PAPER

OPEN ACCESS

Neural subspaces of imagined movements in parietal cortex remain stable over several years in humans

L Bashford[1,2,*] ⓘ, I A Rosenthal[1] ⓘ, S Kellis[1] ⓘ, D Bjånes[1] ⓘ, K Pejsa[1], B W Brunton[3] ⓘ and R A Andersen[1]

[1] Division of Biology and Biological Engineering, and T&C Chen Brain-Machine Interface Center, California Institute of Technology, Pasadena, CA, United States of America
[2] Biosciences Institute, Newcastle University, Newcastle upon Tyne, United Kingdom
[3] Department of Biology, University of Washington, Seattle, WA, United States of America
* Author to whom any correspondence should be addressed.

E-mail: luke.bashford@newcastle.ac.uk

Abstract

Objective. A crucial goal in brain–machine interfacing is the long-term stability of neural decoding performance, ideally without regular retraining. Long-term stability has only been previously demonstrated in non-human primate experiments and only in primary sensorimotor cortices. Here we extend previous methods to determine long-term stability in humans by identifying and aligning low-dimensional structures in neural data. *Approach.* Over a period of 1106 and 871 d respectively, two participants completed an imagined center-out reaching task. The longitudinal accuracy between all day pairs was assessed by latent subspace alignment using principal components analysis and canonical correlations analysis of multi-unit intracortical recordings in different brain regions (Brodmann Area 5, Anterior Intraparietal Area and the junction of the postcentral and intraparietal sulcus). *Main results.* We show the long-term stable representation of neural activity in subspaces of intracortical recordings from higher-order association areas in humans. *Significance.* These results can be practically applied to significantly expand the longevity and generalizability of brain–computer interfaces.
Clinical Trials
NCT01849822, NCT01958086, NCT01964261

1. Introduction

Brain–machine interfaces (BMIs) decode neural activity to reproduce the user's intention and assist individuals with physical and neurological disabilities. In motor BMIs, the user commonly imagines or attempts to make a movement, and the corresponding recorded neural activity is decoded to guide movement in the intended direction, either on a computer or a prosthetic [1, 2]. BMIs can use neural signals acquired at different spatial and temporal resolution, but these have tradeoffs in performance and stability. Whereas single- or multi-unit recordings provide the highest information content, these recordings suffer from non-stationarity—different individual neurons are recorded from day to day or even from morning to afternoon [3, 4]. This variation is caused by

several factors, including movement of the electrodes, changes in the electrode–tissue interface, and degradation of the electrodes. Thus, as the neural features used as inputs to the decoder become more different from the initial training data, the performance of the BMI degrades over time. As BMIs are implanted for increasingly long durations [5, 6], the longitudinal stability of intracortically recorded neural activity is a central challenge to the practical utility of BMI devices. Currently, long-term use of BMI devices is only possible when users perform frequent retraining, often several times in a single day, to maintain desired performance. Identifying a signal with sufficient longitudinal stability could allow future devices to eliminate frequent retraining, and facilitate BMI use in cases where the loss of function over time may eventually prevent the performance of training tasks

(e.g. degenerative diseases such as amyotrophic lateral sclerosis/motor neuron disease ALS/MND).

Changes in the single unit activity recorded on each electrode of an array implanted in a target population that are being recorded over time is a common source of non-stationarity. However, the lower-dimensional subspaces of the same populations neural activity may remain relatively stable [7, 8]; we investigate this intriguing possibility in the context of BMI decoding. Alternative neural signals such as the local field potential have been observed to be more stable over time [9, 10]; however, the tradeoff is a reduction in information content compared to single unit activity, which ultimately limits performance. Therefore, the most promising solution currently being investigated is to use 'latent signals' for BMIs. Latent signals are derived from low-dimensional subspaces of the original high-dimensional single- or multi-unit neural activity (MUA), and they have been shown to preserve information content while minimizing non-stationarity [11–15]. Most stable latent spaces have so far been identified and validated primarily in longitudinal recordings of non-human primate primary sensorimotor cortices. In this paper, we investigate the longitudinal stability of these latent signals in two human participants, for whom neural signals were recorded from higher order areas in parietal cortex over several years of performing a BMI task [16, 17]. Specifically, we demonstrate that the neural subspace of imagined reaches in a center-out task remained remarkably stable in posterior parietal cortex.
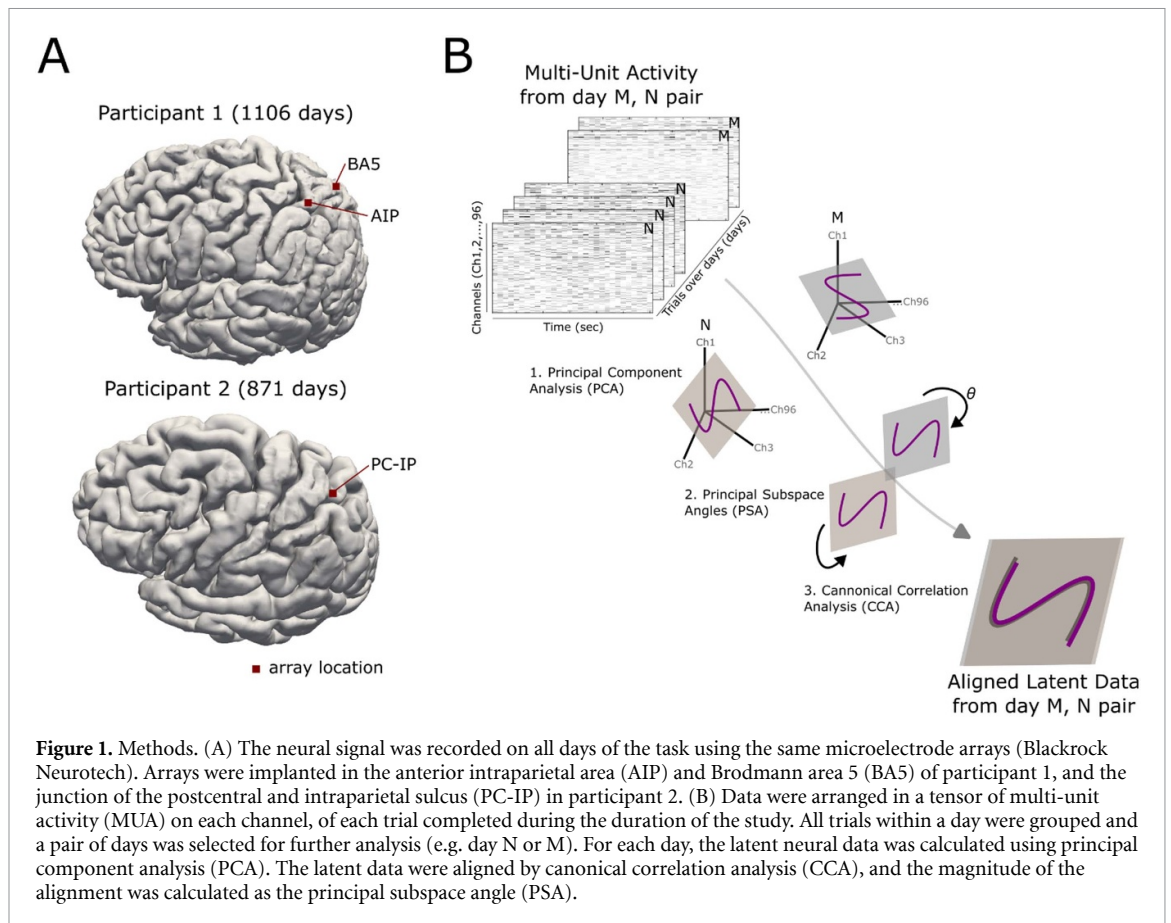
2. Results

Data were collected on 143 and 73 unique days, aggregated over a total period of 1106 and 871 d, for participant 1 [1] and 2 [18], respectively. Participant 1 attempted reaches in 4 directions while MUA was recorded from Brodmann Area 5 (BA5) and the Anterior Intraparietal Area (AIP). Participant 2 attempted reaches in 8 directions while MUA was recorded from the junction of the postcentral and intraparietal sulcus (PC-IP) (figure 1(A)). We use only the 'training' trials for longitudinal analysis, without any decoder present, to ensure the data were directly comparable from day-to-day [19]. During these trials, participants imagined moving their arm to follow the movement of an on-screen cursor. To process the neural data, we adapted methods established in non-human primates [12]. First a latent signal for each day on which the experiment occurred is calculated by performing Principal Component Analysis (PCA) [7] on all trials that day. The latent signal is then aligned for all pairs of days using Canonical Correlation Analysis (CCA) [20]. A Linear Discriminant Analysis (LDA) was used to classify the target locations (figure 1(B)). An LDA model was trained using data from Day N and tested within

day (N on N) using leave one out cross validation (LOOCV). This analysis was then repeated for every possible pair of training day N and testing day M. Further materials and methods information can be found in the Methods section of this manuscript.

When decoding the MUA signal, we observe good decoding accuracy (figure 2(A), MUA—red) within the same day, but this accuracy quickly degrades as the number of days between training and testing day increases (figure 2(A), MUA—black). Intriguingly, aligned latent neural activity space substantially improves the accuracy (figure 2(A), Latent). In particular, across all pairs of days, the decoding performance that can be achieved is higher from the latent signal (mean $\pm$ SD, AIP: $51.2 \pm 8.38\%$, BA5: $63.7 \pm 12.0\%$ and PCIP: $45.8 \pm 8.63\%$) than that achieved with MUA (mean $\pm$ SD, AIP: $35.4 \pm 11.1\%$, BA5: $34.6 \pm 12.1\%$ and PCIP: $25.5 \pm 11.8\%$) (all differences $p < 0.001$, Wilcoxon Sign Rank test, Bonferroni corrected) (figure 2(B)). Further, the across-day training produces a comparable performance compared to within-day using aligned latent data. Across all recording electrode arrays, the correlation between performance and time between the pairs is smaller for latent signals (AIP $r = -0.066$, BA5 $r = 0.020$, PCIP $= -0.033$, Pearson's linear correlation coefficient) than for MUA (AIP $r = -0.12$, BA5 r $= -0.30$ PCIP $r = -0.29$, Pearson's linear correlation coefficient). To summarize these results, we calculate the ratio of performance between all the within-day models and all the across-day models for latent and MUA activity. A ratio of 1 represents a comparable result, a value greater than one would mean that across-day pairs performed better than within-day pairs and vice versa for values below 1 (figure 2(C)). Here we see that in all cases the ratio of latent signals is higher than MUA demonstrating the aligned latent signal across days has significantly increased stability compared to the MUA activity ($p < 0.001$, Wilcoxon signed rank test, Bonferroni corrected. Participant 1 $N = 20\,449$, Participant 2 $N = 5329$). Decoding performance after only performing PCA, without CCA, leads to qualitatively the same results as the decoding performance on the raw MUA. This highlights the CCA step as the major contributor to the longitudinal performance of the latent aligned signal.

In this task, the participants were not required to learn anything novel, so we do not expect neural activity changes related to learning. However, the repetitive nature of the task and the amount of time spent performing it may still alter the brain activity representing the task over time. To investigate the way in which the task is represented in the latent signal over time, we calculated the Principal Subspace Angle (PSA) [21, 22] between all pairs of days in participant 1 (figure 3(A)). This angle reflects the magnitude to which the subspaces in the pair must be rotated to be maximally correlated, which is one way

Figure 1. Methods. (A) The neural signal was recorded on all days of the task using the same microelectrode arrays (Blackrock Neurotech). Arrays were implanted in the anterior intraparietal area (AIP) and Brodmann area 5 (BA5) of participant 1, and the junction of the postcentral and intraparietal sulcus (PC-IP) in participant 2. (B) Data were arranged in a tensor of multi-unit activity (MUA) on each channel, of each trial completed during the duration of the study. All trials within a day were grouped and a pair of days was selected for further analysis (e.g. day N or M). For each day, the latent neural data was calculated using principal component analysis (PCA). The latent data were aligned by canonical correlation analysis (CCA), and the magnitude of the alignment was calculated as the principal subspace angle (PSA).

to measure the extent to which CCA rotates the latent neural data to align day pairs [23]. A smaller PSA indicates a more similar representation between pairs. Because participant 1 had two arrays in different functional areas, we controlled for changes in the health of the array, which were not significantly different for the two arrays over time (see methods figure 1(C)). We divided up the data into early and late periods (methods figure 1(A)), where late was defined as the resumption of the experiment after a significant break (akin to a 'washout' in the familiarity of the task). We did not perform this analysis in participant 2 since there was no prolonged gap (>150 d) in data collection to compare early vs late day pairs (see methods). Over all pairs of days in participant 1, the variability in PSA in BA5 was smaller than for AIP (figure 3(B) left column). We then focused on close pairs of days within the early and late phases of data collection, those with a difference of only up to 10 d (figure 3(B) right column). Interestingly, after the break in the experiment, we find a significant difference in the representation of the data across these relatively close day pairs in only area AIP and not area BA5 ($p < 0.001$, Wilcoxon Signed Rank Test). This finding indicates a more stable intrinsic representation of reaching in area BA5 than AIP.

3. Discussion

Here we have demonstrated the stable representation of neural activity in subspaces of human intracortical recordings over several years. This result validates methods of aligning latent spaces developed in non-human primates during actual reaching [12, 24, 25], here applied successfully in classifying imagined reaches by humans. Furthermore, we have extended the finding of stable subspaces beyond the primary sensorimotor cortices into higher-order association areas in humans. The aligned latent signal performs best in decoding overall in each site, but the magnitude of the improvement reduces as the recordings come from more cognitive brain regions. We see the same effect in the PSA, where the variability in representation increases in the more cognitive brain region. We hypothesize that this trend is due to an increased flexibility in neuronal processing facilitating higher order/conceptual aspects of reaches in AIP compared to a more fundamentally engrained processing in the lower order sensorimotor control of a limb in BA5. Evidence for higher order control in parietal cortex has also recently been demonstrated in non-human primates. Recent comparisons of decoding accuracies in reaching tasks using signals from macaque V6A/PEc/PE areas, which show a

Figure 2. Classification decoder accuracy for two human participants. (A) Left column: the performance of the LDA classifier using the raw multi-unit activity (MUA) from the brain regions AIP (top), BA5 and PCIP (bottom). The red line shows the classification accuracy when the data from the same day (N on N) is used for training and testing with leave-one-out cross validation. The black line shows the classification accuracy when data from session day N is used for testing, but training is always performed on data from a single day (in this example day 1). Shading shows the standard deviation, dotted line shows the chance level for classifications. Right column: the same analysis as the left column but using the latent aligned data. Any breaks in data recording are shown in supplementary methods figure 1. (B) The decoding accuracy of every pair of days for MUA and Latent Aligned activity. (C) The ratio of performance between within day and across day decoding, error bars show standard deviations, stars indicate significance.

Figure 3. Principal subspace angles. (A) The principal subspace angle (PSA) calculated between all pairs of days for AIP (top) and BA5 (bottom). (B) Left column: the PSA between each pair of days, colored according to the early and late period (methods figure 1(A)) Right column: violin plots of only the day pairs where the difference between days is 10 or less, grouped into the early and late period.

lower decoding accuracy, indicating less information, is retrievable from PE with respect to more posterior cortex areas V6A/PEc [26].

There are few human intracortical BMI datasets available to validate the performance of latent signals over substantial periods of time in the same task. Participants enrolled in intracortical clinical trials are typically encouraged to perform a much broader range of tasks, with each requiring little to no training. Consequently, far fewer trials are available for any specific experimental paradigm. Chronic experiments exploring the human cortex offer a unique opportunity to study various changes in neural circuits over extended periods of time. Based on our results, we encourage the design of future studies to facilitate longitudinal task data collection and data collection from cortical sites beyond the traditionally used primary motor and sensory cortices. Our analysis sought to identify and validate stable signals that can be utilized longitudinally, a real time implementation of this method goes beyond the scope of this manuscript. A limitation of this method in real time implementation is the requirement for the entire data set of past and future data to be available. Alternative methods would be required for the stable signals identified here to be implemented in long-term BMIs with minimal retraining, however such methods are being developed [25]. Crucially, assumptions of signal stability based on the evidence presented here can be utilized in future studies to pave the way for BMIs to be used in many more cases; over longer periods, by individuals who lose the ability to retrain due to degenerative condition, or those who suffer injuries that preclude electrode implants in primary sensorimotor cortex. With the identification of such robust features, one promising direction for future work may be to enable the development of generalized BMIs that can be trained on data from individuals other than the eventual intended user [27].

4. Materials and methods

All procedures were approved by the Internal Review Boards of California Institute of Technology, University of Southern California, Rancho Los Amigos National Rehabilitation Center, University of California Los Angeles and Casa Colina Hospital and Centers for Healthcare. Informed consent was obtained from all participants after the nature of the study and possible risks were explained. This work was performed as part of Clinical Trials: NCT01849822, NCT01958086, NCT01964261.

4.1. Participants

Participant 1 was a 32 year-old tetraplegic male at the time of implantation. He was implanted with two microelectrode arrays on 17 April 2013. The electrodes were implanted in BA5 and the AIP. He had a complete lesion of the spinal cord at cervical level C3-4, sustained 10 years earlier, with paralysis of all limbs. Participant 2 was a 59 year-old tetraplegic female at the time of implantation. She was implanted with two arrays but only one was used in this study, at the junction of the post-central and intraparietal gyrus (PC-IP) on 29 August 2014. The other array was not functional. She had a C3-C4 spinal lesion (motor complete) sustained 7 years earlier, and retained movement and sensation in her upper trapezius, without control or sensation in her hands. During their enrollment, the participants performed many different tasks. The data for the analysis presented in this manuscript were collected on 143 and 73 unique days, over a period of 1106 and 871 d, for participant 1 and 2, respectively.

4.2. Task and data collection

The center out task was intended to allow the participants to spatially position a cursor on a computer screen. Targets were presented one at a time on the LCD display. The LCD monitor was positioned approximately 184 cm from the subject's eyes. Stimulus presentation was controlled using the Psychophysics Toolbox for MATLAB. During recording-only sessions, without any decoder, a circular cursor on the screen would move automatically from the center to one of either 4 (participant 1) or 8 (participant 2) targets arranged radially around the center point. Following a 250 ms delay relative to target onset, the cursor moved in a straight line directly to the target with an approximately bell-shaped velocity profile. Each trial lasted 3 s. The number of trials completed by the participants on each day during the study is shown in methods figures 1(A) and (B). Participants were asked to imagine making movements of the arm to mimic the movements observed on the screen.

The NeuroPort System (Blackrock Neurotech, UT, USA), comprising the arrays and neural signal processor (NSP), has received Food and Drug Administration (FDA) clearance for <30 d of acute recordings. For this study we received FDA Investigational Device Exemption clearance for extending the duration of the implant. The health and performance of the arrays was assessed as the mean impedance across all electrodes on each array, recorded on each day of the experiment. Impedance data is available for participant 1 only (methods figure 1(C)). MUA was amplified, digitized, and recorded at 30 kHz with the NeuroPort NSP. The threshold for calculation of MUA spikes was -3.5 * root mean squared voltage, calculated over each recording session. Data was organized into a three-dimensional tensor; The first dimension was the MUA binned into non-overlapping 50 ms windows. The second dimension was the number of electrodes (96 for each array). The third dimension was the index of trials ordered chronologically from the first to last over the entire study period.

4.3. Analysis

The analysis methods used in this manuscript extend the analysis of Gallego and colleagues [12] who demonstrated success in long-term decoding from primate motor cortex recordings. Data from each participant and each array was processed separately. The analysis was completed identically between all pairs of all days in which the participants completed the center out task. For the following description of the analysis two such days are represented as day M and day N.

Initially the same number of trials, containing equal presentation of all targets, are taken on each day. On days with different numbers of trials, we used all the trials on the day with the fewer trials and then randomly selected the same number of trials from the other day (A minimum of 5 and 6, and a maximum of 32 and 18 (participant 1 and participant 2 respectively) repetitions of each of the 4 targets were available from the individual day sessions). To ensure all the trials for a pair of days were used, the entire analysis was repeated 1000 times, each iteration using a different randomly selected set of trials. All electrodes (96) and all time bins were included for all trials. For each day this produced a (electrode $\times$ time $\times$ trials) matrix. The data was concatenated across trials and then dimensionality reduction was performed using PCA ('pca' function, Matlab 2021b). We reduced the data to 10 dimensions, following previous analysis, but confirmed that the results did not qualitatively change using a larger range of values. The result of the PCA analysis was a (10 $\times$ time*trials) matrix. We call this the 'latent data'. The latent data from each day in the pair was then aligned using CCA ('canoncorr' function, Matlab 2021b). This produced a (10 $\times$ time*trials) matrix. We call this the 'aligned latent data'. The data was then split back into individual trials (10 $\times$ time $\times$ trials) and the activity in each trial ('time' dimension) was averaged producing a (10 $\times$ 1 $\times$ trials) matrix for each day M and N. This was then used to calculate a linear regression model for classification ('fitlm' function, Matlab 2021b). The aligned latent data were used as the data and the target labels for each trial were used as the model. For within day calculations of classification accuracy a LOOCV was used to calculate classification accuracy. For calculating classification accuracy across days, the entire data from day M was used to train the LDA model, which was tested on the entire dataset from day N (and vice versa). To calculate the PSA we followed the method presented by Knyazev and colleagues [21] ('subspacea' function, MATLAB Central File Exchange, Matlab 2021b). We present data from the first PSA, but we note the results remained qualitatively consistent when summing over all PSAs calculated from the data.

We perform the PSA analysis on participant 1 due to: (1) Two electrode locations were recorded. In this case AIP and BA5 control for each other in factors related to changes in electrode–tissue interface that could influence the results (see supplementary methods figure 1(C)). We assume that since this crucial metric is consistent between the two arrays analytical differences can be explained by the different neurophysiology of the regions. (2) Two distinct periods of time in the experimental collection regime (early and late—supplementary methods figure 1(A)) can be identified. The break in recording was imposed by changing priorities in experimental data collection during the participants' enrollment. It can also be identified analytically with a threshold of at least 150 d elapsing with no task practice, and the average number of trials per day falling below 40. Because of differences in the data collection of the two subjects these conditions were not met in participant 2.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Code availability

All analyses were implemented using custom Matlab (The Mathworks Inc.) code. Code to replicate the main results are available upon reasonable request.

ORCID iDs

L Bashford ● https://orcid.org/0000-0003-4391-2491
I A Rosenthal ● https://orcid.org/0000-0002-9791-3820
S Kellis ● https://orcid.org/0000-0002-5158-1058
D Bjånes ● https://orcid.org/0000-0002-1208-5916
B W Brunton ● https://orcid.org/0000-0002-4831-3466

References

[1] Aflalo T *et al* 2015 Decoding motor imagery from the posterior parietal cortex of a tetraplegic human *Science* **348** 906–10

[2] Hochberg L R, Serruya M D, Friehs G M, Mukand J A, Saleh M, Caplan A H, Branner A, Chen D, Penn R D and Donoghue J P 2006 Neuronal ensemble control of prosthetic devices by a human with tetraplegia *Nature* **442** 164–71

[3] Dickey A S, Suminski A, Amit Y and Hatsopoulos N G 2009 Single-unit stability using chronically implanted multielectrode arrays *J. Neurophysiol.* **102** 1331–9

[4] Stevenson I H, Cherian A, London B M, Sachs N A, Lindberg E, Reimer J, Slutzky M W, Hatsopoulos N G, Miller L E and Kording K P 2011 Statistical assessment of the stability of neural movement representations *J. Neurophysiol.* **106** 764–74

[5] Simeral J D, Kim S-P, Black M J, Donoghue J P and Hochberg L R 2011 Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array *J. Neural Eng.* **8** 025027

[6] Hughes C L, Flesher S N, Weiss J M, Downey J E, Boninger M, Collinger J L and Gaunt R A 2021 Neural stimulation and recording performance in human sensorimotor cortex over 1500 days *J. Neural Eng.* **18** 045012

[7] Cunningham J P and Yu B M 2014 Dimensionality reduction for large-scale neural recordings *Nat. Neurosci.* **17** 1500–9

[8] Williamson R C, Doiron B, Smith M A and Yu B M 2019 Bridging large-scale neuronal recordings and large-scale network models using dimensionality reduction *Curr. Opin. Neurobiol.* **55** 40–47

[9] Flint R D, Wright Z A, Scheid M R and Slutzky M W 2013 Long term, stable brain machine interface performance using local field potentials and multiunit spikes *J. Neural Eng.* **10** 056005

[10] Flint R D, Scheid M R, Wright Z A, Solla S A and Slutzky M W 2016 Long-term stability of motor cortical activity: implications for brain machine interfaces and optimal feedback control *J. Neurosci.* **36** 3623–32

[11] Pandarinath C *et al* 2018 Inferring single-trial neural population dynamics using sequential auto-encoders *Nat. Methods* **15** 805

[12] Gallego J A, Perich M G, Chowdhury R H, Solla S A and Miller L E 2020 Long-term stability of cortical population dynamics underlying consistent behavior *Nat. Neurosci.* **23** 260–70

[13] Natraj N, Silversmith D B, Chang E F and Ganguly K 2022 Compartmentalized dynamics within a common multi-area mesoscale manifold represent a repertoire of human hand movements *Neuron* **110** 154–74.e12

[14] Trautmann E M *et al* 2019 Accurate estimation of neural population dynamics without spike sorting *Neuron* **103** 292–308.e4

[15] Elsayed G F and Cunningham J P 2017 Structure in neural population recordings: an expected byproduct of simpler phenomena? *Nat. Neurosci.* **20** 1310–8

[16] Andersen R A, Aflalo T, Bashford L, Bjånes D and Kellis S 2022 Exploring cognition with brain–machine interfaces *Annu. Rev. Psychol.* **73** 131–58

[17] Gallego J A, Makin T R and McDougle S D 2022 Going beyond primary motor cortex to improve brain–computer interfaces *Trends Neurosci.* **45** 176–83

[18] Aflalo T, Zhang C Y, Rosario E R, Pouratian N, Orban G A and Andersen R A 2020 A shared neural substrate for action verbs and observed actions in human posterior parietal cortex *Sci. Adv.* **6** eabb3984

[19] Orsborn A and Carmena J M 2013 Creating new functional circuits for action via brain-machine interfaces *Front. Comput. Neurosci.* **7** 157

[20] Bach F R and Jordan M I 2003 Kernel independent component analysis *J. Mach. Learn. Res.* **3** 1–48

[21] Knyazev A V and Argentati M E 2002 Principal angles between subspaces in an A-based scalar product: algorithms and perturbation estimates *SIAM J. Sci. Comput.* **23** 2008–40

[22] Björck Å and Golub G H 1973 Numerical methods for computing angles between linear subspaces *Math. Comput.* **27** 579–94

[23] Williams A H, Kunz E, Kornblith S and Linderman S 2021 Generalized shape metrics on neural representations *Advances in Neural Information Processing Systems* vol 34 (Curran Associates, Inc.) pp 4738–50

[24] Degenhart A D, Bishop W E, Oby E R, Tyler-Kabara E C, Chase S M, Batista A P and Yu B M 2020 Stabilization of a brain–computer interface via the alignment of low-dimensional spaces of neural activity *Nat. Biomed. Eng.* **4** 672–85

[25] Karpowicz B M, Ali Y H, Wimalasena L N, Sedler A R, Keshtkaran M R, Bodkin K, Ma X, Miller L E and Pandarinath C 2022 Stabilizing brain-computer interfaces through alignment of latent dynamics (https://doi.org/10.1101/2022.04.06.487388)

[26] Filippini M, Borra D, Ursino M, Magosso E and Fattori P 2022 Decoding sensorimotor information from superior parietal lobule of macaque via convolutional neural networks *Neural Netw.* **151** 276–94

[27] Peterson S M, Steine-Hanson Z, Davis N, Rao R P N and Brunton B W 2021 Generalized neural decoders for transfer learning across participants and recording modalities *J. Neural Eng.* **18** 026014