fig. 6a). Additionally, the classification accuracy of models using power features from either DLPFC + Cd or M1 + DLPFC was not significantly different from that of models using power features from all ROIs for either subject. At target acquisition, models using

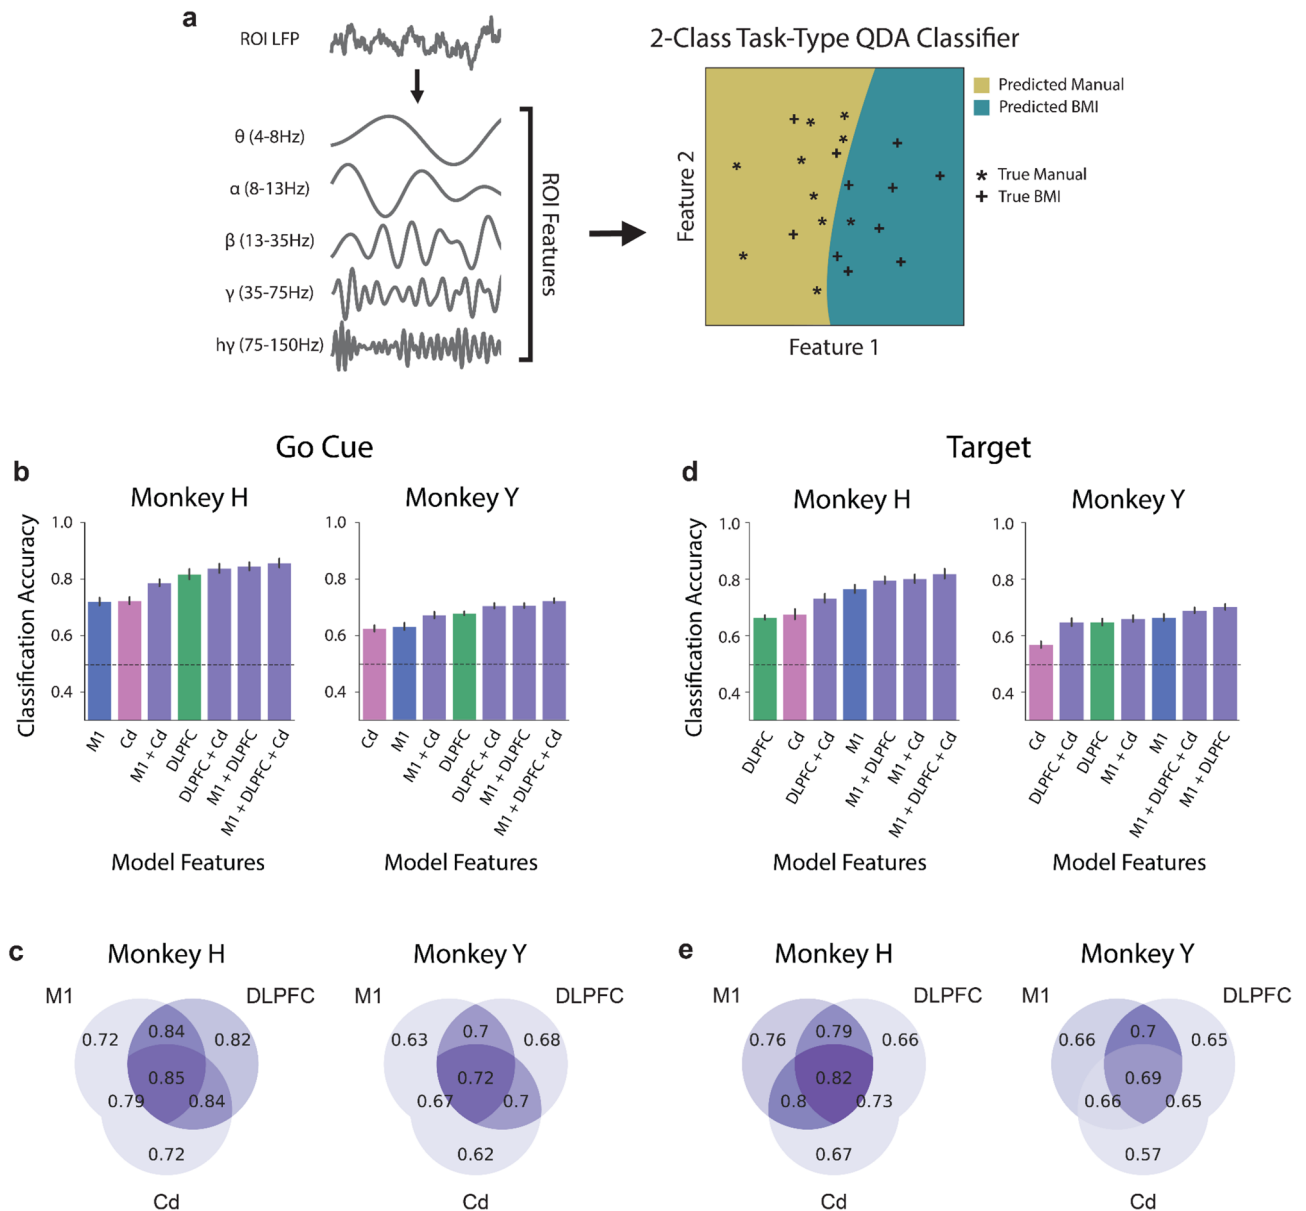


Figure 4. DLPFC activity best distinguishes between BMI and manual control at the go cue, while M1 activity best distinguishes between BMI and manual control at target acquisition. **(a)** LFP was decomposed into 5 distinct frequency bands. The average power in each of these bands in each ROI was used as input to the classifiers. 2-class task-type QDA classifiers were trained to distinguish between BMI control and manual control only. **(b)** Mean tenfold cross-validated classification accuracy across days for a 2-class task-type QDA classifier trained to distinguish between BMI and manual control using all frequency bands from individual ROIs (M1 in blue, DLPFC in green, Cd in pink) or combinations of ROIs (purple) at the go cue for Monkey H (left) and Monkey Y (right). Classifiers are presented in order of ascending accuracy. Chance accuracy shown as a dashed line. Error bars represent standard error mean across days. **(c)** Mean classification accuracy for each combination of ROIs at the go cue. **(d,e)** Same as **(b,c)**, but for power features obtained at target acquisition rather than at the go cue.

power features from M1 resulted in significantly higher classifier performance than the classifiers using power features from DLPFC or Cd for Monkey H and from Cd for Monkey Y (Fig. 4d,e, Supplementary Fig. 6b). Additionally, the classification accuracy of models using power features from either M1 + DLPFC or M1 + Cd was not significantly different from that of models using power features from all ROIs for either subject. Furthermore, both the go cue and target acquisition results were similar across BMI learning stages (Supplementary Fig. 7). These results suggest that most of the information for distinguishing between BMI and manual control at the go cue is present in DLPFC, while most of the information at target acquisition is present in M1. This difference may indicate that cognitive mechanisms known to rely more heavily on DLPFC or M1, such as action-planning or movement-execution, are more involved at the go cue and at target acquisition, respectively.

Though DLPFC and M1 hold key information for distinguishing BMI and manual control at the go cue and at target acquisition, respectively, models using power features from all individual ROIs at both task events yielded above-chance classification accuracy (Fig. 4b and d). This indicates that each individual ROI holds important information for the 2-class task-type model. Additionally, models with features from 2 or 3 ROIs combined consistently yielded the highest classification accuracy (Supplementary Fig. 6). This increase in accuracy obtained by combining information across ROIs suggests that additional ROIs contribute distinct information.

To extend the comparison of neural representations during BMI and manual control in these ROIs, we introduced an 8-class target-direction Linear Discriminant Analysis (LDA) classifier to distinguish between the eight different target-directions in both manual control and BMI control scenarios (Fig. 5a). Similar to previous analyses, power from all frequency bands during the 500 ms after the go cue and before target acquisition from different ROIs was used as input to the models. To ensure a fair comparison between models trained on data from the two control types, an equal number of trials per target-direction under BMI and manual control were incorporated into their respective classifiers. Additionally, the data from the first 5 days were excluded from each model to ensure that there were a sufficient number of successful trials to each target (see “Methods” for details). Under manual control, accuracies of classifiers at the go cue using power features from M1 only were significantly greater than those using power features from DLPFC only and Cd only for both animals (Fig. 5b,c, Supplementary Fig. 8a). Under BMI control, accuracies of classifiers using M1 were significantly greater than models using Cd only for both animals (Fig. 5b,c, Supplementary Fig. 8a). Furthermore, models using power features from all 3 ROIs (M1 + DLPFC + Cd), M1 + DLPFC, and M1 + Cd out-performed the models using DLPFC only and Cd only under both manual and BMI control, suggesting that the addition of M1 power features boosts prediction of target-direction at the go cue during both modes of control (Fig. 5b,c, Supplementary Fig. 8a). While classification accuracies of models including M1 power features tended to be higher under manual control than under BMI control at the go cue, only M1 + Cd and M1 + DLPFC + Cd were significantly greater for both animals (Fig. 5c, Supplementary Table 3a). Conversely, at target acquisition, accuracies from nearly all models using power features from manual control were significantly greater than the models using the same power features from BMI control (Fig. 5e, Supplementary Table 3b). Under manual control, the classifiers containing power features from M1 + DLPFC and M1 + DLPFC + Cd performed the best at target acquisition (Fig. 5d,e). Under BMI control, there were no significant differences between any of the combinations of ROIs at target acquisition and very few models achieved above-chance accuracies (Supplementary Fig. 8b). In contrast, under manual control, all but one model (Cd only for Monkey Y) significantly predicted target-direction at target acquisition (Fig. 5d,e).

Overall, DLPFC appears to best distinguish between BMI and manual control at the go cue, while target-prediction at the go cue is best predicted by M1 within each mode of control. Thus, the differences in neural activity within DLPFC that differentiate BMI and manual control could pertain to factors other than the encoding of target-direction. At target acquisition, M1 best distinguishes between modes of control. M1 and DLPFC seemed to play a large role in predicting target-direction under manual control, while classifiers predicting the target-direction from LFP power features under BMI control perform near chance accuracy.

Information flow to M1 during both BMI and manual control

To investigate how these regions interact with one another during BMI and manual control, we evaluated the directed functional connectivity, or effective connectivity, using Granger causality. Granger causality is a statistical test for determining whether one time-series is useful in predicting another, and therefore provides insight into the direction of information flow between these ROIs. For successful trials under BMI and manual control, we calculated Granger causality in the 500 ms LFP segments following the go cue and those preceding target acquisition. The same calculations were also performed on random non-overlapping 500 ms windows from the baseline period. At both task events, we found significantly above-chance effective connectivity in every ROI → ROI interaction direction during all task types: BMI control, manual control, and baseline (see “Methods” for details). This indicates that every ROI sends information to every other ROI. Because all ROIs exhibit significant reciprocal interactions, even during baseline, we isolated task-specific directed information flow by calculating a normalized Granger causality that compared interactions during BMI and manual control to interactions during baseline. A normalized Granger causality value greater than zero indicates an increase from baseline, while a value less than zero indicates a decrease from baseline. This normalized value provides insight into which interactions are task-specific and accounts for variability in baseline information flow across days. We first analyzed this normalized metric of directed information flow during BMI control. We observed an increase in normalized Granger causality in the DLPFC → M1 and Cd → M1 directions and a decrease in the M1 → Cd direction in both monkeys at both the go cue and target acquisition (Fig. 6a,b). Additionally, there was an increase in the DLPFC → Cd direction for Monkey H and a decrease in the Cd → DLPFC direction for Monkey Y at the go cue (see Supplementary Table 4). These results suggest information flow from DLPFC → M1 and from Cd → M1 during both task events, as well as information flow from Cd → DLPFC at the go cue. Next, we determined the change in directed information flow from baseline during manual control. We observed a

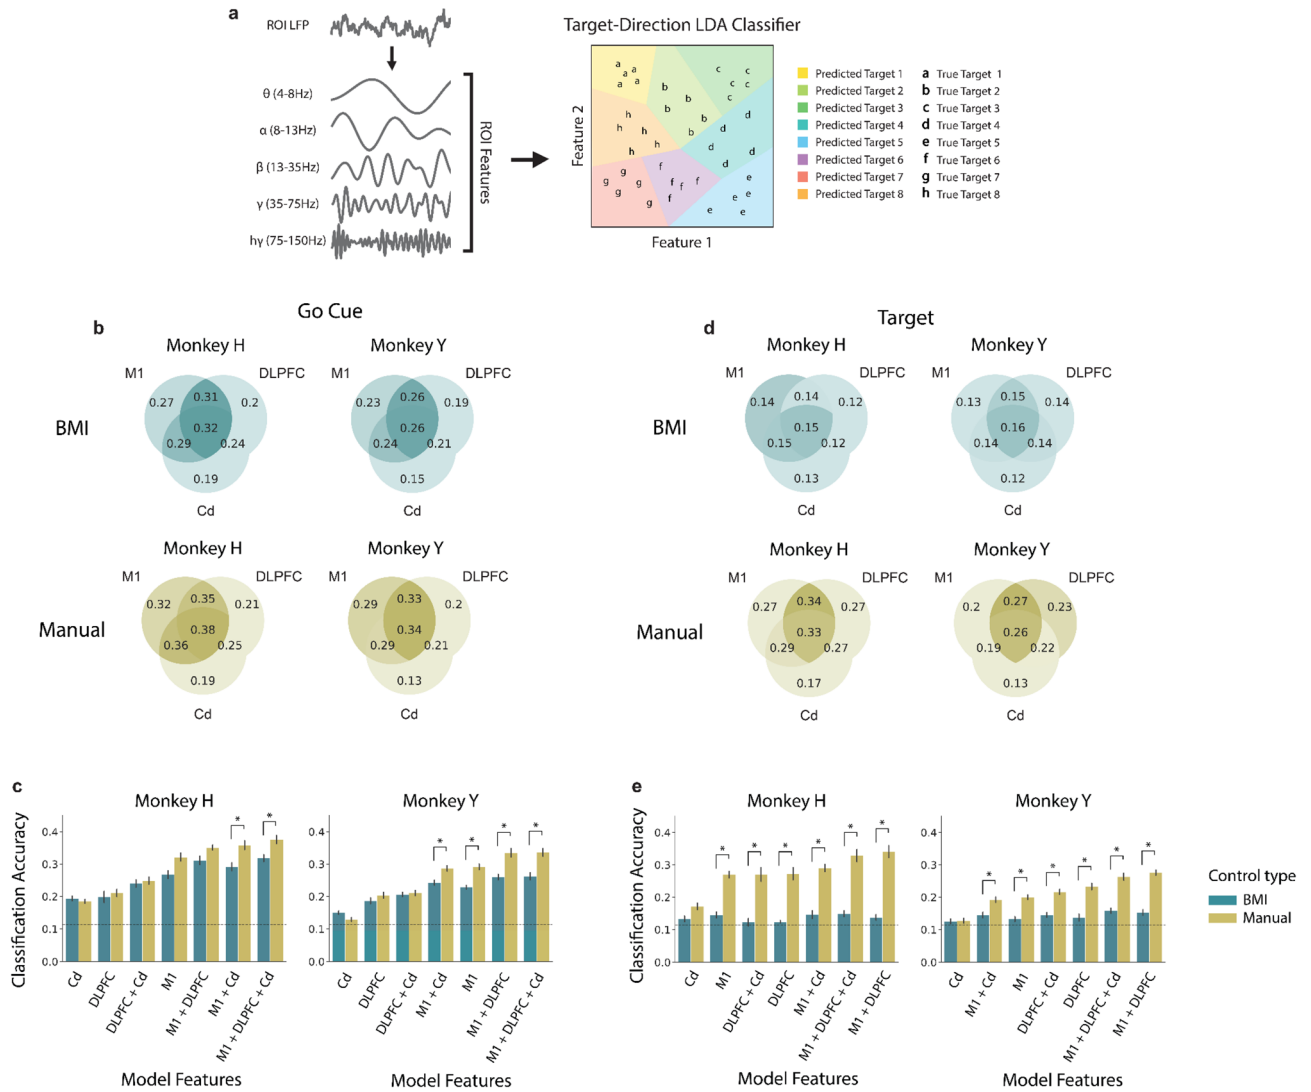


Figure 5. M1 activity best distinguishes between target-direction at the go cue for both BMI and manual control and DLPFC + M1 activity best predicts target-direction under manual control while all ROIs predict target-direction near chance accuracy under BMI control at target acquisition. **(a)** LFP was decomposed into 5 distinct frequency bands. The average power in each of these bands in each ROI was used as input to the classifiers. 8-class target-direction LDA classifiers were trained to distinguish between target-directions under BMI control and manual control separately. **(b)** Mean tenfold cross-validated classification accuracies for each combination of ROIs at the go cue under BMI control (teal) and manual control (yellow) for Monkey H (left) and Monkey Y (right). **(c)** Classification accuracies for the 8-class target-direction LDA classifiers for BMI control and manual control at the go cue. Classifiers are presented in order of ascending accuracy for manual control. Chance accuracy shown as a dashed line. Error bars represent standard error mean across days. Classification accuracies significantly differing between BMI and manual control after Bonferroni correction for multiple comparisons are indicated with an asterisk (Bonferroni corrected p-value < 0.05; see Supplementary Table 3 for exact values). **(d,e)** Same as **(b,c)**, but for power features obtained at target acquisition rather than at the go cue.

similar increase in normalized Granger causality from DLPFC → M1 and from Cd → M1 in both monkeys at both task events (Fig. 6a,b). However, unlike during BMI control, DLPFC ↔ Cd interactions were inconsistent across animals. While there were significant differences between BMI and manual control normalized Granger causality within a subset of interaction directions, these differences were inconsistent across monkeys and mainly differed in amplitude rather than direction (Fig. 6a,b, see Supplementary Table 5). Overall, these results suggest information flow from DLPFC → M1 and Cd → M1 is present during both BMI and manual control.

For both BMI and manual control, most reciprocal interactions underwent opposite changes from baseline. To consolidate these results, we calculated the net normalized Granger causality by taking the difference between each pair of reciprocal directed interactions (Fig. 6c). This metric, which represents the overall predominant direction of task-specific information flow, provides a clear representation of the network between our ROIs. We refer to information flow between two ROIs as significantly directed if we found the difference in normalized



Figure 6. Net flow of information between M1, DLPFC, and Cd relative to baseline is similar across BMI and manual control (a) Normalized Granger causality estimates at the go cue during BMI (teal) and manual (yellow) control. Normalized Granger causality greater than zero indicates an increase from baseline, whereas less than zero indicates a decrease from baseline. Normalized Granger causality estimates significantly differing between BMI and manual control are indicated with an asterisk (Bonferroni corrected p-value < 0.05; see Supplementary Table 5 for exact values). Error bars represent standard error mean across days. (b) Same as (a) but for target acquisition, rather than go cue. (c) Schematic depicting the calculation of net normalized granger causality. The difference in normalized granger causality for reciprocal interactions was used to estimate net flow of information relative to baseline. (d) Arrows represent significantly directed net normalized Granger causality (Bonferroni corrected p-value < 0.05) at the go cue during BMI control (see Supplementary Table 6 for all values of net normalized Granger causality). The absence of an arrow indicates that the net normalized Granger causality was not significantly directed. (e) Same as (d), but for target acquisition during BMI control. (f) Net normalized granger causality at the go cue during manual control. (g) Same as (f), but for target acquisition during manual control.

Granger causality to be significantly different from zero. During BMI control, we found that information was significantly directed from DLPFC → M1 and Cd → M1 for both monkeys at both task events (Fig. 6d,e, Supplementary Table 6a,b). At the go cue, information was also significantly directed from Cd → DLPFC (Fig. 6d, Supplementary Table 6a). Just as with our control-type classification results, we assessed whether this directed network was present across BMI learning stages and found no differences (Supplementary Fig. 9). Under manual control, we found similar directions of net change in information flow from baseline. The net normalized Granger causality was significantly directed, from DLPFC → M1 for both monkeys across both task events (Supplementary Table 6a,b), similar to BMI control (Fig. 6f,g). The interaction from Cd → M1 was significantly directed for Monkey Y, but was not above chance for Monkey H. Interactions between DLPFC and Cd were significantly directed from DLPFC → Cd in Monkey H, rather than Cd → DLPFC as observed during BMI control, but were not significant for Monkey Y. Within each monkey, these results were consistent across task events. Overall, there were many similarities in the net direction of information flow across both BMI and manual control. Under both control types, the net direction of Granger causality increased from DLPFC → M1 relative to baseline at both the go cue and target acquisition. However, the net direction of information flow between Cd and M1 and between DLPFC and Cd was not consistent across monkeys during manual control. Therefore, DLPFC may play an upstream role for both BMI and manual control, while the upstream role of Cd may be more clear for BMI control.

Our results demonstrate the presence of distinct neural representations of BMI and manual control in M1, DLPFC, and Cd. More specifically, we found that neural activity from DLPFC and M1 best distinguish between control types, at the go cue and target acquisition, respectively. Further, we identified net directed information flow from DLPFC → M1 in both BMI and manual control, and from Cd → M1 in BMI control. Together, these results suggest that DLPFC plays an important upstream role in both BMI and manual control tasks, with especially distinct neural representations during a period in the task that requires action planning, selection, and initiation.

Discussion

In this study, single- and multi-unit motor cortical recordings from a semi-chronic microdrive array were used to successfully control a BMI. We leveraged simultaneous LFP recordings from M1, DLPFC, and Cd to investigate differences in distributed activity within cortical and subcortical networks during motor BMI control and during overt motor (manual) control. Our results demonstrate the presence of distinct neural representations of BMI and manual control in M1, DLPFC, and Cd. Notably, we found that neural activity from DLPFC and M1 best distinguishes control types at the go cue and at target acquisition, respectively. Further, we identified net directed information flow from DLPFC → M1 in both BMI and manual control and from Cd → M1 in BMI control. Together, these results suggest that DLPFC plays an important upstream role in both BMI and manual control tasks, with especially distinct neural representations during a period in the task that requires action planning, selection, and initiation. On the contrary, M1 is a net-receiver of information and is better at distinguishing control types during a period in the task that requires more fine effector control. Overall, this work provides evidence of coordinated network activity between M1, DLPFC, and Cd during BMI control that is similar yet distinct from manual control.

We compared spectral power during baseline, BMI control, and manual control and found that theta power in M1, DLPFC, and Cd was significantly higher during the control tasks than during baseline at the go cue. In contrast, there were no significant distinctions in theta power at target acquisition. Previous work has demonstrated the involvement of theta activity in goal-directed information processing^{39,40}, with increased theta activity in the prefrontal cortex observed at decision points where choice-relevant information is under consideration. In the center-out task, the go cue can be thought of as a decision point, where subjects must choose a cursor trajectory, while the target acquisition period involves less choice-relevant information processing. Therefore, this observed increase in theta power across all ROIs at the go cue may be related to trajectory selection and planning. In contrast to theta power, beta power in M1, DLPFC, and Cd at the go cue was significantly lower during the control tasks than at baseline. At target acquisition, only M1 beta power remained distinct between baseline and the two control tasks. Beta power in the prefrontal cortex has been implicated in attention and task-specific rule representation^{41,42}. Beta power in the motor cortex and striatum can be modulated by task-relevant cues, movement preparation, and movement execution^{43–46}. Therefore, beta power within each of our ROIs may play an important role in maintaining task-relevant information, with DLPFC and Cd beta power modulated more specifically during trajectory planning at the go cue. These differences in oscillatory activity in the theta and beta frequency bands between baseline and the two control tasks suggest that this activity, which has previously been observed during planning, movement, and other goal-directed tasks, may be utilized with similar functionality in both BMI and manual control.

We found distinct BMI and manual neural representations in M1, DLPFC, and Cd at both task events when comparing combinations of frequency bands. We also found that control-type classification accuracy generally increased with the addition of neural activity from a second or third ROI, suggesting that each ROI holds some non-overlapping information for differentiating between BMI and manual control. One major difference between our BMI and manual control tasks is the lack of proprioceptive feedback in the BMI task. Many models of how the motor cortex controls movement incorporate sensory feedback as a relevant factor⁴⁷. One potential explanation for observed differences in M1 activity is the difference in sensory feedback between the two control types. The differences between control types in M1 LFP may also be related to changes in the properties of individual M1 neuron activity. Distinct changes in neural activity between BMI and manual control have been previously observed in single- and multi-unit motor cortical activity, in the form of changes in preferred direction and modulation depth^{29,48–51}. The differences we observed between BMI control and manual control in M1 LFP may reflect or drive these distinct single- and multi-unit properties across control types. DLPFC activity also successfully differentiated between BMI and manual control and is largely involved in action planning, selection, and initiation^{13,18,26}. Thus, the distinct neural representations in DLPFC activity may reflect differences between action planning, selection, and initiation between BMI and manual control. Furthermore, we also found that DLPFC activity was better than other ROIs at distinguishing control types at the go cue, a portion of the task that may rely more heavily on action planning, selection, and initiation, further supporting the hypothesis that the differences in DLPFC activity between BMI and manual control are related to differences in these cognitive processes. Interestingly, while DLPFC was best at distinguishing between control types at the go cue, M1 was best at distinguishing between target-directions, suggesting that DLPFC may have a role at the go cue that is distinct from target-direction encoding. The prefrontal cortex has also been implicated in rule representations and in task- and rule-switching¹⁸, as has the basal ganglia^{52,53}. In our experiment, proficient motor BMI and manual control could be viewed as two rules dictating how the subjects must perform the center-out task. Thus, the distinct neural representations in DLPFC and Cd may be involved in differentiating between these two control types as distinct rules.

Behavioral performance is another major difference between BMI and manual control, with BMI behavioral performance significantly changing across days while manual performance remains more stereotyped and consistent. These differences in consistency across trials may explain why predicting target-direction at target

acquisition under BMI control performed worse than under manual control. This also may suggest that BMI performance was more stereotyped at the go cue, where there were fewer significant differences between the modes of control, than at target acquisition. However, despite behavioral performance changes in the BMI control task over days, our control-type classification results did not differ across learning stages, implying that the distinctions in neural representations of control-type do not reflect skill learning and changes in performance. In contrast, previous literature has shown decreases in DLPFC activity and increases in striatal activity across motor learning⁵⁴. Additionally, decreased PFC activity has been observed across learning of a 1-D BMI task in humans^{8,35} as well as an increase in cortico-striatal coherence across learning of a 1-D BMI task in rodents³². However, the 2-D, center-out BMI task used in this study is more challenging than both an overtrained overt motor control task and a 1-D BMI control task. Because the task used in this study required the animal to hold a center target in order to successfully initiate a trial, even trial initiation required some amount of proficient control from each NHP. Thus, any truly initial learning that would need to occur prior to this baseline level of proficiency would not have been captured by our experimental design. Although our results were consistent across early and late learning stages, it is important to note that even at their most proficient BMI control, both subjects performed the task more slowly under BMI control and never achieved the same target acquisition time as manual control. Thus, even proficient BMI control may involve greater online error correction than manual control. Previous studies have discussed the role of the striatum in error-based motor learning and encoding error in movement observation^{55–58}. The distinctions we observe in Cd may represent this difference in online error correction across control types. However, the exact role of Cd in error correction is still debated and comparing the role of other regions implicated in error prediction, such as the cerebellum, is necessary⁵⁹. While further research is necessary to understand which cognitive processes are relevant to each brain region in BMI control, our results expand on prior knowledge of distinct neural representations between BMI and manual control in M1 and also show distinctions in the neural activity of distributed brain regions beyond M1 including in DLPFC and Cd.

To investigate how these regions interact with one another during both BMI and manual control, we quantified directed functional connectivity using Granger causality. Overall, we observed a similar direction of net information flow across ROIs between BMI and manual control despite M1 spikes being used as direct input to the BMI decoder, which may alter the neural dynamics within M1^{30,60} leading to possible changes in LFP synchronization across interconnected brain regions. Our results suggest that these potential changes in neural dynamics during BMI control do not greatly reshape the feedback loops between M1, PFC, and Cd that already exist under overt motor control. Instead, we observed net directed information flow from DLPFC → M1 during both BMI and manual control. Previous studies have identified structural and functional connections between prefrontal and motor cortices, with the prefrontal cortex sending information to the motor cortex regarding both motor regulatory functions^{13,61} and goal-directed behaviors²⁵. Our results support this upstream information processing model during overt motor control and suggest that a similar model is present during BMI control. The direction of net information flow was consistent across the go cue and target acquisition, suggesting that the information-processing network is stable despite shifts in primary cognitive demands throughout an individual trial and may be representative of a more general state of goal-directed behavior.

We also observed net directed information flow from Cd → M1 during BMI control. Much of the previous literature on cortico-striatal connections has focused on motor execution directed from M1 to the striatum due to the existence of direct structural connections between them²⁰. However, recent studies in mice have confirmed the presence of a cortico-striatal thalamo-cortical feedback loop that is relevant in motor control, suggesting that indirect modulation of the motor cortex by the striatum also plays an important role in motor control^{22,62–64}. Additionally, a recent study in humans suggests that basal ganglia may contribute to the regulation of sensorimotor cortical regions⁶⁵. While the bottom-up indirect feedback from the striatum to the motor cortex has not received as much attention, particularly in NHPs, our results suggest that this directed functional connection may be an important area of future study in BMI control.

Altogether, this study leverages simultaneous cortical and subcortical recordings during the same task under BMI control and under overt motor control to identify control-specific neural representations within M1, DLPFC, and Cd and to uncover an information-processing network from DLPFC and Cd to M1 during BMI control. Our results contribute to a growing body of work elucidating the neural mechanisms underlying BMI control, improving our understanding of BMIs as a scientific tool and as a therapeutic device. This study provides insight into the cortical and subcortical circuits supporting both motor and high-level cognitive functions associated with BMI control.

Methods

Surgery

Two rhesus macaques were implanted with recording chambers. Chamber positions were calculated based on images obtained from 3 T magnetic resonance imaging (MRI) (Siemens Medical Solutions, Malvern, PA) scans of each subject's brains. Regions of interest, including primary motor cortex (M1), dorsolateral prefrontal cortex (DLPFC), and caudate (Cd), were manually traced in 3D Slicer⁶⁶ using the Paxinos primate atlas as a reference⁶⁷. The resulting neuroanatomical models were used to decide on stereotaxic coordinates for implantation. All procedures were conducted in compliance with the National Institute of Health (NIH) Guide for the Care and Use of Laboratory Animals and were approved by the University of California at Berkeley Institutional Animal Care and Use Committee.

Large-scale semi-chronic microdrive

We used custom large-scale semi-chronic microdrive arrays for recording in both animals (Gray Matter Research, MT). These arrays allowed us to simultaneously record from regions of interest at different depths using independently-moveable single microelectrodes ($n = 124$). Electrodes consisted of both glass-coated Tungsten electrodes (Alpha Omega, Nof HaGalil (Nazareth Illit), Israel) and Platinum-Iridium electrodes (MicroProbes for Life Sciences, Gaithersburg, MD). Throughout implantation and recovery, electrodes were stored inside of the microdrive. The impedance of the electrodes was monitored using a TDT NanoZ (Tucker-Davis Technologies, Alachua, Florida) while advancing them into the brain. First, electrodes were lowered out of the microdrive until they penetrated the dura. While most electrodes successfully entered the brain (Monkey H: 79/124; Monkey Y: 64/124), a subset broke upon penetrating the dura (Monkey H: 45/124; Monkey Y: 60/124). These recording channels were excluded from all analyses. Electrodes that successfully entered the brain were advanced until their respective neural targets were reached and unit activity was detected. The estimated target depth of each electrode was calculated using the neuroanatomical models built in 3D Slicer. Ultimately, we successfully lowered a subset of electrodes into each region of interest (Monkey H: M1 = 27, PMd = 10, DLPFC = 22, Cd = 8; Monkey Y: M1 = 12, PMd = 12, DLPFC = 16, Cd = 8).

Intracortical Recording

Neural data were recorded using the OmniPlex Neural Recording Data Acquisition System (Plexon Inc, Dallas, TX). Single- and multi-unit activity was sorted prior to beginning recording sessions using an online sorting application (Sort Client, Plexon Inc, Dallas, TX). Wideband activity was recorded at 5 kHz. LFP activity was obtained by low-pass filtering at 250 Hz, notch filtering at 60 Hz and 120 Hz, and down-sampling to 1 kHz. LFP activity was common median referenced by first z-scoring activity within each channel (x_i where $i : \{1, 2, \dots, n\}$) by subtracting the mean and dividing by the standard deviation in each recording session (s)

$$Z_{is} = \frac{x_{is} - \mu_{is}}{\sigma_{is}}$$

and then subtracting the median value across all channels at each time point (t).

$$Z'_{it} = Z_{it} - \text{median}(\{Z_{1t}, Z_{2t}, \dots, Z_{nt}\})$$

After subtracting the median to remove common signals, likely low frequency noise, across channels, the LFP activity was multiplied by the standard deviation and the mean was added to restore the LFP to its original scaling.

$$x'_{is} = (Z'_{is} * \sigma_{is}) + \mu_{is}$$

Common median referencing was used instead of common average referencing to avoid influence from outliers. The common median referenced LFP activity was then z-scored using the mean and standard deviation within channel across all recording sessions within a day (1, 2, ...S).

$$Z''_{is} = \frac{x'_{is} - \mu'_{i1,2,\dots,S}}{\sigma'_{i1,2,\dots,S}}$$

Center-out task

Subjects performed a self-paced, center-out, reaching task to eight targets. Trials were initiated by moving a cursor to the central target. A successful trial required a short hold at the center, moving to the peripheral target within a specified time, and a brief hold at the target. Successful trials resulted in a juice reward; failed trials were repeated up to 10 times before a new target was presented. Target directions were presented in a pseudo-randomized order.

Subjects were first overtrained in the center-out task performed with arm movements before starting BMI. In this manual control (MC) version of the task, the subject's arm moved in a KINARM exoskeleton (BKIN Technologies, Kingston, ON, Canada) that restricted movements to the horizontal plane. Neural activity recorded during MC was used to train a BMI decoder. Using this BMI decoder, the animals performed the same task under BMI control (BC). During BC, the animals' arms were restricted in a fixed position within the exoskeleton and the animals were required to move the cursor to the target by modulation of motor cortex activity.

Brain-machine interface

Subjects learned to control a two-dimensional BMI cursor in real-time using a fixed velocity Kalman Filter (KF) decoder^{68–70}. The KF assumes two linear models:

$$x_{t+1} = Ax_t + w_t$$

$$y_{t+1} = Cx_t + q_t$$

where x_t and y_t are the cursor state and neural activity at time t , respectively. The first equation represents the state-transition model, which describes the state of the cursor over time. It is specified by the state-transition matrix and additive Gaussian noise term. The second equation represents the observation model and describes the relationship between neural activity and cursor state. It is parameterized by the observation matrix and additive Gaussian noise. Neural activity was input as a vector of spike counts in 100 ms bins from the selected direct units.

Decoder parameters were initialized from neural and cursor kinematic data collected during the MC version of the task at the beginning of each recording day. Maximum likelihood estimation methods were used to fit initial parameters. Neural data from 10 to 20 single- and multi-units recorded from motor cortex were selected for BMI control each day (i.e., direct units). The population of direct units was highly overlapping from day to day, but there was some variability as units dropped out or new ones appeared. For Monkey Y, closed-loop decoder adaptation (CLDA) was performed using the SmoothBatch algorithm before fixing the BMI decoder. This algorithm uses knowledge of task goals (i.e., reaching targets) to infer a subject's intent. The intended kinematics and observed neural activity during closed-loop BMI were used to re-estimate KF parameters. The SmoothBatch algorithm^{71–73} re-estimates the observation model of the KF (matrices C and Q), and updates were constrained to enforce smoothness. CLDA was typically run for 2–5 min to provide the subject with adequate performance to allow successful reaches to all targets. For Monkey H, the initial decoder trained from MC data was used.

Neural data analysis

All analyses were performed in Python with custom-written routines utilizing publicly available software packages including scipy⁷⁴, numpy⁷⁵, sklearn⁷⁶, and statsmodels⁷⁷. Unless otherwise specified, analyses involving activity from DLPFC and Cd were performed on the average LFP signal across all channels within DLPFC and Cd, respectively. Analyses involving M1 were performed on the average LFP signal across all channels that were used to record single- or multi-unit activity that was input to the BMI decoder.

Task-type classifiers

In order to understand which of our ROIs were most predictive of control type, we trained classifiers to predict between task-types using power from different ROIs and frequency bands as input features. We used quadratic discriminant analysis (QDA) models using singular value decomposition (SVD) from the sklearn toolbox in Python. We chose QDA as our classification model as it does not assume equal variances, unlike other linear classifiers, which was essential for our 3-class task-type input data.

Power in the theta (4–8 Hz), alpha (8–13 Hz), beta (13–35 Hz), gamma (35–75 Hz), and highgamma (75–150 Hz) frequency bands was calculated using Welch's method on segments of LFP recorded during each control type. For BMI and manual control, power was computed in the 500 ms window following the go cue and in the 500 ms window preceding target acquisition on successful trials. Power estimates were averaged across all channels in each ROI prior to being used as input features. While no individual channels were notable outliers within each ROI, we chose to use the channel-average in our analyses in order to prevent unequal representation of ROIs with more channels than others (Supplementary Fig. 10). The same calculations were also performed in random non-overlapping 500 ms windows during the baseline period. To avoid negative effects on imbalanced data between classes, the number of trials selected for each task type was matched to the minimum number of trials per individual task type each day by selecting a random subset of trials from the larger class to match that minimum. The corresponding number of windows were selected from the baseline period as well.

Results for each model were validated using a tenfold cross-validation within each recording day. For each round of cross-validation, QDA classifiers were trained on 90% of data on each day (3-class: Monkey H: 151–696 trials, Monkey Y: 224–699 trials; 2-class: Monkey H: 100–464 trials, Monkey Y: 149–466 trials) and test on 10% of data on each day (3-class: Monkey H: 16–77 trials, Monkey Y: 24–77 trials; 2-class: Monkey H: 11–51 trials, Monkey Y: 16–51 trials). To validate whether classification accuracy results were significantly above chance, we used a paired t-test to compare each result to a chance accuracy. For each frequency band and ROI combination, chance accuracy was calculated by randomly shuffling class labels and rerunning the QDA classification (see above for ranges of training and test data set sizes) 1000 times for each day.

To determine whether particular frequency bands were useful for distinguishing between task types, we computed the QDA classification accuracy using each individual frequency band from each ROI as well as using all frequency bands from each ROI as input (Supplementary Fig. 3). We determined that all frequency bands in all ROIs could predict the task type significantly above chance and that the classifier trained on power features from all 5 frequency bands outperformed the classifiers using individual bands, so all subsequent analyses use classifiers trained on all 5 frequency bands from each ROI. This resulted in 5 features per ROI included in a QDA classifier, with models ranging from 5 to 15 features in total depending on the input ROI combination.

Target-direction classifiers

To further compare neural representations under BMI and manual control, we also trained classifiers to predict between the eight target-directions in the task under BMI control and under manual control separately. As in our task-type classifiers, we used power from different ROIs and frequency bands as input features (same as described above). We used linear discriminant analysis (LDA) models using SVD from the sklearn toolbox in Python. Rather than employing QDA, as we did for the task-type classifiers, we used LDA models for target-direction prediction due to training data for these classification containing one-eighth of the number of samples per class, which could result in overfitting with QDA models. Variances across target-directions were equal, therefore satisfying LDA's equal variances assumption.

Data from the first 5 recording days was excluded from these analyses to ensure that there were enough successful trials per each target-direction under BMI control to train a classifier without overfitting. Because there were an unequal number of successful trials under manual control and BMI control, a random selection of equal number of BMI and manual trials per target-direction per day was selected for input into the classifiers. This ensured that all targets and control types were equally represented, allowing for a fair comparison of accuracies between the two modes of control. Results for each model were validated using a tenfold cross-validation within each included recording day. For each round of cross-validation, LDA classifiers were trained on 90%

of data (8-class: Monkey H: 165–230 trials, Monkey Y: 136–230 trials) and tested on the remaining 10% of data (8-class: Monkey H: 18–25 trials, Monkey Y: 15–25 trials). To validate whether classification accuracy results were significantly above chance, we used a paired t-test to compare each result to a chance accuracy. For each frequency band and ROI combination, chance accuracy was calculated by randomly shuffling class labels and rerunning the LDA classification (see above for ranges of training and test data set sizes) 1000 times for each day.

Granger causality

Granger causality was used to estimate the directional functional connectivity among all pairs of regions of interest. Granger causality relies on an autoregressive (AR) modeling framework, in which future values of a time series are modeled as a weighted combination of past values of time series. The quality of an AR-model is assessed by quantifying the variance of the model's residuals. If the variance of the AR-model's residuals is reduced by the inclusion of past measurements from a second time series, then the second time series is said to Granger-cause or G-cause the first^{78,79}. Applying this logic, we obtained Granger causality estimates using simultaneously recorded LFP signals. We compared Bayesian information criterion (BIC) of AR-models of different orders, p , for each combination of LFP signals obtained from all pairs of regions during each trial from all recording sessions. We selected $p = 13$ to minimize BIC in the average case, allowing for the signals to be sufficiently long enough to capture the data structure without over-parameterization.

For successful trials under BMI and manual control, we calculated Granger causality in the 500 ms average LFP segments following the go cue and those preceding target acquisition. The same calculations were also performed on random non-overlapping 500 ms average LFP segments from the baseline period. To obtain an estimate of spurious interactions, we calculated Granger causality between the average LFP during baseline and the average LFP during BMI control. Because these signals were recorded at separate times in different tasks, any interactions that occur between regions would be artifactual. True interactions between regions within baseline and BMI control were considered significant if they were statistically different from this null distribution. The average Granger causality value across BMI control trials within each day was determined and compared to that of the null distribution. A null distribution for manual trials was calculated and compared using the same protocol.

We compared ROI → ROI Granger causality ($G_{Y \rightarrow X}$) in each interaction direction to the null distribution via an unpaired t-test and found that every interaction was significantly greater than the null distribution during all task types: BMI control, manual control, and baseline. In order to isolate task-relevant interactions, we computed the normalized Granger causality for BMI and manual control as follows:

$$(norm.G)_{Y \rightarrow X, BMI} = \frac{G_{Y \rightarrow X, BMI} - G_{Y \rightarrow X, Baseline}}{G_{Y \rightarrow X, Baseline}}$$

$$(norm.G)_{Y \rightarrow X, Manual} = \frac{G_{Y \rightarrow X, Manual} - G_{Y \rightarrow X, Baseline}}{G_{Y \rightarrow X, Baseline}}$$

Here, a normalized Granger causality greater than zero indicates an increase from baseline, whereas a normalized Granger causality less than zero indicates a decrease from baseline. To further assess the direction of information flow, we also calculated the net normalized Granger causality by subtracting the value of normalized Granger causality between ROI → ROI reciprocal interactions.

Quantification and statistical analyses

All analyses were performed within a single recording day and error is depicted across days. For analyses comparing two distributions or comparing a single distribution to a value, two-sample or one-sample t-tests were used, respectively. Bonferroni correction was used post hoc to correct for multiple comparisons. Significance is reported after correction for multiple comparisons.

Ethical approval

All procedures were conducted in compliance with the NIH Guide for the Care and Use of Laboratory Animals and reporting in the manuscript follows the recommendations in the ARRIVE guidelines. All procedures were approved by the University of California at Berkeley Institutional Animal Care and Use Committee under protocol ID AUP-2014-09-6720-2.

Data availability

The datasets analyzed during the current study are publicly available on Figshare at https://figshare.com/projects/Shvartsman_et_al_2023/177054.

Received: 31 May 2023; Accepted: 7 October 2023

Published online: 19 October 2023

References

1. Fetz, E. E. Volitional control of neural activity: implications for brain–computer interfaces: Volitional control of neural activity. *J. Physiol.* 579(3), 571–579. <https://doi.org/10.1113/jphysiol.2006.127142> (2007).
2. Carmena, J. M. *et al.* Learning to control a brain–machine interface for reaching and grasping by primates. *PLoS Biol.* 1(2), e42. <https://doi.org/10.1371/journal.pbio.0000042> (2003).

3. Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G. & Moran, D. W. A brain–computer interface using electrocorticographic signals in humans. *J. Neural Eng.* **1**(2), 63–71. <https://doi.org/10.1088/1741-2560/1/2/001> (2004).
4. Hochberg, L. R. *et al.* Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* **442**(7099), 164–171. <https://doi.org/10.1038/nature04970> (2006).
5. Golub, M. D., Yu, B. M. & Chase, S. M. Internal models for interpreting neural population activity during sensorimotor control. *eLife* **4**, e10015. <https://doi.org/10.7554/eLife.10015> (2015).
6. Collinger, J. L. *et al.* High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet* **381**(9866), 557–564. [https://doi.org/10.1016/S0140-6736\(12\)61816-9](https://doi.org/10.1016/S0140-6736(12)61816-9) (2013).
7. Green, A. M. & Kalaska, J. F. Learning to move machines with the mind. *Trends Neurosci.* **34**(2), 61–75. <https://doi.org/10.1016/j.tins.2010.11.003> (2011).
8. Wander, J. D. *et al.* Distributed cortical adaptation during learning of a brain–computer interface task. *Proc. Natl. Acad. Sci.* **110**(26), 10818–10823. <https://doi.org/10.1073/pnas.1221127110> (2013).
9. Vyas, S. *et al.* Neural population dynamics underlying motor learning transfer. *Neuron* **97**(5), 1177–1186.e3. <https://doi.org/10.1016/j.neuron.2018.01.040> (2018).
10. Hoshi, T., Zagotta, W. & Aldrich, R. Biophysical and molecular mechanisms of Shaker potassium channel inactivation. *Science* **250**(4980), 533–538. <https://doi.org/10.1126/science.2122519> (1990).
11. Graybiel, A. M., Aosaki, T., Flaherty, A. W. & Kimura, M. The basal ganglia and adaptive motor control. *Sci. New Ser.* **265**(5180), 1826–1831 (1994).
12. Groenewegen, H. J. The basal ganglia and motor control. *Neural Plast.* **10**(1–2), 107–120. <https://doi.org/10.1155/NP.2003.107> (2003).
13. Goldman-Rakic, P. S. Motor control function of the prefrontal cortex. In *Novartis Foundation Symposia* (eds Bock, G. *et al.*) 187–200 (Wiley, 2007). <https://doi.org/10.1002/9780470513545.ch12>.
14. Wallis, J. D., Anderson, K. C. & Miller, E. K. Single neurons in prefrontal cortex encode abstract rules. *Nature* **411**(6840), 953–956. <https://doi.org/10.1038/35082081> (2001).
15. Muhammad, R., Wallis, J. D. & Miller, E. K. A comparison of abstract rules in the prefrontal cortex, premotor cortex, inferior temporal cortex, and striatum. *J. Cogn. Neurosci.* **18**(6), 974–989. <https://doi.org/10.1162/jocn.2006.18.6.974> (2006).
16. Badre, D., Kayser, A. S. & D'Esposito, M. Frontal cortex and the discovery of abstract action rules. *Neuron* **66**(2), 315–326. <https://doi.org/10.1016/j.neuron.2010.03.025> (2010).
17. Antzoulatos, E. G. & Miller, E. K. Differences between neural activity in prefrontal cortex and striatum during learning of novel abstract categories. *Neuron* **71**(2), 243–249. <https://doi.org/10.1016/j.neuron.2011.05.040> (2011).
18. Miller, E. K. & Cohen, J. D. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* **24**(1), 167–202. <https://doi.org/10.1146/annurev.neuro.24.1.167> (2001).
19. Haber, S. N. The primate basal ganglia: Parallel and integrative networks. *J. Chem. Neuroanat.* **26**(4), 317–330. <https://doi.org/10.1016/j.jchemneu.2003.10.003> (2003).
20. Haber, S. N. Corticostriatal circuitry. *Dialogues Clin. Neurosci.* **18**(1), 21 (2016).
21. Yeterian, E. H., Pandya, D. N., Tomaiuolo, F. & Petrides, M. The cortical connectivity of the prefrontal cortex in the monkey brain. *Cortex* **48**(1), 58–81. <https://doi.org/10.1016/j.cortex.2011.03.004> (2012).
22. Guadalupe Cruz, K. *et al.* Cortical-subcortical interactions in goal-directed behavior. *Physiol. Rev.* **103**(1), 347–389. <https://doi.org/10.1152/physrev.00048.2021> (2023).
23. Costa, R. M., Cohen, D. & Nicolelis, M. A. L. Differential corticostriatal plasticity during fast and slow motor skill learning in mice. *Curr. Biol.* **14**(13), 1124–1134. <https://doi.org/10.1016/j.cub.2004.06.053> (2004).
24. Antzoulatos, E. G. & Miller, E. K. Increases in functional connectivity between prefrontal cortex and striatum during category learning. *Neuron* **83**(1), 216–225. <https://doi.org/10.1016/j.neuron.2014.05.005> (2014).
25. Buschman, T. J. & Miller, E. K. Goal-direction and top-down control. *Philos. Trans. R. Soc. B Biol. Sci.* **369**(1655), 20130471. <https://doi.org/10.1098/rstb.2013.0471> (2014).
26. Tanji, J. & Hoshi, E. Behavioral planning in the prefrontal cortex. *Curr. Opin. Neurobiol.* **11**(2), 164–170. [https://doi.org/10.1016/S0959-4388\(00\)00192-6](https://doi.org/10.1016/S0959-4388(00)00192-6) (2001).
27. Pearson, K. Motor systems. *Curr. Opin. Neurobiol.* **10**(5), 649–654. [https://doi.org/10.1016/S0959-4388\(00\)00130-6](https://doi.org/10.1016/S0959-4388(00)00130-6) (2000).
28. Golub, M. D., Chase, S. M., Batista, A. P. & Yu, B. M. Brain–computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Curr. Opin. Neurobiol.* **37**, 53–58. <https://doi.org/10.1016/j.conb.2015.12.005> (2016).
29. Ganguly, K., Dimitrov, D. F., Wallis, J. D. & Carmena, J. M. Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat. Neurosci.* **14**(5), 662–667. <https://doi.org/10.1038/nn.2797> (2011).
30. Zippi, E. L., You, A. K., Ganguly, K. & Carmena, J. M. Selective modulation of cortical population dynamics during neuroprosthetic skill learning. *Sci. Rep.* **12**(1), 15948. <https://doi.org/10.1038/s41598-022-20218-3> (2022).
31. Gallego, J. A., Makin, T. R. & McDougle, S. D. Going beyond primary motor cortex to improve brain–computer interfaces. *Trends Neurosci.* **45**(3), 176–183. <https://doi.org/10.1016/j.tins.2021.12.006> (2022).
32. Koralek, A. C., Jin, X., Long, J. D. II., Costa, R. M. & Carmena, J. M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**(7389), 331–335. <https://doi.org/10.1038/nature10845> (2012).
33. Koralek, A. C., Costa, R. M. & Carmena, J. M. Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron* **79**(5), 865–872. <https://doi.org/10.1016/j.neuron.2013.06.047> (2013).
34. Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M. & Carmena, J. M. Volitional modulation of primary visual cortex activity requires the basal ganglia. *Neuron* **97**(6), 1356–1368.e4. <https://doi.org/10.1016/j.neuron.2018.01.051> (2018).
35. Wander, J. D. *et al.* Cortico-cortical interactions during acquisition and use of a neuroprosthetic skill. *PLoS Comput. Biol.* **12**(8), e1004931. <https://doi.org/10.1371/journal.pcbi.1004931> (2016).
36. Casimo, K., Weaver, K. E., Wander, J. & Ojemann, J. G. BCI use and its relation to adaptation in cortical networks. *IEEE Trans. Neural Syst. Rehabil. Eng.* **25**(10), 1697–1704. <https://doi.org/10.1109/TNSRE.2017.2681963> (2017).
37. Dotson, N. M., Hoffman, S. J., Goodell, B. & Gray, C. M. A large-scale semi-chronic microdrive recording system for non-human primates. *Neuron* **96**(4), 769–782.e2. <https://doi.org/10.1016/j.neuron.2017.09.050> (2017).
38. Qiao, S., Brown, K. A., Orsborn, A. L., Ferrantino, B. & Pesaran, B. Development of semi-chronic microdrive system for large-scale circuit mapping in macaque mesolimbic and basal ganglia systems. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (IEEE, 2016) 5825–5828. <https://doi.org/10.1109/EMBC.2016.7592052>.
39. Womelsdorf, T., Vinck, M., Leung, L. S. & Everling, S. Selective theta-synchronization of choice-relevant information subserves goal-directed behavior. *Front. Hum. Neurosci.* <https://doi.org/10.3389/fnhum.2010.00210> (2010).
40. Phillips, J. M., Vinck, M., Everling, S. & Womelsdorf, T. A long-range fronto-parietal 5- to 10-Hz network predicts ‘top-down’ controlled guidance in a task-switch paradigm. *Cereb. Cortex* **24**(8), 1996–2008. <https://doi.org/10.1093/cercor/bht050> (2014).
41. Voloh, B. & Womelsdorf, T. Cell-type specific burst firing interacts with theta and beta activity in prefrontal cortex during attention states. *Cereb. Cortex* **28**(12), 4348–4364. <https://doi.org/10.1093/cercor/bhx287> (2018).
42. Buschman, T. J., Denovellis, E. L., Diogo, C., Bullock, D. & Miller, E. K. Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron* **76**(4), 838–846. <https://doi.org/10.1016/j.neuron.2012.09.029> (2012).
43. O’Leary, J. G. & Hatsopoulos, N. G. Early visuomotor representations revealed from evoked local field potentials in motor and premotor cortical areas. *J. Neurophysiol.* **96**(3), 1492–1506. <https://doi.org/10.1152/jn.00106.2006> (2006).

44. Saleh, M., Reimer, J., Penn, R., Ojakangas, C. L. & Hatsopoulos, N. G. Fast and slow oscillations in human primary motor cortex predict oncoming behaviorally relevant cues. *Neuron* **65**(4), 461–471. <https://doi.org/10.1016/j.neuron.2010.02.001> (2010).
45. Khanna, P. & Carmena, J. M. Neural oscillations: Beta band activity across motor networks. *Curr. Opin. Neurobiol.* **32**, 60–67. <https://doi.org/10.1016/j.conb.2014.11.010> (2015).
46. Kilavik, B. E., Zaepffel, M., Brovelli, A., MacKay, W. A. & Riehle, A. The ups and downs of beta oscillations in sensorimotor cortex. *Exp. Neurol.* **245**, 15–26. <https://doi.org/10.1016/j.expneurol.2012.09.014> (2013).
47. Omrani, M., Kaufman, M. T., Hatsopoulos, N. G. & Cheney, P. D. Perspectives on classical controversies about the motor cortex. *J. Neurophysiol.* **118**(3), 1828–1848. <https://doi.org/10.1152/jn.00795.2016> (2017).
48. Ganguly, K. & Carmena, J. M. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.* **7**(7), e1000153. <https://doi.org/10.1371/journal.pbio.1000153> (2009).
49. Canolty, R. T., Ganguly, K. & Carmena, J. M. Task-dependent changes in cross-level coupling between single neurons and oscillatory activity in multiscale networks. *PLoS Comput. Biol.* **8**(12), e1002809. <https://doi.org/10.1371/journal.pcbi.1002809> (2012).
50. Orsborn, A. L. *et al.* Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron* **82**(6), 1380–1393. <https://doi.org/10.1016/j.neuron.2014.04.048> (2014).
51. Zhang, Y. *et al.* Changes in modulation characteristics of neurons in different modes of motion control using brain–machine interface. In *Annu Int Conf IEEE Eng Med Biol Soc*, no. 2021, 6445–6448 (2021) <https://doi.org/10.1109/EMBC46164.2021.9630212>.
52. Hikosaka, O. & Isoda, M. Switching from automatic to controlled behavior: Cortico-basal ganglia mechanisms. *Trends Cogn. Sci.* **14**(4), 154–161. <https://doi.org/10.1016/j.tics.2010.01.006> (2010).
53. Yehene, E., Meiran, N. & Soroker, N. Basal ganglia play a unique role in task switching within the frontal-subcortical circuits: Evidence from patients with focal lesions. *J. Cogn. Neurosci.* **20**(6), 1079–1093. <https://doi.org/10.1162/jocn.2008.20077> (2008).
54. Dayan, E. & Cohen, L. G. Neuroplasticity subserving motor skill learning. *Neuron* **72**(3), 443–454. <https://doi.org/10.1016/j.neuron.2011.10.008> (2011).
55. Houk, J. C. *et al.* Action selection and refinement in subcortical loops through basal ganglia and cerebellum. *Philos. Trans. Biol. Sci.* **362**(1485), 1573–1583 (2007).
56. Tunik, E., Houk, J. C. & Grafton, S. T. Basal ganglia contribution to the initiation of corrective submovements. *NeuroImage* **47**(4), 1757–1766. <https://doi.org/10.1016/j.neuroimage.2009.04.077> (2009).
57. Schiffer, A.-M. & Schubotz, R. I. Caudate nucleus signals for breaches of expectation in a movement observation paradigm. *Front. Hum. Neurosci.* <https://doi.org/10.3389/fnhum.2011.00038> (2011).
58. Seidler, R. D., Kwak, Y., Fling, B. W. & Bernard, J. A. Neurocognitive mechanisms of error-based motor learning. In *Progress in Motor Control, Advances in Experimental Medicine and Biology* Vol. 782 (eds Richardson, M. J. *et al.*) 39–60 (Springer New York, 2013). https://doi.org/10.1007/978-1-4614-5465-6_3.
59. Turner, R. S. & Desmurget, M. Basal ganglia contributions to motor control: A vigorous tutor. *Curr. Opin. Neurobiol.* **20**(6), 704–716. <https://doi.org/10.1016/j.conb.2010.08.022> (2010).
60. Athalye, V. R., Ganguly, K., Costa, R. M. & Carmena, J. M. Emergence of coordinated neural dynamics underlies neuroprosthetic learning and skillful control. *Neuron* **93**(4), 955–970.e5. <https://doi.org/10.1016/j.neuron.2017.01.016> (2017).
61. Selemon, L. & Goldman-Rakic, P. Common cortical and subcortical targets of the dorsolateral prefrontal and posterior parietal cortices in the rhesus monkey: Evidence for a distributed neural network subserving spatially guided behavior. *J. Neurosci.* **8**(11), 4049–4068. <https://doi.org/10.1523/JNEUROSCI.08-11-04049.1988> (1988).
62. Haber, S. N. & Calzavara, R. The cortico-basal ganglia integrative network: The role of the thalamus. *Brain Res. Bull.* **78**(2–3), 69–74. <https://doi.org/10.1016/j.brainresbull.2008.09.013> (2009).
63. Oldenburg, I. A. & Sabatini, B. L. Antagonistic but not symmetric regulation of primary motor cortex by basal ganglia direct and indirect pathways. *Neuron* **86**(5), 1174–1181. <https://doi.org/10.1016/j.neuron.2015.05.008> (2015).
64. Foster, N. N. *et al.* The mouse cortico–basal ganglia–thalamic network. *Nature* **598**(7879), 188–194. <https://doi.org/10.1038/s41586-021-03993-3> (2021).
65. Kasahara, K., DaSalla, C. S., Honda, M. & Hanakawa, T. Basal ganglia–cortical connectivity underlies self-regulation of brain oscillations in humans. *Commun. Biol.* **5**(1), 712. <https://doi.org/10.1038/s42003-022-03665-6> (2022).
66. Fedorov, A. *et al.* 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magn. Reson. Imaging* **30**(9), 1323–1341. <https://doi.org/10.1016/j.mri.2012.05.001> (2012).
67. Paxinos, G., Huang, X.-F. & Toga, A. W. *The Rhesus Monkey Brain in Stereotaxic Coordinates* (Academic Press, 2000).
68. Gilja, V. *et al.* A high-performance neural prosthesis enabled by control algorithm design. *Nat. Neurosci.* **15**(12), 1752–1757. <https://doi.org/10.1038/nn.3265> (2012).
69. Kim, S.-P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P. & Black, M. J. Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *J. Neural Eng.* **5**(4), 455–476. <https://doi.org/10.1088/1741-2560/5/4/010> (2008).
70. 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**(1), 80–118. <https://doi.org/10.1162/089976606774841585> (2006).
71. Dangi, S. *et al.* Continuous closed-loop decoder adaptation with a recursive maximum likelihood algorithm allows for rapid performance acquisition in brain–machine interfaces. *Neural Comput.* **26**(9), 1811–1839. https://doi.org/10.1162/NECO_a_00632 (2014).
72. Gowda, S., Orsborn, A. L., Overduin, S. A., Moorman, H. G. & Carmena, J. M. Designing dynamical properties of brain–machine interfaces to optimize task-specific performance. *IEEE Trans. Neural Syst. Rehabil. Eng.* **22**(5), 911–920. <https://doi.org/10.1109/TNSRE.2014.2309673> (2014).
73. Orsborn, A. L., Dangi, S., Moorman, H. G. & Carmena, J. M. Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions. *IEEE Trans. Neural Syst. Rehabil. Eng.* **20**(4), 468–477. <https://doi.org/10.1109/TNSRE.2012.2185066> (2012).
74. Virtanen, P. *et al.* SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2> (2020).
75. Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**(7825), 357–362. <https://doi.org/10.1038/s41586-020-2649-2> (2020).
76. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *Mach. Learn. PYTHON*.
77. Seabold, S. & Perktold, J. Statsmodels: Econometric and Statistical Modeling with Python. In *Presented at the Python in Science Conference, Austin, Texas*, 92–96 (2010) <https://doi.org/10.25080/Majora-92bf1922-011>.
78. Bressler, S. L. & Seth, A. K. Wiener–Granger causality: A well established methodology. *NeuroImage* **58**(2), 323–329. <https://doi.org/10.1016/j.neuroimage.2010.02.059> (2011).
79. Ding, M., Chen, Y. & Bressler, S. L. Granger causality: Basic theory and application to neuroscience. In *Handbook of Time Series Analysis: Recent Theoretical Developments and applications* 437–460 (2006).

Acknowledgements

We thank Maki Kitano for assistance with NHP care, handling, and training, Paul Botros for methodological advice, and Charles Gray for help with surgical procedures. This work was supported by the National Science Foundation Graduate Research Fellowship (to E.L.Z.), the National Institute of Mental Health Award R01MH117763 (to J.D.W.), and the National Institute of Health Award R01NS106094 (to J.M.C.).

Author contributions

E.L.Z., G.F.S., J.D.W., and J.M.C. designed and performed research. E.L.Z. and G.F.S. analyzed the data. E.L.Z., G.F.S., and N.V-L. interpreted the data. E.L.Z. and G.F.S. wrote the manuscript text and prepared the manuscript figures. N.V-L., J.D.W., and J.M.C. edited the manuscript. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.


Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-023-44405-y>.

Correspondence and requests for materials should be addressed to J.M.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

 **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023