nature human behaviour

Article

https://doi.org/10.1038/s41562-023-01804-5

Motor cortex retains and reorients neural dynamics during motor imagery

Received: 24 January 2023

Accepted: 13 December 2023

Published online: 29 January 2024

Check for updates

Brian M. Dekleva[1,2,3], Raeed H. Chowdhury[3,4], Aaron P. Batista[3,4], Steven M. Chase[3,5,6], Byron M. Yu[3,5,7], Michael L. Boninger[1,2,4] & Jennifer L. Collinger[1,2,4,5]

The most prominent characteristic of motor cortex is its activation during movement execution, but it is also active when we simply imagine movements in the absence of actual motor output. Despite decades of behavioural and imaging studies, it is unknown how the specific activity patterns and temporal dynamics in motor cortex during covert motor imagery relate to those during motor execution. Here we recorded intracortical activity from the motor cortex of two people who retain some residual wrist function following incomplete spinal cord injury as they performed both actual and imagined isometric wrist extensions. We found that we could decompose the population activity into three orthogonal subspaces, where one was similarly active during both action and imagery, and the others were active only during a single task type—action or imagery. Although they inhabited orthogonal neural dimensions, the action-unique and imagery-unique subspaces contained a strikingly similar set of dynamic features. Our results suggest that during motor imagery, motor cortex maintains the same overall population dynamics as during execution by reorienting the components related to motor output and/or feedback into a unique, output-null imagery subspace.

As people prepare to execute a skilled action, they often pause to mentally rehearse and visualize it. For example, a tennis player might imagine hitting an upcoming serve, or a pianist might imagine playing a difficult sequence prior to performance. This type of covert motor imagery is constrained to the same performance limits as overt execution. One study showed that the speed with which people were able to imagine performing a sequence of finger movements was limited to their actual overt performance[1]. Motor imagery is similarly impacted by neurologic impairment; a lesion in motor cortex leads to an equal slowing of both executed and imagined movements[2]. Conversely, imagery-based practice can improve actual motor function, in some instances offering a comparable performance benefit as standard overt training[3–6]. This tight coupling between imagery and actual motor function suggests similar central mechanisms, so that information and experience gleaned from one modality can usefully inform the other.

Primary motor cortex, known mainly for its role in the generation of volitional movement, is also active during covert motor imagery[7–11]. In fact, many movement-related brain areas are also active during the mental rehearsal of imagined movements. Premotor and supplementary motor cortices[12,13], anterior cingulate areas[13] and parietal areas[13,14] all display modulated activity during covert motor imagery. Despite the clear link between imagery and action, we know little about how cortical population activity differs between the two. Previous studies in both monkeys[15] and humans[10,11] have shown evidence that motor cortex activity is somewhat consistent across volitional states but exhibits clear differences as well. In particular, the study in monkeys by

[1]Rehab Neural Engineering Labs, University of Pittsburgh, Pittsburgh, PA, USA. [2]Department of Physical Medicine & Rehabilitation, University of Pittsburgh, Pittsburgh, PA, USA. [3]Center for the Neural Basis of Cognition, Pittsburgh, PA, USA. [4]Department of Bioengineering, University of Pittsburgh, Pittsburgh, PA, USA. [5]Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA. [6]Neuroscience Institute, Carnegie Mellon University, Pittsburgh, PA, USA. [7]Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA. e-mail: collinger@pitt.edu

Jiang et al.[15] found that a portion of motor cortex activity could be partitioned into distinct subspaces containing unique responses during either overt hand control of an on-screen cursor or passive observation of cursor movements. This segmentation of activity into orthogonal subspaces appears to be a common motif in motor cortex, the most established example being that of movement preparation[16,17]. Preparation is certainly distinct from motor imagery; movement preparation involves 'readying' for an imminent overt action, while motor imagery is the covert rehearsal of a complete action. However, during both imagery and preparation, the motor system faces a similar objective: to engage in movement-related processing while avoiding the activation of descending control pathways. In the case of movement preparation, the cortical implementation appears to be well explained by the coordination of a small number of population activity patterns[16,17]. Within this framework, activity in motor and/or premotor cortices activates orthogonal neural subspaces during preparation and execution. This orthogonality allows preparatory activity to evolve while avoiding the dimensions that would engage descending pathways and cause overt movement.

Given that imagery, by definition, does not involve overt movement, it seems reasonable to assume that imagery activity in motor cortex, like preparatory activity, somehow avoids dimensions responsible for downstream control. One possibility is that imagery exists only in a subset of the neural dimensions active during overt action. Another possibility is that imagery exists in dimensions completely orthogonal to those for action. However, both of these possibilities are unlikely given the results from Jiang et al.[15], which show strong evidence for a large degree of overlap across volitional states, plus additional subspaces specific to overt (arm movement) and covert (cursor observation) volitional states. Related work by Vargas-Irwin et al.[11] and Rastogi et al.[10] also shows that despite broad similarities between attempted and imagined movements in the motor cortex of paralysed humans, different volitional states are still highly discriminable. However, the population analyses used in these studies did not identify the underlying geometry driving the separation in volitional state. The present study sets out to characterize the population-level organization across volitional states and to identify the dynamic features that are shared as well as those that are unique to either overt motor action or covert motor imagery.

Here we use an isometric wrist extension task to examine the relationship between imagery and action in motor cortex. Two participants with tetraplegia due to spinal cord injury participated in the study. Despite having no hand or lower-extremity function, both retained residual proximal arm and wrist extension control. We recorded from intracortical microelectrode arrays implanted in the hand and arm areas of motor cortex as the participants performed either real or imagined isometric wrist extensions to achieve low and high force targets. After reducing the recorded population activity to a low-dimensional manifold, we found that it contained three distinct subspaces: (1) a shared space, in which responses were nearly identical during action and imagery; (2) an action-unique subspace that was only modulated during actual force production; and (3) an imagery-unique subspace that was only modulated during imagery. Strikingly, we found that the neural dynamics in the imagery-unique subspace during imagery closely resembled those observed in the action-unique subspace during execution. These unique subspace dynamics also contained elements that did not exist in the shared subspace. From this, we conclude that motor cortex maintains the same overall neural dynamics during imagery as during overt action. However, since the population activity must avoid output dimensions (and lacks modulation along feedback-related dimensions) during imagery, cortex reorients output-related or feedback-related responses—contained within the action-unique subspace—into an orthogonal, imagery-unique (output-null) subspace. We propose that the retention of overall neural dynamics structure during imagery provides the motor system with a useful proxy for overt practice.

Results

Motor cortex is active during actual and imagined force

We asked the participants (P2 and P3) to perform isometric wrist extensions within an immobile frame affixed with a force sensor to control the height of a line trace displayed on a monitor in front of them (Fig. 1a). During 'action' trials, they observed a horizontal bar indicating the required force (either low (~5 N) or high (~40 N)) and then were required to apply the appropriate force such that the line trace matched the vertical position of the bar, holding the target force for approximately four seconds. During 'imagery' trials, the participants kept their hand within the force-sensing apparatus and were asked to imagine producing the same wrist extension forces without actually doing so. On imagery trials, the line trace automatically increased to the target force and then returned to zero. For each session, we collected alternating 12-trial blocks of action and imagery, resulting in 36 total trials of action and 36 trials of imagery. We then removed trials with force profiles that deviated substantially from the average (see 'Experimental setup' in Methods), resulting in approximately 31 ± 4 action trials and 28 ± 4 imagery trials per session (six sessions for P2 and three sessions for P3).

Both participants successfully achieved and maintained the requested force targets during action trials (Fig. 1c, top, and Supplementary Figs. 1 and 2) and produced no appreciable force during imagery trials (Fig. 1c, bottom, and Supplementary Figs. 1 and 2). Throughout the experimental sessions, we recorded activity from the hand and arm areas of motor cortex (Fig. 1b). Despite the stark difference in force output between action and imagery trials, we observed modulation in the overall population-wide firing rates for both task types (Fig. 1d, top versus bottom, and Supplementary Table 1). Individual channels displayed a wide variety of responses, including a mix of preferential activation during action or imagery (Fig. 1e). Some channels displayed similar modulation during both action and imagery (for example, Fig. 1e, channel 183), while others appeared uniquely active during only one task type (for example, Fig. 1e, channels 188 and 19). For each channel, we calculated the maximum modulation during action and imagery by calculating the difference between the 5th-percentile and 95th-percentile firing rates for each task type (imagery or action) and force level. We found that for P2, the average low-force imagery modulation was 75 ± 6% that of low-force action modulation (the bounds represent 95% confidence intervals (CIs) obtained via bootstrapping with 10,000 resamples). High-force imagery modulation was 57 ± 4% that of action (Fig. 1f, left). For P3, low-force imagery modulation was 59 ± 5% that of action, and high-force imagery modulation was 53 ± 6% that of action (Fig. 1f, right).

Latent space contains distinct action and imagery subspaces

As a first step towards characterizing the differences between the neural representations of action and imagery, we asked whether a portion of the population activity could be separated into unique dimensions containing only action or imagery activity. For simple motor behaviours, the measured dimensionality of motor cortical activity is typically far lower than the number of recorded neurons[18–23]. Thus, as an initial step, we reduced the activity from our recorded populations (176 channels for P2 and 192 channels for P3) to a lower-dimensional (~36D) latent space. To do this, we first performed principal component analysis (PCA) separately on action and imagery trials, keeping enough dimensions to explain >99% of the variance for each task. We then combined them into a single, unified space (see 'Dimensionality reduction' in Methods for details). Within this combined space, we observed an incomplete overlap between the dimensions containing action and imagery variance. For example, by performing singular value decomposition on action data from P2, we found that the leading 11 dimensions explained >99% of action variance but only 66% of imagery variance (Fig. 2a). Similarly, the leading 12 imagery dimensions explained >99% of imagery variance but only 70% of action variance (Fig. 2b). This discrepancy in the set of dimensions containing meaningful variance for each

https://doi.org/10.1038/s41562-023-01804-5

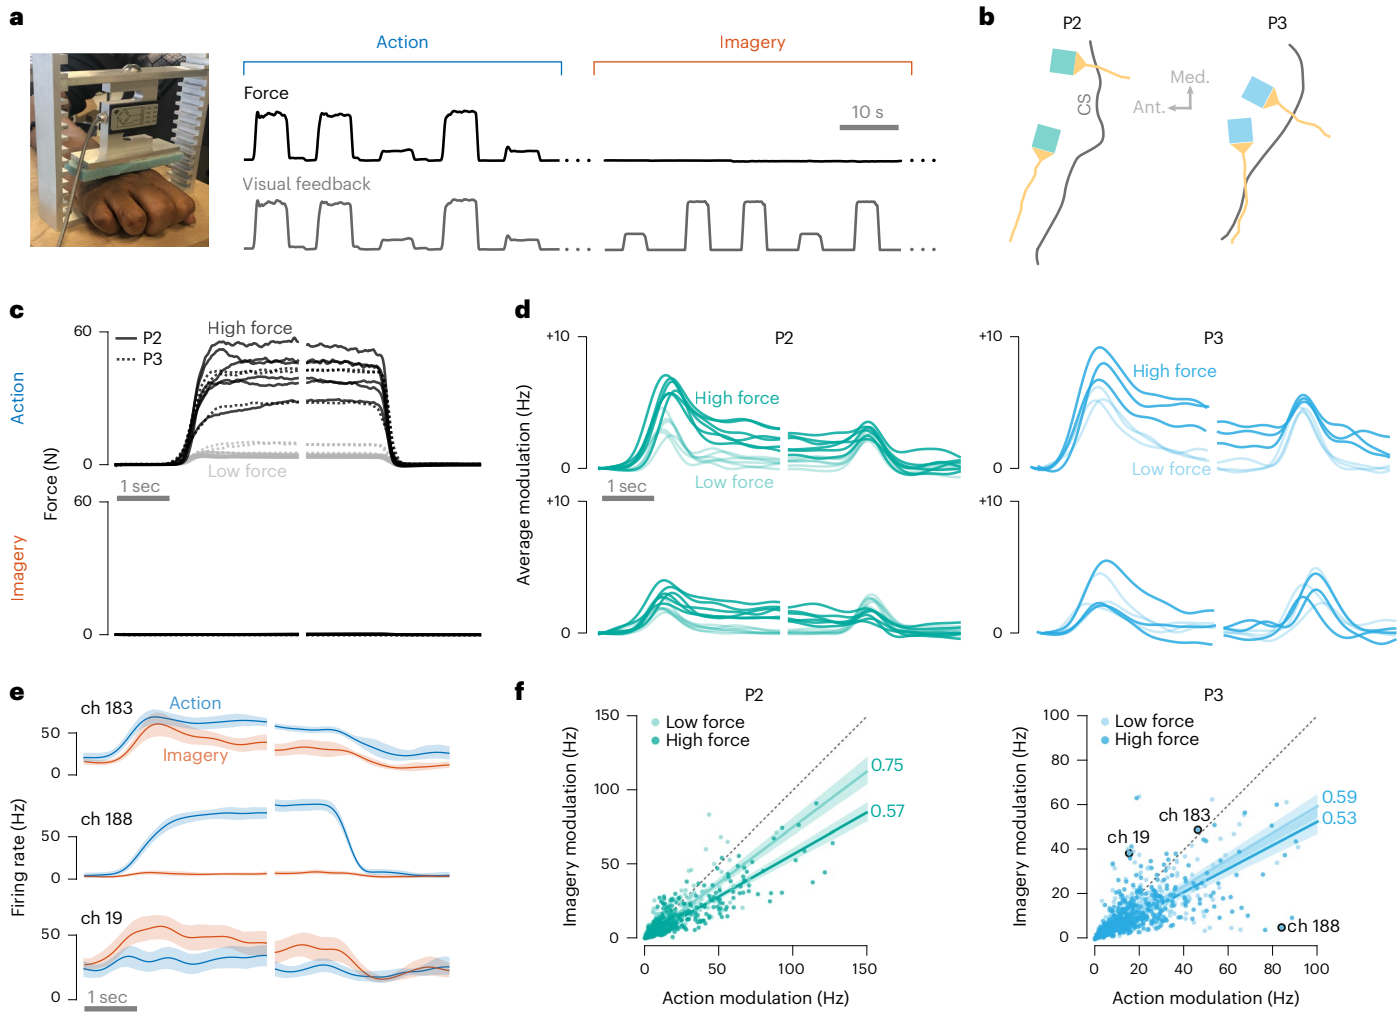

Fig. 1 | Motor cortex is active during both imagery and action. a, The participants placed their hands on a board beneath a load cell and produced either real or imagined wrist extension forces. For all trials, they received visual feedback of either their actual produced force (action trials) or an automated proxy (imagery trials). **b**, Locations of microelectrode arrays implanted in the motor cortices. CS, central sulcus. **c**, Average low-force (grey) and high-force (black) traces for all sessions (solid lines indicate P2; dashed lines indicate P3) during action trials (top) and imagery trials (bottom). **d**, Average population firing rate modulation in motor cortex (M1) during action trials (top) and imagery trials (bottom). Here modulation is calculated as the change in firing rate from the trial start. Each trace corresponds to a single session. Lighter traces represent low-force trials, and darker traces represent high-force trials. **e**, Average activity for three example channels (ch) (P3) during high-force action (blue) and high-force imagery (red) trials. The shading represents 95% confidence bounds, calculated by bootstrapping with 1,000 resamples. **f**, Maximum modulation during low-force (light) and high-force (dark) action and imagery for all recorded channels and all sessions (left, P2; right, P3). The number next to each line is the slope. The shaded bounds represent 95% CIs, calculated by bootstrapping with 1,000 resamples.

task suggests that the neural subspaces involved in the two tasks were not fully aligned. We summarized this partial alignment in Fig. 2d using the subspace alignment index metric[16], which ranges from 0 (when subspaces are orthogonal) to 1 (when subspaces are fully aligned). The action and imagery tasks showed only moderate alignment, well below what we would expect to be able to resolve if they were actually aligned given the trial-by-trial variability (Fig. 2d; alignments calculated across 1,000 shuffled datasets were uniformly larger than for the true dataset; $P < 0.001$; two-sided t-test). These results suggest that some portion of the activity inhabited distinct orthogonal subspaces unique to both action and imagery.

On the basis of this observation, we developed a method to jointly identify the action-unique, imagery-unique and shared (common) subspaces (see 'Subspace separation' in Methods). The identified transformations from the latent space to the three subspaces were constrained to be fully orthogonal, such that the three subspaces fully spanned the original space (that is, the percentages in Fig. 2c sum to 100%). The final transformation thus provided a different view of the same underlying latent space such that activity clustered into discrete subspaces with unique task-related variance characteristics (Fig. 2e). The action-unique subspace contained high variance responses only during the action task, the imagery-unique subspace contained high variance only during the imagery task and the shared subspace contained both action and imagery variance. The subspace splitting procedure can be considered an extension of the alignment index concept[16]—the proportion of variance explained by the shared subspace is essentially equivalent to the alignment index value. In line with Jiang et al.[15], the subspaces did not correspond to distinct subpopulations within the recordings (unimodal distribution of channel contributions to each subspace; P2, $P = 0.99$, P3, $P = 0.99$; Hartigan's dip test; Supplementary Fig. 3). Rather, individual channels exhibited mixed selectivity, with the clear condition-specific subspace structure appearing only at the broader population level[24].

In addition to total variance explained, we examined the dimensionality of each subspace (Fig. 2f,g). To do this, we first performed a varimax rotation on the activity within each subspace and ranked

https://doi.org/10.1038/s41562-023-01804-5

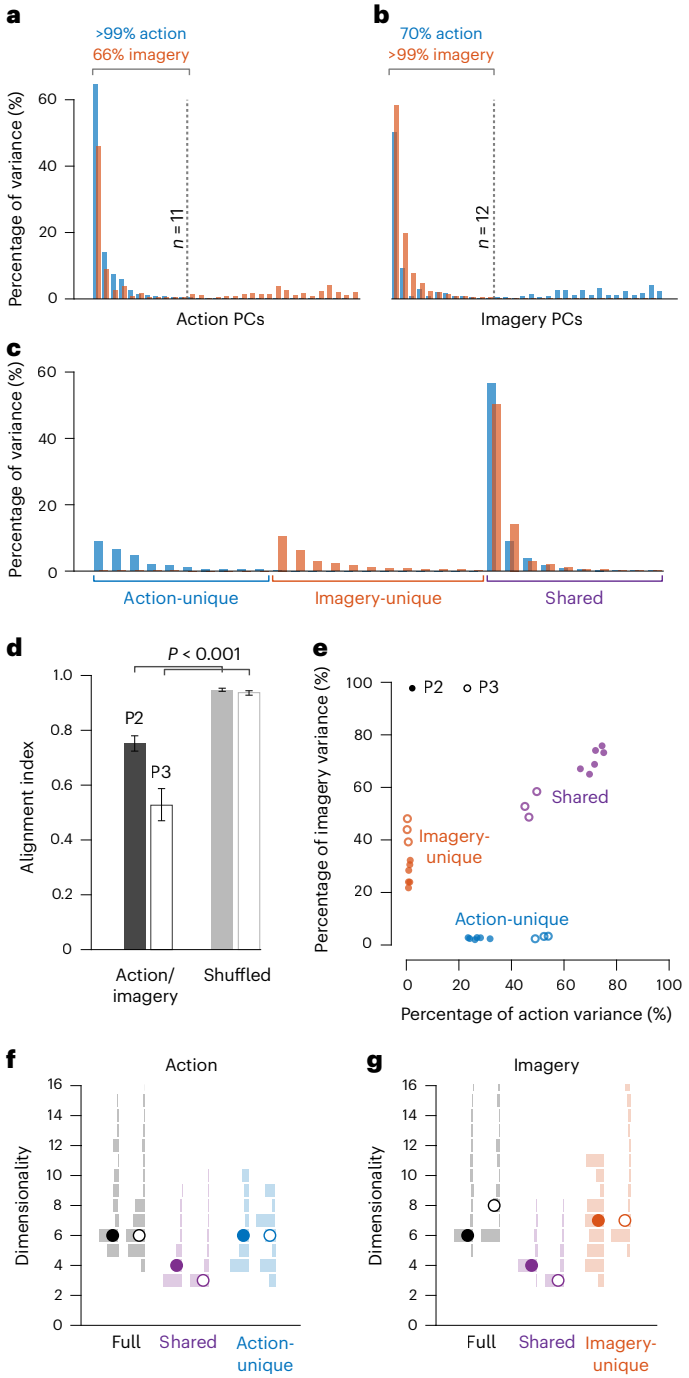

Fig. 2 | Population activity contains distinct action, imagery, and shared subspaces. a, An example from P2 showing the percentage of total action (blue) and imagery (red) variance explained by principal components (PCs) computed on action trials for an example session. **b**, Same as in **a**, but for PCs computed on imagery trials. **c**, The same space as in **a**,**b** after an orthogonal rotation to isolate distinct action and imagery variances into unique subspaces. **d**, The alignment index between action and imagery, as well as a label-shuffled control. The alignment index reflects the proportion of variance for one condition that is captured by the leading PCs computed from the opposite condition. The error bars represent 95% CIs through bootstrapping with 1,000 resamples. For both participants, the action/imagery alignments were significantly lower than the shuffled label control ($P < 0.001$, two-sided t-test). **e**, Percentage of action and imagery variance captured by the action-unique, imagery-unique and shared subspaces following the procedure in **c**. **f**, Estimates of the dimensionality of the full latent space, shared subspace and action-unique subspace during action. **g**, Estimates of the dimensionality of the full latent space, shared subspace and imagery-unique subspace during imagery. The shaded vertical histograms in **f**,**g** represent the distributions of dimensionalities across 1,000 threshold values from 0.5% to 2%.

the resulting components by variance. Compared with PCA, varimax provided a more severe 'elbow' in the ranked variances, which in turn gave a more consistent estimate of dimensionality. We then counted the number of dimensions that accounted for more than 1% of the total task variance on all 1,000 bootstrapped resamples across trials (see 'Feature extraction and subspace dimensionality estimation' in Methods). We found that the estimated dimensionalities of the action and imagery subspaces were mostly larger than the dimensionality of the shared subspace for both P2 (action-unique versus shared, $P = 0.122$; imagery-unique versus shared, $P < 0.001$; bootstrap with 1,000 resamples) and P3 (action-unique versus shared, $P = 0.002$; imagery-unique versus shared, $P < 0.001$; bootstrap with 1,000 resamples). Additionally, the estimated dimensionalities of the full space during each task were significantly less than the sum of the shared and unique dimensionalities (P2 action, $P < 0.001$; P2 imagery, $P = 0.001$; P3 action, $P < 0.001$; P3 imagery, $P = 0.013$). This is because the unique subspaces are characterized by their cross-task variance properties. However, within a single task, there exists a more compact representation of the total population response.

Common dynamics within the shared subspace

The previous analysis of the neural latent space revealed a shared subspace containing activity for both tasks, and two separate subspaces that were differentially active depending on whether the participant exerted force or imagined exerting force. Next, we turn to an examination of how neural activity evolved in time within each of these subspaces.

Projecting trial-averaged activity from both tasks into the shared subspace revealed visually similar temporal profiles between action and imagery for each shared dimension (Fig. 3b). To assess the extent of this correlation irrespective of chosen dimension, we performed a Monte Carlo method (see 'Monte Carlo sampling' in Methods) in which we sampled 10,000 random unit vectors from the shared subspace (Fig. 3a, top). On each draw, we computed the correlation between the action and imagery activity along that dimension. Across all sampled dimensions, we found median correlations of 0.93 (P2; 95% CI, (0.75, 0.97)) and 0.88 (P3; 95% CI, (0.64, 0.94)), indicating that the shared subspace activity was universally well matched between the two task types.

Common dynamics across action-unique and imagery-unique subspaces

The comparison of temporal components between the action-unique and imagery-unique subspaces is less straightforward than for the shared subspace. By construction, the unique subspaces are orthogonal to each other and contain meaningful activity only during their respective tasks, making it futile to compare activity across tasks along a single dimension. Instead, we first found a rotation of the imagery subspace axes that aligned the multidimensional responses observed in the imagery-unique subspace during imagery to those in the action-unique subspace during action (see 'Action–imagery subspace alignment' in Methods). After this alignment procedure, we observed that the action and imagery spaces appeared to comprise a similar set of temporal components (Fig. 3c,d), despite existing in orthogonal subspaces of the population space.

We quantified the overall similarity between the multidimensional action-unique and imagery-unique responses by calculating correlations between the action-unique activity on randomly chosen dimensions in the action-unique subspace and the imagery-unique activity along the corresponding (aligned) imagery-unique dimensions (Fig. 3a, bottom). Across 10,000 randomly sampled dimensions, we found median correlations of 0.93 (P2; 95% CI, (0.84, 0.97)) and 0.86 (P3; 95% CI, (0.76, 0.94)). To provide context for these values, we also found, for each randomly selected action-unique dimension, the maximally correlated imagery-unique dimension that did not incorporate the

https://doi.org/10.1038/s41562-023-01804-5

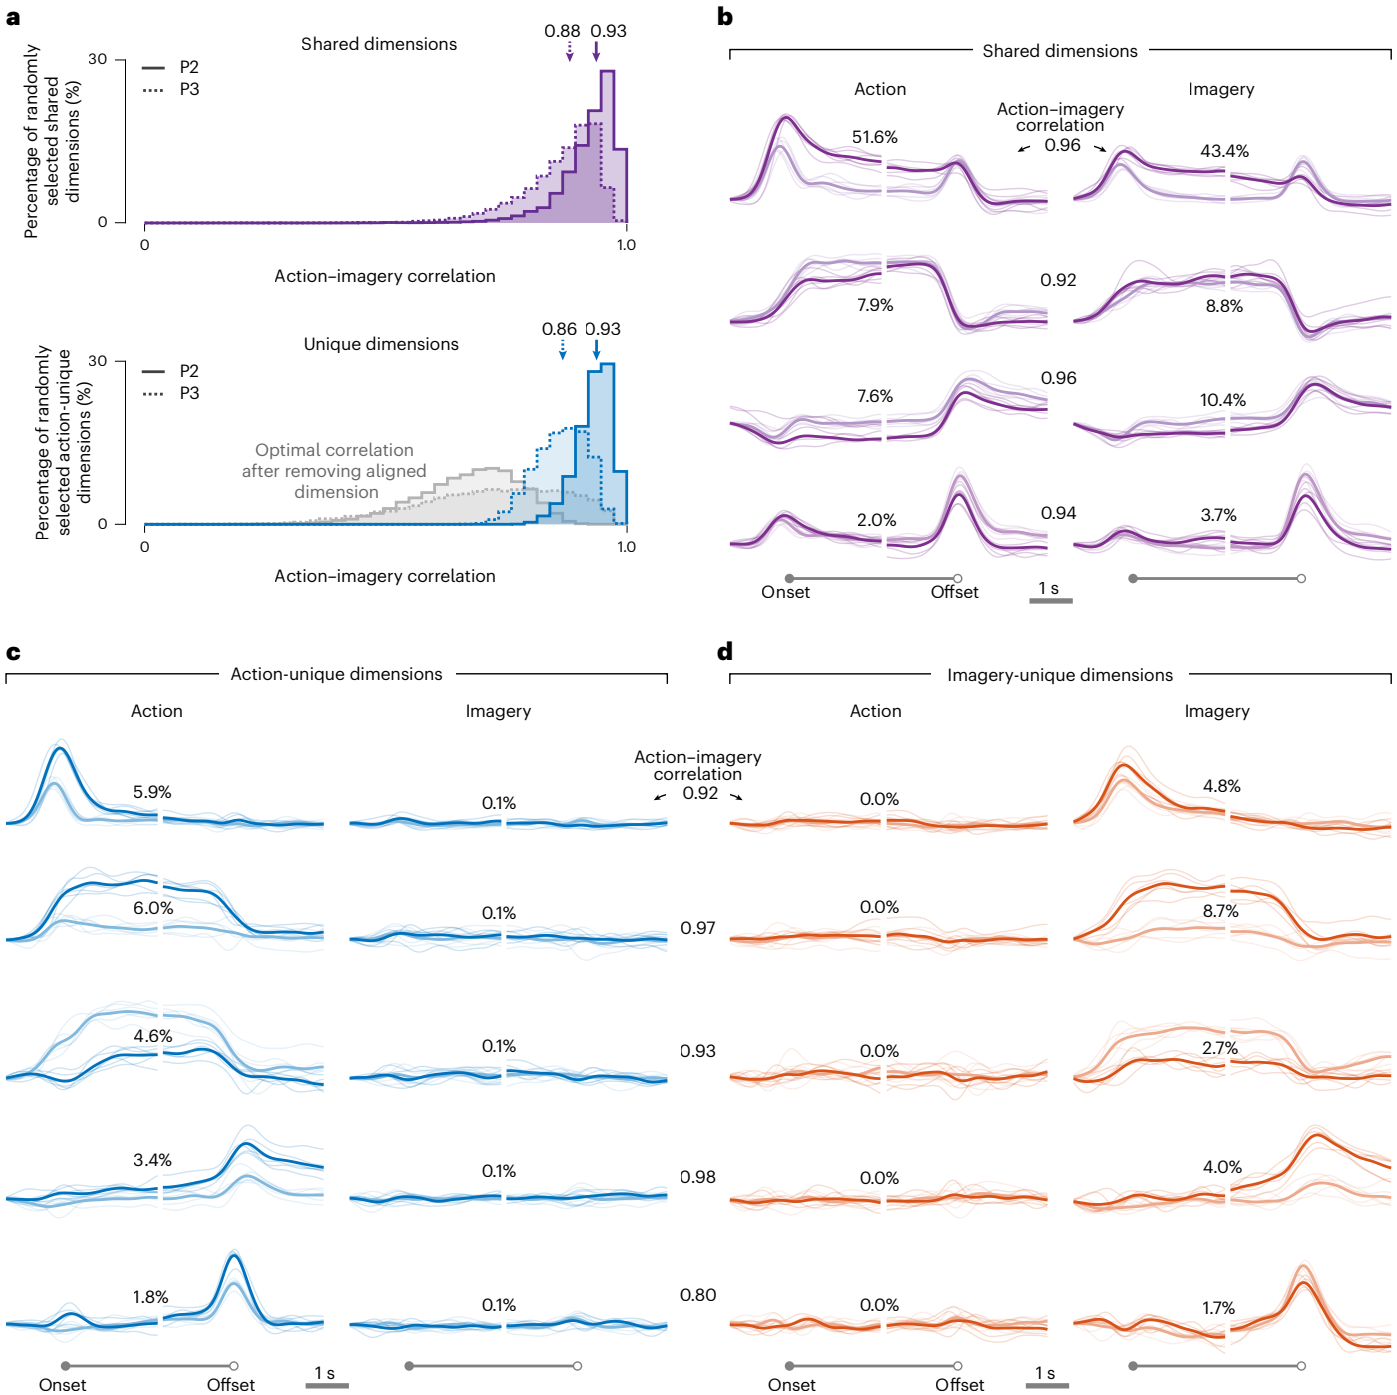

Fig. 3 | Temporal components of action and imagery are similar within the shared subspace and across unique subspaces. a, Top, correlations between action and imagery responses for 10,000 random dimensions selected from the shared subspace. The arrows indicate the median correlations for each participant. Bottom, correlations between aligned action-unique and imagery-unique components for 10,000 randomly selected dimensions. The arrows indicate the median correlations for each participant. The grey distributions represent the maximal possible correlations if the corresponding aligned imagery dimension is removed on each random draw. **b**, Activity within the shared subspace during action (left) and imagery (right). **c**, Activity within the action-unique subspace during action (left) and imagery (right). **d**, Activity within the imagery-unique subspace during action (left) and imagery (right), following an orthogonal transformation to align to the action-unique responses (see 'Action–imagery subspace alignment' in Methods). For **b**–**d**, the thin traces correspond to individual sessions (P2), and the thick lines correspond to cross-session averages. Light and dark traces represent low- and high-force conditions, respectively. The action–imagery correlation values in **c**,**d** correspond to the correlations between action-unique dimensions during action and imagery-unique dimensions during imagery.

aligned dimension (Fig. 3a, bottom, grey histograms). These secondary dimensions represent the next-best correlation, indicating the degree of triviality in the original alignment. Temporal components reflecting non-specific task timing, for example, could be fairly ubiquitous, and correlates might exist on multiple dimensions. However, we found that removing the aligned imagery-unique dimension consistently and significantly reduced the maximum possible correlation that could be achieved from all other dimensions (P2, $P < 0.001$; P3, $P < 0.001$;

https://doi.org/10.1038/s41562-023-01804-5

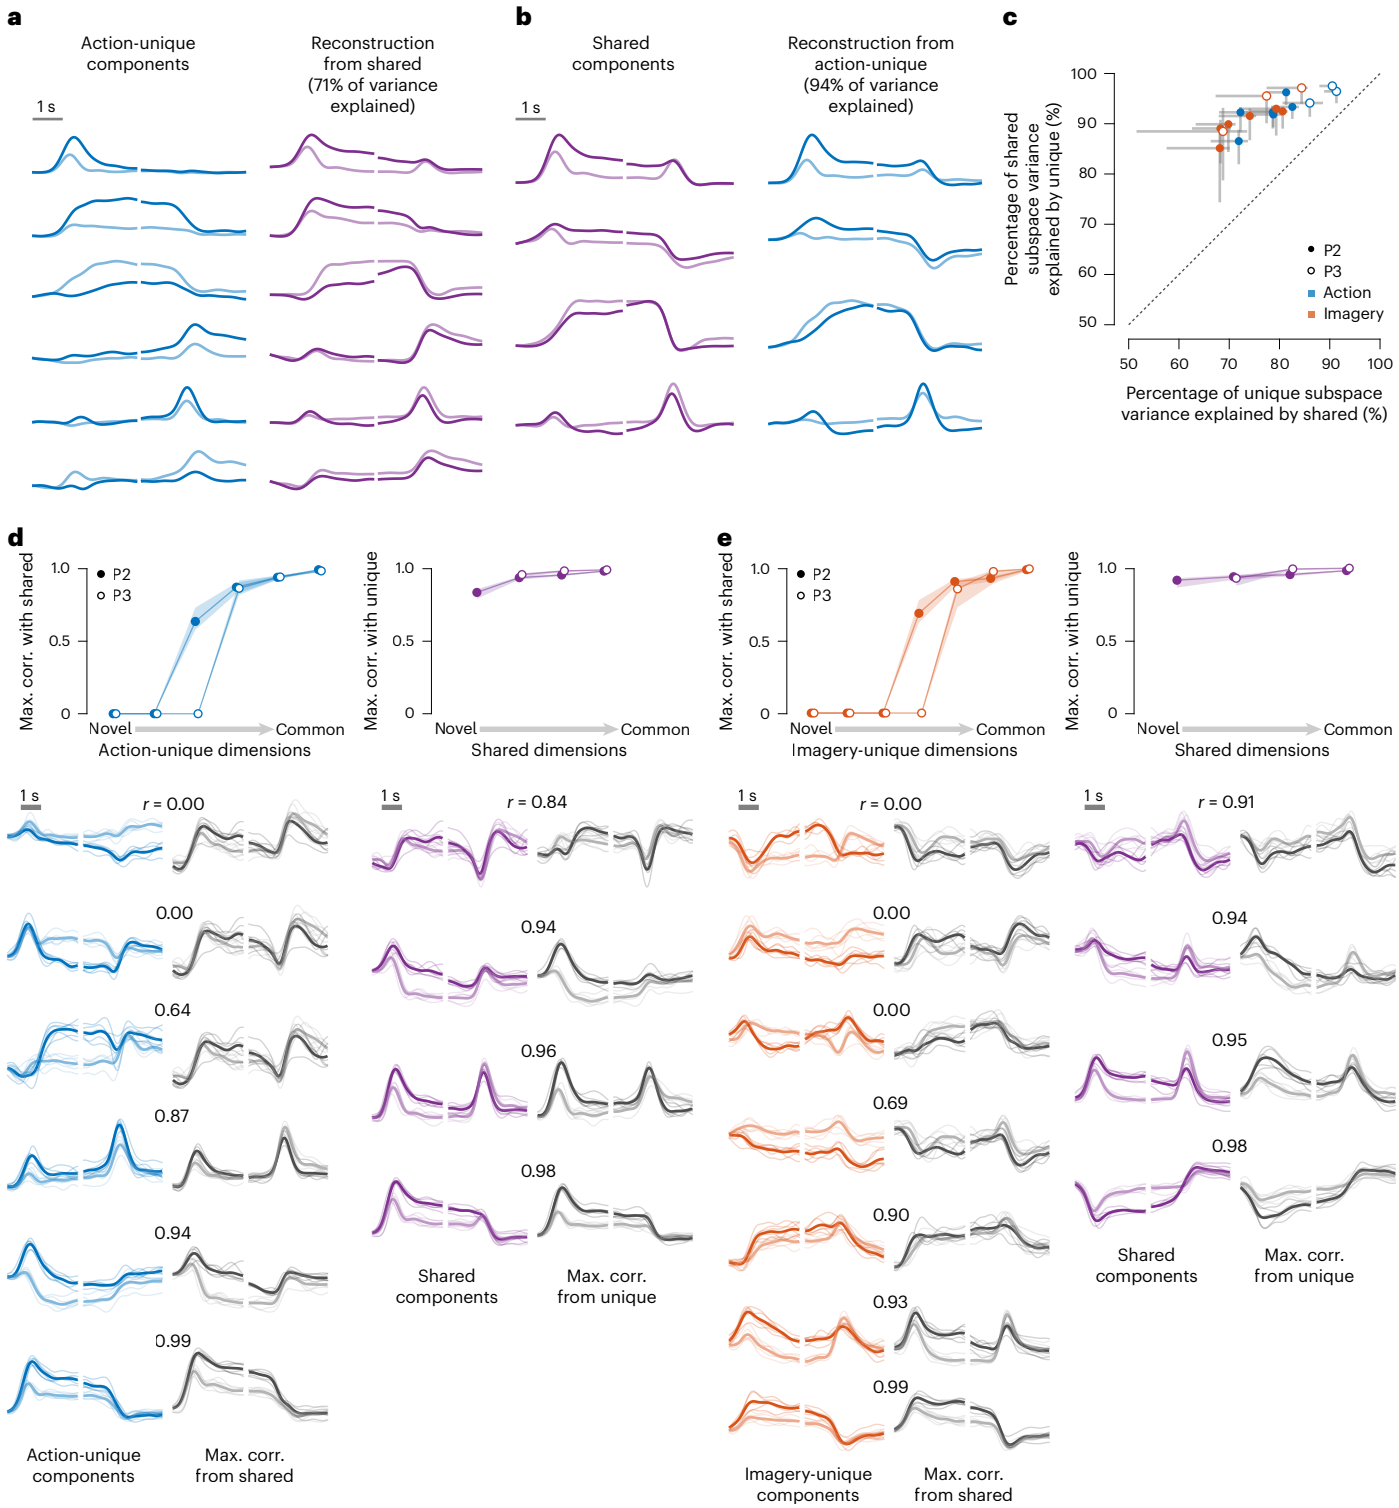

Fig. 4 | Unique subspaces contain novel dynamic features. a, Average action-unique subspace components for P2 (left) and the corresponding best linear reconstruction of each component from the shared subspace responses (right). **b**, Average shared subspace components for P2 (left) and the corresponding best linear reconstruction of each component from the action-unique subspace responses (right). **c**, The *x* axis shows the percentage of unique subspace variance than can be explained via linear combinations of shared subspace components. The *y* axis shows the percentage of shared subspace variance that can be explained via linear combinations of unique subspace components. Each dot corresponds to an individual session. Grey bars represent 95% CIs obtained via bootstrapping with 1,000 resamples. **d**, Left, action-unique dimensions ordered by increasing maximal correlation (max. corr.) within the shared subspace during action. Right, shared dimensions ordered by increasing maximal correlation within the action-unique subspace during action. The top plots show the ranked correlations for both participants, and the bottom plots show the actual component traces for P2. For each subplot below, the coloured traces on the left show the target subspace component, and the grey traces on the right show the maximally correlated response from the opposite subspace. **e**, As in **d** for the imagery-unique and shared subspaces during imagery. For the top sections of **d**,**e**, the shaded regions represent 95% CIs obtained via bootstrapping with 1,000 resamples. In the bottom sections of **d**,**e**, the thin traces correspond to individual sessions, and the thick lines correspond to cross-session averages. Light and dark traces represent low- and high-force averages, respectively.

paired *t*-test on 10,000 randomly chosen dimensions). As such, activity on aligned imagery dimensions correlated non-trivially with the corresponding action-unique dimension activity. This suggests that the same set of temporal components found in the action-unique subspace also exist during mental imagery, but within an orthogonal imagery-unique subspace.

Distinct unique subspace dynamics

Examining the components across subspaces in Fig. 3 revealed that, in addition to the high correlation between action-unique and imagery-unique subspaces, there appeared to be some similarity between responses in the unique subspaces and those in the shared subspace. A strong similarity between these subspaces could mean that the unique subspaces simply recapitulate responses from the shared subspace. We thus set out to determine the degree of this dynamic similarity between the unique and shared subspaces.

To assess the degree to which temporal components of activity in the unique subspaces simply recapitulated those from the shared subspace, we used linear regression to reconstruct unique subspace activity from the shared subspace (Fig. 4a). Similarly, we found a linear reconstruction of the shared subspace activity from the unique subspaces (Fig. 4b). If the unique subspaces were merely reflections of the shared activity, we would expect the quality of these reconstructions to be fairly equivalent. Instead, we found that the reconstructions of the unique subspaces were uniformly worse than the reconstructions of shared subspace activity from the unique subspaces (Fig. 4c). This argues against the possibility of the unique spaces being simple 'readouts' of the shared subspace; rather, they contain novel dynamic components.

To examine the novel dynamics within the unique spaces a bit further, we performed a second analysis in which we identified specific components that exemplified this dynamic novelty (see 'Assessing dynamic novelty' in Methods). The result of this analysis showed that for each unique subspace, there existed at least one dimension that had no correlate within the shared subspace (Fig. 4d,e). However, even the most novel response from the shared subspace could be fit relatively well from unique-subspace activity.

Unique subspace activity exhibits more complex dynamics

From the analysis in Fig. 4, it appeared that the unique subspace components that were least correlated with the shared subspace tended to display a large force-dependent effect. We suspected that the discrepancy in dynamics between the shared and unique subspaces arose largely from differences in the force-related information (Fig. 3 also suggests more pronounced and varied force-dependent effects in the unique subspaces than in the shared). To test this, we first isolated the force-dependent response in each subspace by subtracting the mean response across both force levels. We then estimated the dimensionality of this resulting force-dependent response in each subspace (see 'Force-specific responses' in Methods) and found that the dimensionality of the force-dependent activity in the shared subspace was lower than in the unique spaces (Fig. 5a,b)—though the difference was statistically significant in only one of the four cases on the basis of bootstrapping with 1,000 resamples (P2 action-unique versus shared, $P = 0.064$; P2 imagery-unique versus shared, $P = 0.17$; P3 action-unique versus shared, $P = 0.11$; P3 imagery-unique versus shared, $P < 0.001$). The unique subspaces appeared to contain both transient and tonic force-dependent responses, whereas the shared space seemingly only exhibited a single tonic response throughout the entirety of the trial duration (Fig. 5c,d).

Action-unique activity contains downstream motor commands

On the basis of the core property of the action-unique subspace—that it is active only during overt action—we hypothesized that it at least in part reflected communication of descending control signals. The second component displayed in Fig. 3c, for example, resembles the recorded forces across the action and imagery conditions. To more directly test the relationship between each subspace and the executed force, we attempted to decode moment-by-moment force from both the action-unique and shared subspaces (see 'Force decoding' in Methods). The example traces in Fig. 6a show the recorded force (black) and predictions (on held-out data) from the action-unique (blue) and shared subspace (purple) decoders for eight consecutive action trials from P2 (session five). Across all trials from all sessions, the action-unique subspace decoder outperformed the shared subspace decoder ($P < 0.001$; bootstrap across trials with 1,000 resamples), providing further evidence that the action-unique subspace is more closely linked to motor execution.

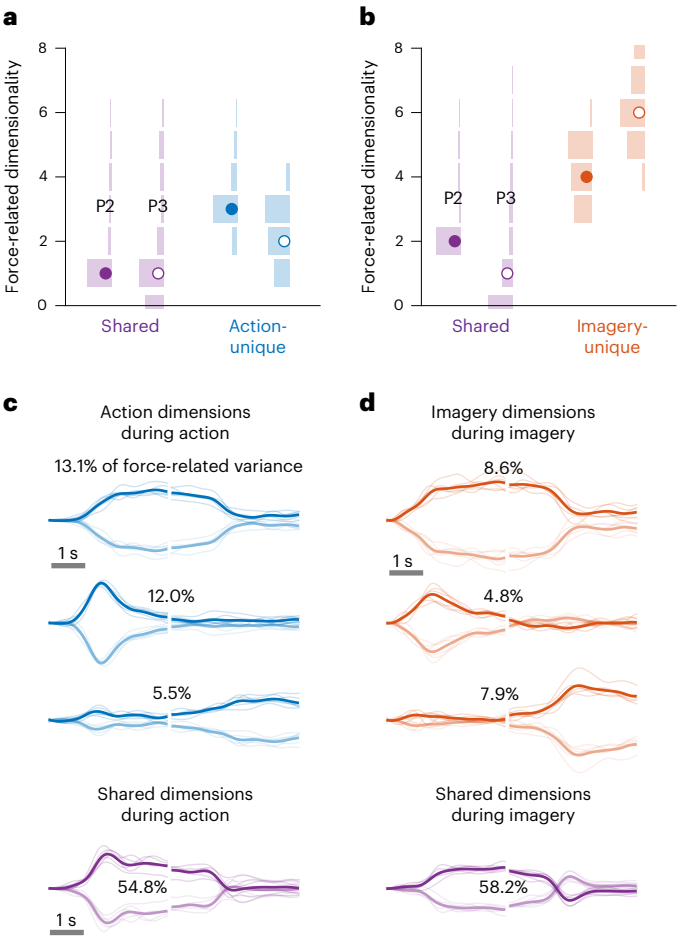

Fig. 5 | Action-unique and imagery-unique subspaces contain more complex force-specific responses. a, Estimates of force-related dimensionality in the shared and action-unique subspaces during action. **b**, Estimates of force-related dimensionality in the shared and imagery-unique subspaces during imagery. **c**, Leading force-dependent components (mean-subtracted) in the action-unique subspace (top) and shared subspace (bottom) during action (P2). **d**, Leading force-dependent components (mean-subtracted) in the imagery-unique subspace (top) and shared subspace (bottom) during imagery (P2). The shaded vertical histograms in **a**,**b** represent the distributions of dimensionalities for 1,000 threshold values from 0.5% to 2%.

Discussion

In this study, we examined the relationship between population activity in motor cortex during isometric force production and corresponding covert motor imagery. We found that the low-dimensional manifold activity comprised three orthogonal subspaces: a shared

https://doi.org/10.1038/s41562-023-01804-5

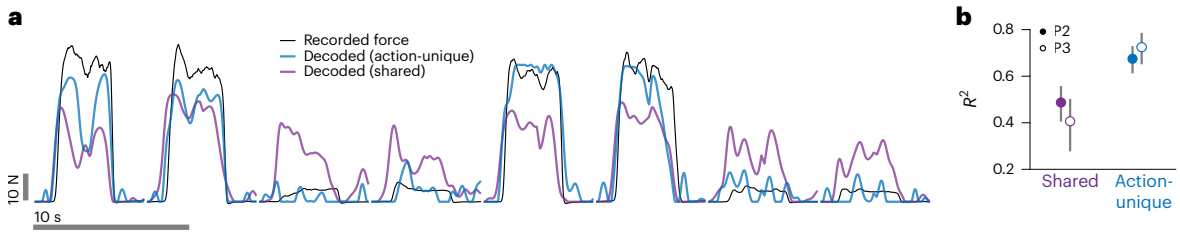

Fig. 6 | Moment-by-moment force can be decoded more accurately from the action-unique subspace than from the shared subspace. a, Eight example action trials (successive) from the fifth session by P2. The black traces correspond to the actual recorded wrist extension force. The blue and purple traces show cross-validated predictions from decoders (Wiener cascade) trained on action-unique and shared subspace responses, respectively. **b**, Total R^2 values for the action-unique and shared force predictions. The error bars represent the 95% CIs computed across all trials/sessions.

subspace, an action-unique subspace and an imagery-unique subspace. Activity in the shared subspace accounted for approximately half of the total variance and was nearly identical during both action and imagery. Activity in the action and imagery subspaces, though constructed from completely orthogonal correlation patterns, also contained well-matched sets of temporal responses. Furthermore, the action-unique and imagery-unique activity contained dynamically novel components over the shared space activity, and these novel components appear related to actual or imagined motor output (Supplementary Fig. 4).

Because the action-unique subspace we identified modulates only during action (and not during motor imagery), we hypothesized that it is directly involved in generating motor output and possibly receiving sensory feedback. During imagery, output dimensions must be avoided, and there is no incoming somatosensory feedback. In theory, motor cortex could satisfy the constraint of avoiding output dimensions by restricting imagery activity to a lower-dimensional subspace, such that there existed only a shared subspace and an action-unique subspace. However, we instead found that imagery engaged an additional, separate subspace, containing temporal components equivalent to those in the action-unique subspace. This suggests that covertly imagining an action does not simply suppress output activity but instead reorients into dimensions that do not generate muscle activity (Fig. 7). There are multiple potential explanations for why the act of motor imagery should involve creating 'dummy' output-related responses in motor cortex. One reason might be that during imagery, cortex is practising to generate output commands, even though those commands do not actually make it downstream. Even for this one-dimensional task, the action-unique subspace contained approximately six dimensions (Fig. 2f), which suggests an output-related space that is more complex than the eventual muscle activity[25]. In addition to muscle-like responses (Fig. 3c, second component), the action-unique subspace also contains transient responses (Fig. 3c, first component), which could reflect indirect control through subcortical areas. There is evidence that downstream motor structures are able to integrate brief, transient activity from cortex to generate sustained muscle output[26,27]. The ability to practise producing this multidimensional control signal within a motor-output-null space before generating the actual output commands might explain why mental rehearsal improves subsequent performance on overt motor tasks[3,4,6,28–30].

A separate possible reason for the existence of output-like components during imagery is that they are necessary for maintaining the dynamic structure of the entire motor cortical ensemble. There is ample evidence that motor cortex activity operates as a dynamic system[23,31,32]. The multidimensional population response unfolds predictably from an initial neural state, often dictated by preparatory activity in premotor cortex[16], presumably reflecting intrinsic motor cortical or broader synaptic connectivity. From our results, it appears that activity in the action-unique subspace is dynamically

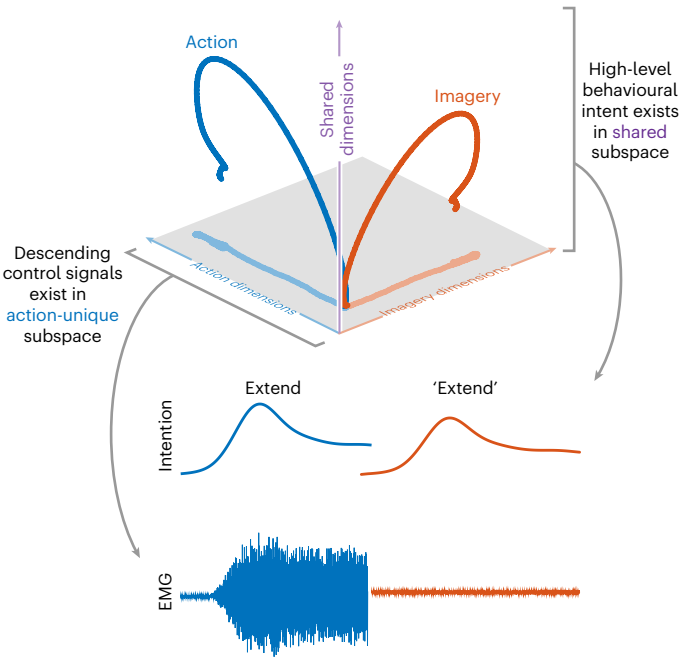

Fig. 7 | Hypothesized population-level architecture of motor cortex. Hypothesized population-level architecture of motor cortex that enables both action and imagery. Both action and motor imagery comprise similar sets of latent dynamics, which evolve in partially overlapping subspaces. We propose that the abstract motor goal exists within a consistent, shared subspace (likely as input from more frontal brain areas), while the commands related to descending control are confined to dimensions unique to volitional state.

distinct from activity in the shared subspace and thus is likely to be important for maintaining the dynamic structure of neural population activity (Fig. 4). Simply suppressing action-unique activity entirely during imagery would lead to the loss of this dynamic structure in motor cortex. However, recapitulating those components in an orthogonal subspace suppresses output while preserving its dynamic properties, which might be important for stabilizing the behaviour of the broader sensorimotor network across different volitional states.

The orthogonality between the action and imagery subspaces presumably functions similarly to the orthogonality observed between movement preparation and movement execution[16,17]. For both preparation and imagery, restricting the population activity to uniquely non-output dimensions prevents unwanted movement. However, the two processes are distinct in their dynamic relationship to the action. If enough time is allowed, preparatory activity in premotor areas appears to settle into a static neural state[33]. From the dynamic systems perspective, this represents a set point, which dictates how

the subsequent multidimensional response will unfold during movement execution[34]. The process of motor imagery, in contrast, is not an imminent preparation for movement but rather a rehearsal of the entire action. We did not observe strong preparatory responses in this experiment (perhaps due to array placement; Fig. 1b), so we do not know how preparation for imagery might relate to preparation for action. Uncovering the full population-level organization of preparation-, imagery- and action-related activity could help elucidate how cortex uses both overt and covert processes to improve motor skill.

The concept of covert motor imagery also invokes the related function of action observation. When people or animals observe others performing a motor action, it engages motor-related brain areas in a way similar to self-initiated movement[15,35–42]. This correlation between observation and action exists even in the activity of individual cortical neurons[12,35,42–44], which supports the notion of a 'mirror neuron' network through which the motor system can presumably learn new skills by observing others[45]. It is tempting to assume that observation and imagery/rehearsal are equivalent processes, and that observing an action triggers a person (or animal) to imagine performing the action themselves. Since the vast majority of intracortical observation-based experiments are performed with monkeys, it is often impossible to resolve the degree to which the animal is actively involved in motor imagery. However, recent work in humans with tetraplegia found that observation and imagery are actually not equivalent[10,11] and that it is possible to distinguish those two volitional states from population-level activity in motor cortex. Our observation of orthogonal subspaces containing action-unique and imagery-unique activity mirrors results from Jiang et al.[15], who observed separate subspaces containing action-unique and observation-unique activity. This suggests that although observation and imagery can be considered distinct volitional states[10,11], they might employ similar population-level mechanisms (that is, orthogonal subspaces for output and non-output conditions). Because the task used in our study was isometric, we could not include a meaningful observation-only condition. In the future, including kinematic limb movements could help identify the degree of overlap (common versus unique subspaces) across a larger range of volitional conditions, including observation, imagery, action and perhaps even replay during sleep[46].

While the task-dependent nature of the action and imagery subspaces provides clear insight into their functional roles (for example, the action-unique subspace includes output-related activity), interpretation of the shared subspace is more difficult. Activity in this subspace was nearly identical for both tasks, suggesting that it represents some sort of higher-level, abstracted task objective. The separation of force levels within the shared subspace argues against the interpretation that it is highly non-specific and reflects broad subject state processes such as arousal or engagement[47,48]. The shared responses also do not seem driven by visual feedback, as they consistently lead the executed force (and subsequent visual feedback; Supplementary Figs. 5 and 6). Instead, the shared subspace appears to contain information related to the specific task goal (that is, the force level). We speculate that activity in the shared subspace corresponds to goal-oriented or 'task-intention' signals from higher-order brain areas.

The shared subspace in particular may also be responsible for allowing the transfer of learning from covert practice (imagery) to actual motor performance. Work in monkeys found that activity in preparatory subspaces shared across volitional states can indeed help facilitate covert-to-overt learning of visuomotor rotations[49]. Similarly, a recent study in humans showed that movement preparation alone (without overt execution) can drive motor adaptation[50]. From an ethological perspective, the ability to imagine movements is useful only if it can meaningfully inform or assist overt motor control. Our results suggest that such transfer may be made possible in motor cortex by maintaining the full repertoire of population responses during imagery as is present during action. The responses related to motor output are reoriented into an orthogonal subspace, which allows the system to suppress actual motor output without changing the overarching dynamics. Further investigation of these faux-output responses—and their interactions with shared components—may give insight into how covert imagery can be used to drive skill learning or enhance rehabilitation following injury or disease.

Methods

Participants

Two participants (P2 and P3) took part in this study. The participants provided informed consent prior to performing any study-related procedures and were each compensated for the time committed to the ongoing clinical trial. The study was conducted under an Investigational Device Exemption from the Food and Drug Administration and approved by the Institutional Review Board at the University of Pittsburgh (Pittsburgh, Pennsylvania), registered at ClinicalTrials.gov (NCT01894802). The clinical trial is an early feasibility study with a primary outcome of evaluating the safety of an intracortical brain–computer interface for long-term neural recording and stimulation. The work presented here is a scientific effort to understand how neural activity during motor imagery, the basis of brain–computer interface devices, relates to that of overt movement. P2 is a 35-year-old man with tetraplegia caused by C5 motor/C6 sensory ASIA B spinal cord injury. P3 is a 30-year-old man with tetraplegia caused by incomplete C6/C7 ASIA B spinal cord injury. Both participants retain some residual upper arm and wrist control, but no hand function. The manual muscle test scores for wrist extension[51] were 4- for both participants (full range of motion against gravity/mild resistance). Using ASIA exam sensory testing of the right side[52], P2 had normal sensation from C2 to C4, altered sensation for C5, no sensation for C6 and altered sensation for C7 to T1. P3 had normal sensation from C2 to C6 and altered sensation from C7 to T2.

Both participants had two microelectrode arrays (Blackrock Microsystems, Inc., Salt Lake City, UT) implanted in the hand and arm areas of motor cortex (two 88-channel arrays for P2 and two 96-channel arrays for P3). They also had two 64-channel arrays implanted in somatosensory cortex[53], which were not used for this study. Data collection for P2 occurred approximately four years post-implant, and collection for P3 approximately one year post-implant.

Experimental setup

For each experimental session, the participants placed their pronated right hands on a board on their lap. We then secured a load cell in a frame attached to the board, positioning it such that it made gentle contact with the top of the hand. At the start of each session, we asked the participants to perform one maximal voluntary contraction. We then set initial low and high force targets based on the peak force observed during maximal voluntary contraction (10% and 60%) and asked the participants to practise by attempting each force level a few times. If the participant reported concern that the high force target was too high and he would be unable to perform the task without pain or fatigue, we lowered it until it reached a comfortable level. This resulted in average low forces of 5%, 8%, 4%, 4%, 4% and 4% (P2) and 11%, 11% and 7% (P3), and high forces of 50%, 66%, 50%, 47%, 48% and 48% (P2) and 51%, 52% and 38% (P3) maximal voluntary contraction. Once the force targets were set, we began the experiment, alternating blocks (12 trials each for P2 and 10 trials each for P3) of action and imagery. P2 performed three blocks of action and three blocks of imagery for all sessions. P3 performed four blocks of each for sessions one and two, and five blocks for session three. During both action and imagery blocks, the participants' hands remained positioned within the force-sensing apparatus (Fig. 1a). On each block, we randomly interleaved low and high forces. An audible cue ('gentle' or 'firm') cued the upcoming target, and then at the 'go' time, a red bar appeared at the target force level. The participants then attempted to achieve and maintain the target force,

using the line trace of exerted force as feedback. Each session always began with action to provide a reference for the subsequent imagery. We excluded trials that contained force traces that deviated substantially from the cued profile. On average, we excluded 6 ± 2 action trials and 8 ± 4 imagery trials (due to non-zero force output) per session.

Data acquisition

We collected neural data via digital NeuroPlex E headstages connected via fibre-optic cable to two synced Neural Signal Processors (Blackrock Microsystems). The neural signals were filtered using a fourth-order 250 Hz high-pass filter, logged as threshold crossings (−4.5 root mean square (RMS)) and subsequently binned at 50 Hz. These binned counts were then convolved offline with a Gaussian kernel (σ = 200 ms) to provide a smoothed estimate of firing rate.

Dimensionality reduction

We sought to reduce the dimensionality of the neural population recordings by projecting the activity into a low-dimensional space. PCA is a common approach for reducing dimensionality, but simply applying PCA to the combined action + imagery dataset could bias later results; the leading dimensions would preferentially capture action-related variance, since the action task contained higher overall variance. To ensure that the low-dimensional space encompassed both action and imagery responses equally, we first performed PCA separately on the condition-averaged firing rates for each dataset, which resulted in two orthonormal weight matrices, $W_{\text{action}} \in \mathbb{R}^{N \times D_{\text{action}}}$ and $W_{\text{imagery}} \in \mathbb{R}^{N \times D_{\text{imagery}}}$. Here N is the number of channels (N = 176 for P2 and N = 192 for P3), and D_{action} and D_{imagery} are the dimensions needed to capture at least 99% of the corresponding task-related variance. For all sessions of P2, D_{action} = [17, 16, 16, 15, 15] and D_{imagery} = [20, 19, 16, 21, 22, 20]. For P3, D_{action} = [19, 18, 17] and D_{imagery} = [31, 22, 31]. We then concatenated these two matrices into a new matrix $W_{\text{action+imagery}} \in \mathbb{R}^{N \times (D_{\text{action}} + D_{\text{imagery}})}$ and performed singular value decomposition:

$$W_{\text{action+imagery}} = U \Sigma V^{\text{T}}$$

This procedure provided $U \in \mathbb{R}^{N \times (D_{\text{action}} + D_{\text{imagery}})}$, an orthonormal basis spanning both W_{action} and W_{imagery}. The resulting space overestimated the actual dimensionality of the combined action and imagery tasks, since there ended up being a great deal of overlap between the two (Fig. 2). This approach represented the most conservative choice, allowing for mild dimensionality reduction (which aided in cross-session alignment and computational savings in future analyses) while ensuring that the space contained all (>99%) meaningful variance for both tasks. We projected all data from both tasks into this space by simply multiplying the firing rate estimates by U.

Alignment index shuffled control

To provide a comparison for the alignment index values computed between action and imagery, we performed a shuffled control analysis. For all trials of a given force condition (low or high), we randomly reassigned the action/imagery labels and recomputed the alignment index between the new scrambled 'action' and 'imagery' conditions. We repeated this 10,000 times for each dataset. Without exception, the shuffled control alignments were always higher than the action/imagery alignments (100%), suggesting that the amount of overlap between action and imagery was lower than what could be trivially explained by trial-by-trial variability.

Cross-session alignment

For each dataset, we performed PCA to reduce the dimensionality from channel space (176 channels for P2 and 192 channels for P3) to a lower-dimensional latent space (see 'Dimensionality reduction'). However, there is evidence that such low-dimensional (manifold) representations of the population activity remain consistent for a single behaviour[20]. To combine datasets from different sessions, it is necessary to first align the low-dimensional spaces[54]. We chose to align the low-dimensional latent activity across sessions for each participant using generalized Procrustes analysis[55]. Generalized Procrustes analysis iterates to find a multidimensional response common to all datasets and returns the axis transformation (an orthonormal rotation with uniform scaling) necessary to align each dataset with that common response. We found that this approach successfully aligned the responses, achieving a high degree of correlation across sessions (Supplementary Fig. 7). We used the cross-session average response to perform the subspace separation as described below.

Subspace separation

We aimed to identify, if possible, subspaces within the population activity that contained wholly task-specific variance (variance only during action or only during imagery). To achieve this, we implemented an optimization method that extends the concept of the alignment index[16] to identify orthogonal subspaces containing the 'unaligned' responses (that is, variance during one task that appears in the trailing PCs of the opposite task).

The intuition behind the approach is as follows. Given a dataset containing two tasks (for example, A and B), we can perform PCA on the neural data from just one of the tasks, $Z_{\text{A}} \in \mathbb{R}^{M \times N}$ (M time points and N latent dimensions) and identify the leading D_{potentA} dimensions that capture the vast majority of task A variance (we chose 99% as the cut-off, but other reasonable choices provide nearly equivalent results; Supplementary Fig. 8). This also results in $D_{\text{nullA}} = N - D_{\text{potentA}}$ dimensions, which combined contain insignificant (for example, <1%) variance for task A and can be considered task-A-null. However, we also can project activity from task B into this task A null space. If the null dimensions of task A contain a meaningful proportion of task B variance, we consider the activity within that subspace to be task-B-unique. We can perform an equivalent procedure starting with PCA on task B activity to identify task-A-unique activity.

The main challenge with this approach is that the null spaces from each task—that is, A_{null} (containing B-unique activity) and B_{null} (containing A-unique activity)—are computed by performing PCA separately on data from different task conditions, and so will not be mathematically orthogonal (due to noise introduced by low-variance components). However, we know that they are in fact functionally orthogonal. This is because A-unique activity resides in the potent space of task A (since it contains meaningful variance during task A) but the null space of task B. Likewise, B-unique activity resides in the null space of task A but the potent space of task B. Therefore, since A-unique exists in the potent space of A and B-unique exists in the null-space of A, they must be orthogonal. We were able to obtain orthogonal subspaces containing the unique responses through a simple optimization, as described below for our specific action/imagery case.

We begin with latent activity $L \in \mathbb{R}^{M \times D_{\text{latent}}}$, containing M time points and D_{latent} dimensions (see 'Dimensionality reduction'). For convenience, we also split this activity into task-specific data matrices $L_{\text{action}} \in \mathbb{R}^{M_{\text{action}} \times D_{\text{latent}}}$ (M_{action} time points from the action task) and $L_{\text{imagery}} \in \mathbb{R}^{M_{\text{imagery}} \times D_{\text{latent}}}$ (M_{imagery} time points from the action task).

We identified the imagery-null subspace $U_{\text{imagery-null}} \in \mathbb{R}^{D_{\text{latent}} \times D_{\text{imagery-null}}}$ by performing PCA on L_{imagery} and keeping only the $D_{\text{imagery-null}}$ trailing dimensions containing a total of <1% of L_{imagery} variance. We then performed PCA on $L_{\text{action}} U_{\text{imagery-null}}$ to obtain the subspace $V_{\text{action,imagery-null}} \in \mathbb{R}^{D_{\text{imagery-null}} \times D_{\text{action-unique}}}$ by keeping the leading $D_{\text{action-unique}}$ dimensions (discarding trailing dimensions that accounted for a total of <1% L_{action} variance). We then multiplied $U_{\text{imagery-null}}$ and $V_{\text{action,imagery-null}}$ to obtain a single orthonormal subspace $Z_{\text{action-unique}} \in \mathbb{R}^{D_{\text{latent}} \times D_{\text{action-unique}}}$, which contained meaningful variance during action and no meaningful variance during imagery. Projecting activity from both tasks, L, into this subspace gave data matrix

$Y_{\text{action-unique}} \in \mathbb{R}^{M \times D_{\text{action-unique}}}$, a representation of the action-unique responses. We then used the same method to obtain imagery-unique responses, $Y_{\text{imagery-unique}} \in \mathbb{R}^{M \times D_{\text{imagery-unique}}}$.

We used gradient descent via the Manopt toolbox[56] to identify two orthonormal unique subspaces $Q_{\text{action-unique}} \in \mathbb{R}^{D_{\text{latent}} \times D_{\text{action-unique}}}$ and $Q_{\text{imagery-unique}} \in \mathbb{R}^{D_{\text{latent}} \times D_{\text{imagery-unique}}}$, which minimized the sum of squared residuals between $[Y_{\text{action-unique}}, Y_{\text{imagery-unique}}]$ and $[L \bullet Q_{\text{action-unique}}, L \bullet Q_{\text{imagery-unique}}]$, subject to $Q_{\text{action-unique}} \perp Q_{\text{imagery-unique}}$ (the brackets indicate concatenation across dimensions). This optimization reconstructed the action-unique and imagery-unique responses but ensured that they were contained within orthogonal subspaces.

The optimization above resulted only in subspaces containing task-unique responses. We defined the remaining subspace not spanned by the combination of these two unique subspaces to be the shared space—that is

$$Q_{\text{shared}} \in \mathbb{R}^{D_{\text{latent}} \times (D_{\text{latent}} - D_{\text{action-unique}} - D_{\text{imagery-unique}})}$$

$$Q_{\text{shared}} \perp [Q_{\text{action-unique}}, Q_{\text{imagery-unique}}]$$

This shared subspace necessarily contains only dimensions for which meaningful variance exists during both tasks or no meaningful variance exists for either task. Together, all of the subspaces Q can be concatenated to form a single transformation $Q \in \mathbb{R}^{D_{\text{latent}} \times D_{\text{latent}}}$ that represents an orthonormal transformation of the original space. This approach to subspace identification contrasts with existing methods—for example, demixed PCA[57]—that identify subspaces containing cross-condition variance (Supplementary Fig. 9).

Feature extraction and subspace dimensionality estimation

With the latent responses split into shared and unique subspaces, we then found a compact representation of the underlying temporal components within each subspace using a varimax rotation. That is, we performed a varimax rotation on the baseline-centred multidimensional common response (cross-session average) for each subspace. We defined 'baseline' as the average response during a 200 ms window at the start of the trial. We then ranked the varimax-rotated version of the multidimensional response by variance. As an orthogonal transformation, this varimax procedure—like PCA—does not change the underlying nature of the multidimensional responses but rather highlights the separable temporal features within each subspace. For example, the first, second and fifth dimensions in the action subspace (Fig. 3c) are readily interpretable as onset, sustained and offset responses, respectively.

For dimensionality estimation, as in Fig. 2f,g, we calculated the number of dimensions within each subspace that explained more than 1% of the variance during the appropriate task (action or imagery) for all of 1,000 bootstrapped resamples of the trials. That is, for each bootstrapping run, we recalculated the condition means using the bootstrapped selection of trials (all sessions), aligned the within-participant averages and identified unique and shared subspaces. We performed a varimax rotation on each subspace response (we found that the ranked variances following the varimax rotation displayed a greater discontinuity in slope (that is, 'elbow') than did those resulting from PCA) and ranked the dimensions in decreasing order of variance. The reported dimensionality represents the number of dimensions for which the condition (action or imagery) variance exceeded 1% for all bootstrapped runs. True dimensionality is difficult to estimate and may differ markedly from the values reported here. However, here we were mostly interested in a qualitative comparison of dimensionalities across subspaces, rather than actual values. Different cut-offs and dimensionality estimation approaches provided similar main results—namely, that the unique subspaces consistently displayed higher dimensionality than did the shared subspace.

Action–imagery subspace alignment

Just as the low-dimensional subspace responses had to be aligned to make cross-session comparisons, so too did the action and imagery subspaces. We wanted to compare the similarity of the multidimensional temporal responses between the action and imagery subspaces. However, we could not simply compare, for example, the first dimension of the action subspace with the first dimension of the imagery subspace, since the two subspaces existed in orthogonal subspaces (see 'Subspace separation').

To compare the imagery and action subspace responses, we found an orthonormal transformation, $Z_{\text{im-act}}$, of the imagery subspace that maximized the sum of the squared covariance between action responses in the action subspace and imagery responses in the imagery subspace:

$$\max_{Z_{\text{im-act}}} \text{Tr}\left(\left(X_{\text{action}}^{\text{T}} X_{\text{imagery}} Z_{\text{im-act}}\right)^2\right)$$

We performed this optimization using the Manopt toolbox in MATLAB 2023b[56]. Unlike the more commonly used canonical correlation analysis for aligning temporal components[20,58], this rotation is orthonormal and represents a middle ground between maximizing correlation and returning leading components that explain a large amount of within-subspace variance.

Monte Carlo sampling

For the temporal response comparisons in Fig. 3a, we wanted to quantify the similarity of the multidimensional responses in a way that did not depend on the chosen coordinate frame. For example, we provide in Fig. 3b–d the specific correlation values for the displayed dimensions, but an orthogonal rotation of the space—which preserves the actual multidimensional relationship—would result in a different set of correlations.

To provide a coordinate-frame-agnostic quantification of two multidimensional responses, we performed a Monte Carlo sampling-based procedure in which we calculated the correlation between responses on randomly selected dimensions.

For multidimensional data X_1 and X_2 (dimensionality d):

(1) Generate a random d-dimensional unit vector $\mathbf{u}_{\text{rand}}$.
(2) Project X_1 and X_2 onto $\mathbf{u}_{\text{rand}}$ and calculate the correlation between the resulting projections $c_i = \text{corr}(X_1\mathbf{u}_{\text{rand}}, X_2\mathbf{u}_{\text{rand}})$.
(3) Repeat for 10,000 random vectors.

The resulting distribution of correlations c provides an overall picture of the multidimensional correspondence between X_1 and X_2.

Action–imagery correlation control distributions

The Monte Carlo sampling-based method provided a quantification of the overall correlation between responses in the action and imagery subspaces following the imagery–action alignment (see 'Action–imagery subspace alignment'). However, we also wanted to include an additional reference distribution that would help provide context for the resulting distribution of correlations (Fig. 3a, bottom). Ideally, we aimed to clarify whether the relatively high correlations observed between the action and imagery subspaces were unique to the specific alignment (indicating a true alignment of similar components), or whether they simply reflected broad, non-specific modulation throughout the multidimensional space as a result of, for example, task timing.

The core question that we addressed with our control procedure was: for any given dimension, what is the maximum possible correlation between action and imagery if we remove the aligned imagery dimension? To implement this, we performed an additional optimization on each draw of the Monte Carlo routine (see 'Monte Carlo sampling').

(1) Project imagery subspace activity X_{imagery} into the $(d-1)$-dimensional space orthogonal to $\mathbf{u}_{\text{rand}}$ (that is, in the null space of $\mathbf{u}_{\text{rand}}$) to obtain $X_{\text{im-V-null}}$:

$$\underset{\mathbf{m}_{\text{im-V-null}}}{\text{maximize}}\ \text{corr}(X_{\text{im-V-null}}\mathbf{m}_{\text{im-V-null}}, X_{\text{action}}\mathbf{u}_{\text{rand}})^2$$

(2) Find the unit vector $\mathbf{m}_{\text{im-V-null}}$ that maximizes the correlation between $X_{\text{im-V-null}}$ and $X_{\text{action}}\mathbf{u}_{\text{rand}}$.
(3) Save the resulting (positive) correlation:

$$c_{\text{null},i} = |\text{corr}\,(X_{\text{im-V-null}}\mathbf{m}_{\text{im-V-null}}, X_{\text{action}}\mathbf{u}_{\text{rand}})|$$

The distribution of c_{null} is therefore a strong control, since (unlike c) each element results from an independent optimization routine. However, even with that additional freedom, the values of c_{null} were consistently lower than those in c. This indicates that the temporal structure of action activity along any given action dimension is uniquely mirrored by imagery activity along the corresponding matched imagery dimension.

Assessing dynamic novelty

To determine the dynamic overlap between unique and shared subspaces (Fig. 4), we compared the extent to which the responses within one subspace could be fit via linear combinations of the responses in the other subspace. For example, in Fig. 4a, we performed a simple linear regression to find a best fit to the action-unique components from shared subspace responses during action. The condition-specific variance explained by those fits (action variance for action-shared fits, imagery variance for imagery-shared fits) is reported for each session in Fig. 4c. We obtained 95% CIs by bootstrapping over trials with 1,000 resamples.

To provide a more in-depth analysis of the novel dynamics that seemed to exist within the action- and imagery-unique subspaces, we ranked the components within each unique subspace with increasing maximal correlation with the shared subspace. To do this, we first performed the following optimization:

$$\max_{u} \frac{\sum \left(X_{\text{shared}}\left(X_{\text{shared}}{}^{\text{T}}X_{\text{shared}}\right)^{-1}X_{\text{shared}}{}^{\text{T}}X_{\text{act}} \bullet u - X_{\text{act}} \bullet u\right)^2}{\sum (X_{\text{act}} \bullet u)^2}$$

Intuitively, this finds the dimension u within the action-unique subspace activity X_{act} (mean-centred) that results in the poorest linear fit from X_{shared} to $X_{\text{act}} \bullet u$. The normal equation in the numerator identifies the optimal fit from X_{shared} to each projected component $X_{\text{act}} \bullet u$, and the maximization finds u for which that fit is least successful (accounts for the lowest variance). We then projected the action-unique activity X_{act} into the null space of u and repeated. By doing this, we were able to assemble a full orthonormal transformation of the action-unique activity in which each successive transformed dimension reflected the next most novel dynamic response (the poorest fit from X_{shared}). We performed this entire procedure for both unique subspaces as above, as well as for the shared subspace during action and imagery, optimizing with respect to the action-unique and imagery-unique subspaces, respectively.

Force-specific responses

To identify the force-specific responses within each subspace, we first subtracted the mean response across force conditions and then performed varimax on the resulting trajectories. We estimated the dimensionality in a similar way as for the full responses (see 'Feature extraction and subspace dimensionality estimation'). However, instead of using a cut-off based on the percentage of force-specific variance explained, we instead used the same actual variance cut-off as identified from the full condition response. That is, for the action condition, we calculated the number of force-specific action-unique and shared dimensions that explained at least 1% of the total action variance.

Force decoding

To probe the relationship between the action-unique/shared subspaces and executed force, we employed a simple Wiener cascade decoding model[59]. The Wiener cascade involves a simple linear model with a static nonlinearity and has been used previously to decode force and EMG responses[60]. The predictions were fully cross-validated using a leave-one-out approach for each session; that is, on each trial, the model used to predict force was trained on all other trials from that session. We then concatenated all of the cross-validated predictions to report total R^2.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Given the potential sensitivity concerns, deidentified data from this study are posted on DABI, a repository for data related to the National Institutes of Health Brain Research Through Advancing Neurotechnologies Initiative. The data for this specific sub-project can be found at https://doi.org/10.18120/70gm-a975 and are available upon request. A portion of the data included in this paper (action conditions only) was used in a previous publication[61].

Code availability

The code central to the results presented in this manuscript is publicly available at https://github.com/pitt-rnel/action_imagery.

References

1. Sirigu, A. et al. The mental representation of hand movements after parietal cortex damage. *Science* **273**, 1564–1568 (1996).
2. Sirigu, A. et al. Congruent unilateral impairments for real and imagined hand movements. *NeuroReport* **6**, 997–1001 (1995).
3. Clark, L. V. Effect of mental practice on the development of a certain motor skill. *Res. Q.* **31**, 560–569 (1960).
4. Frank, C., Land, W. M., Popp, C. & Schack, T. Mental representation and mental practice: experimental investigation on the functional links between motor memory and motor imagery. *PLoS ONE* **9**, e95175 (2014).
5. Ladda, A. M., Lebon, F. & Lotze, M. Using motor imagery practice for improving motor performance—a review. *Brain Cogn.* **150**, 105705 (2021).
6. Yue, G. & Cole, K. J. Strength increases from the motor program: comparison of training with maximal voluntary and imagined muscle contractions. *J. Neurophysiol.* **67**, 1114–1123 (1992).
7. Hotz-Boendermaker, S. et al. Preservation of motor programs in paraplegics as demonstrated by attempted and imagined foot movements. *NeuroImage* **39**, 383–394 (2008).
8. Jeannerod, M. *Motor Cognition: What Actions Tell the Self* (Oxford Univ. Press, 2006).
9. Kilteni, K., Andersson, B. J., Houborg, C. & Ehrsson, H. H. Motor imagery involves predicting the sensory consequences of the imagined movement. *Nat. Commun.* **9**, 1617 (2018).
10. Rastogi, A. et al. Neural representation of observed, imagined, and attempted grasping force in motor cortex of individuals with chronic tetraplegia. *Sci. Rep.* **10**, 1429 (2020).
11. Vargas-Irwin, C. E. et al. Watch, imagine, attempt: motor cortex single-unit activity reveals context-dependent movement encoding in humans with tetraplegia. *Front. Hum. Neurosci.* **12**, 450 (2018).
12. Cisek, P. & Kalaska, J. F. Neural correlates of mental rehearsal in dorsal premotor cortex. *Nature* **431**, 993–996 (2004).

13. Stephan, K. M. et al. Functional anatomy of the mental representation of upper extremity movements in healthy subjects. *J. Neurophysiol.* **73**, 373–386 (1995).
14. Aflalo, T. et al. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science* **348**, 906–910 (2015).
15. Jiang, X., Saggar, H., Ryu, S. I., Shenoy, K. V. & Kao, J. C. Structure in neural activity during observed and executed movements is shared at the neural population level, not in single neurons. *Cell Rep.* **32**, 108006 (2020).
16. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.* **7**, 13239 (2016).
17. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014).
18. Elsayed, G. F. & Cunningham, J. P. Structure in neural population recordings: an expected byproduct of simpler phenomena? *Nat. Neurosci.* **20**, 1310–1318 (2017).
19. Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural manifolds for the control of movement. *Neuron* **94**, 978–984 (2017).
20. Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A. & Miller, L. E. Long-term stability of cortical population dynamics underlying consistent behavior. *Nat. Neurosci.* **23**, 260–270 (2020).
21. Remington, E. D., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* **98**, 1005–1019.e5 (2018).
22. Sadtler, P. T. et al. Neural constraints on learning. *Nature* **512**, 423–426 (2014).
23. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
24. Kaufman, M. T. et al. The implications of categorical and category-free mixed selectivity on representational geometries. *Curr. Opin. Neurobiol.* **77**, 102644 (2022).
25. Russo, A. A. et al. Motor cortex embeds muscle-like commands in an untangled population response. *Neuron* **97**, 953–966.e8 (2018).
26. Shalit, U., Zinger, N., Joshua, M. & Prut, Y. Descending systems translate transient cortical commands into a sustained muscle activation signal. *Cereb. Cortex* **22**, 1904–1914 (2012).
27. Albert, S. T. et al. Postural control of arm and fingers through integration of movement commands. *Elife* **9**, e52507 (2020).
28. Ryan, E. D. & Simons, J. Efficacy of mental imagery in enhancing mental rehearsal of motor skills. *J. Sport Exerc. Psychol.* **4**, 41–51 (1982).
29. Schack, T., Essig, K., Frank, C. & Koester, D. Mental representation and motor imagery training. *Front. Hum. Neurosci.* **8**, 328 (2014).
30. Sheahan, H. R., Ingram, J. N., Žalalytė, G. M. & Wolpert, D. M. Imagery of movements immediately following performance allows learning of motor skills that interfere. *Sci. Rep.* **8**, 14330 (2018).
31. Churchland, M. M. et al. Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
32. Perich, M. G. et al. Motor cortical dynamics are shaped by multiple distinct subspaces during naturalistic behavior. Preprint at *bioRxiv* https://doi.org/10.1101/2020.07.30.228767 (2020).
33. Cisek, P. & Kalaska, J. F. Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–814 (2005).
34. Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* **68**, 387–400 (2010).
35. Dushanova, J. & Donoghue, J. Neurons in primary motor cortex engaged during action observation. *Eur. J. Neurosci.* **31**, 386–398 (2010).
36. Hari, R. et al. Activation of human primary motor cortex during action observation: a neuromagnetic study. *Proc. Natl Acad. Sci. USA* **95**, 15061–15065 (1998).
37. Holmes, P., Collins, D. & Calmels, C. Electroencephalographic functional equivalence during observation of action. *J. Sports Sci.* **24**, 605–616 (2006).
38. Muthukumaraswamy, S. D. & Johnson, B. W. Primary motor cortex activation during action observation revealed by wavelet analysis of the EEG. *Clin. Neurophysiol.* **115**, 1760–1766 (2004).
39. Papadourakis, V. & Raos, V. Neurons in the macaque dorsal premotor cortex respond to execution and observation of actions. *Cereb. Cortex* **29**, 4223–4237 (2019).
40. Rizzolatti, G., Fadiga, L., Gallese, V. & Fogassi, L. Premotor cortex and the recognition of motor actions. *Cogn. Brain Res.* **3**, 131–141 (1996).
41. Stefan, K. et al. Formation of a motor memory by action observation. *J. Neurosci.* **25**, 9339–9346 (2005).
42. Tkach, D., Reimer, J. & Hatsopoulos, N. G. Congruent activity during action and action observation in motor cortex. *J. Neurosci.* **27**, 13241–13250 (2007).
43. Mazurek, K. A., Rouse, A. G. & Schieber, M. H. Mirror neuron populations represent sequences of behavioral epochs during both execution and observation. *J. Neurosci.* **38**, 4441–4455 (2018).
44. Vigneswaran, G., Philipp, R., Lemon, R. N. & Kraskov, A. M1 corticospinal mirror neurons and their role in movement suppression during action observation. *Curr. Biol.* **23**, 236–243 (2013).
45. Rizzolatti, G., Fogassi, L. & Gallese, V. Neurophysiological mechanisms underlying the understanding and imitation of action. *Nat. Rev. Neurosci.* **2**, 661–670 (2001).
46. Rubin, D. B. et al. Learned motor patterns are replayed in human motor cortex during sleep. *J. Neurosci.* **42**, 5007–5020 (2022).
47. Hennig, J. A. et al. Learning is shaped by abrupt changes in neural engagement. *Nat. Neurosci.* **24**, 727–736 (2021).
48. Kaufman, M. T. et al. The largest response component in the motor cortex reflects movement timing but not movement type. *eNeuro* **3**, ENEURO.0085-16.2016 (2016).
49. Vyas, S. et al. Neural population dynamics underlying motor learning transfer. *Neuron* **97**, 1177–1186.e3 (2018).
50. Kim, O. A., Forrence, A. D. & McDougle, S. D. Motor learning without movement. *Proc. Natl Acad. Sci. USA* **119**, e2204379119 (2022).
51. Kendall, F. P., McCreary, E. K. & Provance, P. G. *Muscles: Testing and Function* 4th edn (Williams & Wilkins, 1993).
52. Kirshblum, S. C. et al. International standards for neurological classification of spinal cord injury (revised 2011). *J. Spinal Cord Med.* **34**, 535–546 (2011).
53. Flesher, S. N. et al. Intracortical microstimulation of human somatosensory cortex. *Sci. Transl. Med.* **8**, 361ra141 (2016).
54. Degenhart, A. D. et al. Stabilization of a brain–computer interface via the alignment of low-dimensional spaces of neural activity. *Nat. Biomed. Eng.* **4**, 672–685 (2020).
55. Gower, J. C. Generalized Procrustes analysis. *Psychometrika* **40**, 33–51 (1975).
56. Boumal, N., Mishra, B., Absil, P.-A. & Sepulchre, R. Manopt, a Matlab toolbox for optimization on manifolds. *J. Mach. Learn. Res.* **15**, 1455–1459 (2014).
57. Kobak, D. et al. Demixed principal component analysis of neural population data. *eLife* **5**, e10989 (2016).
58. Jude, J., Perich, M., Miller, M. & Hennig, M. Robust alignment of cross-session recordings of neural population activity by behaviour via unsupervised domain adaptation. In *Proc. 39th International Conference on Machine Learning, PMLR* **162**, 10462–10475 (2022).

59. Glaser, J. I. et al. Machine learning for neural decoding. *eNeuro* **7**, ENEURO.0506-19.2020 (2020).
60. Westwick, D. T., Pohlmeyer, E. A., Solla, S. A., Miller, L. E. & Perreault, E. J. Identification of multiple-input systems with highly coupled inputs: application to EMG prediction from multiple intracortical electrodes. *Neural Comput.* **18**, 329–355 (2006).
61. Balasubramanian, K., Arce-McShane, F. I., Dekleva, B. M., Collinger, J. L. & Hatsopoulos, N. G. Propagating motor cortical patterns of excitability are ubiquitous across human and non-human primate movement initiation. *iScience* **26**, 106518 (2023).

Acknowledgements

We thank N. Copeland and Mr. Dom for their continued efforts and commitment to this study. We also thank the research team, especially D. Harrington for regulatory management as well as C. Schoenewald, J. Ting, D. Sarma, A. Sethi and J. Weiss for their help with data collection. The research reported in this publication was supported by the National Institute of Neurological Disorders and Stroke of the National Institutes of Health under award numbers UH3NS107714 and U01NS108922. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

B.M.D. designed the task and collected the data. B.M.D., R.H.C. and S.M.C. contributed to the analysis. All authors contributed to the interpretation of the results. B.M.D. wrote the manuscript, with input from all authors.

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/s41562-023-01804-5.

Correspondence and requests for materials should be addressed to Jennifer L. Collinger.

Peer review information *Nature Human Behaviour* thanks Jonathan Michaels, Daniel O'Shea and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

nature portfolio

Corresponding author(s): Jennifer Collinger

Last updated by author(s): 10/31/2023

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Data were acquired using custom Matlab, Python, and C code.
Data analysis	Data analysis was completed with Matlab and is shared on GitHub.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

Provide your data availability statement here.

Research involving human participants, their data, or biological material

Policy information about studies with human participants or human data. See also policy information about sex, gender (identity/presentation), and sexual orientation and race, ethnicity and racism.

Reporting on sex and gender	Both subjects were male, which is consistent with the represented patient population (spinal cord injury, SCI) being approximately 80% male.
Reporting on race, ethnicity, or other socially relevant groupings	Given that there are only two participants, race and ethnicity are not reported as they are not central to the work being presented.
Population characteristics	Both participants were 28 at time of implant and one was 35 at the time of data collection, while the other was 30. No effects of age were expected or found. Both participants had tetraplegia due to cervical spinal cord injury (SCI), as required by the inclusion criteria.
Recruitment	Participants were recruited through voluntary research registries for people with SCI and through clinician referrals.
Ethics oversight	This study was conducted under an IDE from the FDA and with approval from the University of Pittsburgh IRB.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	Sample size of 2 subjects. Sample sizes for Brain-Computer Interface (BCI) studies in humans are limited due to the necessity of surgery and limited size of the eligible population. Most related papers only include 1 or 2 subjects.
Data exclusions	We excluded trials that contained force traces that significantly deviated from the cued profile. On average, we excluded 6 ± 2 action trials and 8 ± 4 imagery trials (due to non-zero force output) per session. This is described in the manuscript.
Replication	The results were replicated in two participants across multiple sessions.
Randomization	No randomization was possible for the intervention given the required surgery to implant electrodes.
Blinding	There was only a single group so the investigators were not blinded.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☐	☒ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Clinical data

Policy information about clinical studies

All manuscripts should comply with the ICMJE guidelines for publication of clinical research and a completed CONSORT checklist must be included with all submissions.

Clinical trial registration	NCT01894802
Study protocol	Trial protocol is unavailable as it is part of an ongoing study (IDE G130082).
Data collection	Data were collected in the Rehab Neural Engineering Labs at the University of Pittsburgh. Data were collected approximately 7 years after implant for P2, and approximately 2 years after implant for P3.
Outcomes	These results are are tangential to the primary and secondary outcomes of there overall clinical trial. Interim publication has been approved by the DSMB. The primary outcome of the ongoing trial is that the implant is safe for at least one year; all enrolled participants have exceeded this goal. The secondary outcome was functional use of the device; assessment of this outcome is still active.

Plants

Seed stocks	*Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.*
Novel plant genotypes	*Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.*
Authentication	*Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosiacism, off-target gene editing) were examined.*