Months-long high-performance fixed LSTM decoder for cursor control in human intracortical brain-computer interfaces

Thomas Hosman*, Tsam Kiu Pun*, Anastasia Kapitonava, John D. Simeral†, Leigh R. Hochberg†

Abstract— **Intracortical brain-computer interfaces (iBCIs) enable high performance cursor control for people with tetraplegia by inferring motor intentions from neural recordings. However, current methods rely on frequent decoder recalibrations to reduce performance fluctuations attributable to instability in neural recordings. Towards clinical translation, iBCIs must sustain high performance over long periods of time with minimal interruptions to the user. Recent non-human primate (NHP) studies indicate that recurrent neural network (RNN) decoders are more robust to neural variability. Here, we demonstrate that an RNN variant, a long short-term memory (LSTM) neural decoder, provides online long-term, stable two-dimensional cursor control for a participant with tetraplegia enrolled in the BrainGate2 clinical trial. An LSTM decoder was trained with multiple days of the participant's historical intracortical motor cortex recordings spanning seventy days. The LSTM decoder was then fixed and evaluated online as the participant used the iBCI to control a computer cursor during a center out and back task for 15 sessions across four months. The LSTM demonstrated high performance for the first three months without recalibration or adaptive parameter updates with an average performance of 93.8% of targets acquired. This longitudinal study suggests that a nonlinear RNN-based decoder can provide stable, intuitive control of 2-D kinematics by humans with tetraplegia.**

Keywords—BCI, LSTM, Stable

* These authors contributed equally to this work.
† These authors contributed equally to this work.

** CAUTION: Investigational Device. Limited by Federal Law to Investigational Use. Research supported by the Office of Research and Development, Rehabilitation R&D Service, Dept of Veterans Affairs (N2864C, A2295R), NIH NIMH (T32MH115895), NIH NINDS (UH2NS095548), NIH NIDCD (U01DC017844), the Croucher Foundation. The content is solely the responsibility of the authors and does not necessarily represent the official views of NIH or the Department of Veterans Affairs or the United States Government. The MGH Translational Research Center has a clinical research support agreement with Neuralink and Synchron, for which LRH provides consultative input. MGH is a subcontractor on an NIH SBIR with Paradromics.

T. Hosman (thomas_hosman@brown.edu), T. K. Pun (tsam_kiu_pun@brown.edu), J. D. Simeral (john_simeral@brown.edu) and L. R. Hochberg (leigh_hochberg@brown.edu) are with the School of Engineering and Carney Institute for Brain Science, Brown University, Providence, RI. T.H., J.D.S., and L.R.H. are also with the VA Rehabilitation R&D Center for Neurorestoration and Neurotechnology, Dept. of VA Medical Center, Providence, RI.

L.R.H. is with the Center for Neurotechnology and Neurorecovery, Department of Neurology at Massachusetts General Hospital (MGH), Boston, MA, and with the Dept. of Neurology, Harvard Medical School, Boston, MA.

A. Kapitonava (ak4326@drexel.edu) was Center for Neurotechnology and Neurorecovery, Department of Neurology at Massachusetts General Hospital (MGH), Boston, MA. She is now with the Drexel School of Medicine, Philadelphia, PA.

I. INTRODUCTION

Intracortical brain-computer interfaces (iBCIs) have enabled individuals with tetraplegia to control external devices by decoding movement intentions from neural activity. Recent human iBCI research has demonstrated rapid decoder calibration [1] and achieved high bitrates for communication [2], [3]. However, a major challenge hindering clinical translation of iBCI systems is the inability to reliably decode recorded neural activity across various recording instabilities. The information-rich neural activity underlying consistent behaviors varies over time due to changes in recording conditions, individual's physiological and cognitive states, and other possible factors. Performance declines when the decoder is unable to adapt or be robust against such within- and across-day instabilities [4], [5]. Nonlinear recurrent neural network (RNN) decoders can achieve higher performance iBCI cursor control relative to traditional linear filters for nonhuman primates (NHPs) [6]–[8]. An LSTM network is a variant RNN architecture with multiplicative gates to better account for long-term temporal dependencies [9]. In offline simulations with chronic human intracortical recordings, LSTM-based RNNs trained on multiple days of data outperformed a state-of-the-art Kalman Filter in decoding accuracy and speed [10]. Non-linear RNNs trained on additional historical data theoretically have a higher modeling capacity to infer consistent underlying neural population dynamics and should be more robust against recording instabilities. Here, we demonstrate that an LSTM decoder can provide stable, high performance online cursor kinematic control by a participant with tetraplegia for up to three months without the need for explicit recalibration or other parameter updates.

II. METHODS

The Institutional Review Boards of Mass General Brigham, Brown University, and VA Providence Healthcare System granted permission for this study.

A. Participant

Intracortical neural activity was recorded from participant T11, a 37-year-old male enrolled in the Braingate2 pilot clinical trial** (NCT00912041) who had two 96-channel microelectrode arrays (Blackrock Microsystems, Salt Lake City, UT) placed in the dominant (left) hand/arm knob area of the precentral gyrus [11]. T11 has tetraplegia due to spinal cord injury (C4 AIS-B) that occurred approximately 11 years prior to study enrollment.

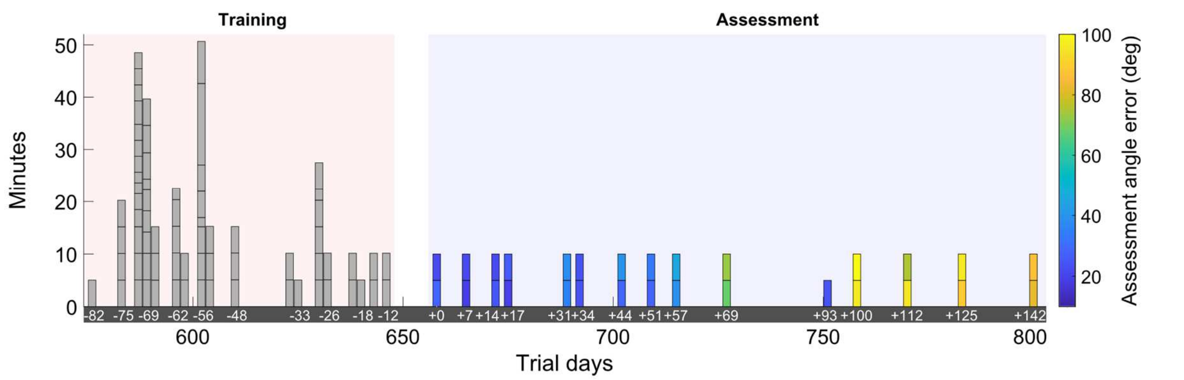

Fig. 1. Number of blocks in sessions and data duration for training and assessing the fixed LSTM. For each block on each trial day in the LSTM training data set (red region, gray bars), the participant completed one of several closed-loop point-and-select cursor tasks. Each block in the assessment phase (blue region) used the fixed LSTM on the same center-out-and-back cursor task. Each assessment block is colored based on the median angle error of the block.

B. *Study Design*

This study examined 33 sessions, spanning 224 days, in which T11 completed a variety of iBCI closed-loop point-and-select cursor tasks (a center-out-and-back task, a random target task, and a grid target task) in 2-5 minute blocks. T11 consistently attempted thumb joystick imagery to continuously move the neural cursor and imagined discrete digit actions that were decoded by a linear discriminant classifier [5] into clicks. Neural data from task blocks in the first part of this study (trial days 576-646, red region of Fig. 1) were used to train an LSTM kinematic decoder that was then repeatedly assessed in a closed-loop 2D cursor target task in sessions spanning trial days 658-800 (blue region of Fig. 1).

C. *Human Intracortical Neural Recordings*

Intracortical neural activity was acquired from 192 microelectrodes via a wireless broadband iBCI interface [12] as T11 completed a variety of closed-loop 2D cursor tasks. Two neural features (multi-unit threshold-crossing spike rates and power in the 250–5000 Hz spike band) were extracted from each electrode in 20ms non-overlapping bins. Each feature was z-scored, initially using its mean and variance from the previously recorded block of data and then adaptively updated using the mean and variance from a 3-minute rolling average window. For closed-loop 2D cursor control, the 384 normalized neural features were decoded into movement velocities by a real-time LSTM decoder.

D. *LSTM Decoder*

The LSTM decoder in this study followed the same architecture as our previous study [10] with adjusted training hyperparameters (Table I). The input, z-scored neural features, were passed directly to the RNN layer whose outputs went to three densely connected activation functions, decoding the x- and y-velocity and the normalized velocity.

The data used to train the decoder came from T11's 18 most recent 2D cursor sessions preceding the first assessment session. The training data spanned 70 days (trial days 576 to 646, Fig. 1) and contained 331 minutes across 8,441 acquisition trials.

TABLE I. LSTM TRAINING HYPERPARAMETERS

Hidden units	Batch size	Learning rate	Unrolled steps	# Features	Drop out	Loss
100	1024	0.0005	25	384	50%	M.S.E.

Individual trials with a median angle error (angle between the instantaneous cursor direction and the center of the target) greater than 50° were excluded from the training dataset. The labels (instantaneous cursor-to-target vectors) were normalized by their per-block 99th percentile to prevent outliers from compressing the label ranges while maintaining a magnitude range of roughly ±1.

For the entire LSTM training procedure, the data were randomly split by trials into training (70%) and validation (30%) datasets. LSTM training used the Adam optimizer and the Keras library (TensorFlow backend). The LSTM coefficients were updated using the training dataset (one epoch). After each epoch, the mean squared error (MSE) was assessed on the held-out validation data. LSTM training iterated until the held-out error did not improve.

E. *Assessment Task*

To assess longitudinal performance of the decoder for BCI kinematic control, T11 performed a closed-loop 2D point-and-select center-out-and-back assessment task using the same decoder for 15 sessions spanning 142 calendar days (trial days 658 to 800, Fig. 1). In each trial, T11 attempted thumb joystick imagery to continuously move the neural cursor from the center target to one of the eight pseudo-randomly selected peripheral targets. After the cursor reached a target, T11 attempted a right index finger down action to generate a "click" decoded by a linear discriminant analysis classifier calibrated at the beginning of each session. A trial was successful when the cursor contacted the cued target and a click action was decoded. Otherwise, a trial was unsuccessful after a 10-second timeout. Each assessment session collected two 5-min blocks except day 93 where only one 5-min block was collected. A total of 1840 assessment trials were collected during 145 minutes of cumulative task time. Neural recordings, cursor-to-target vectors and decoder outputs were logged.

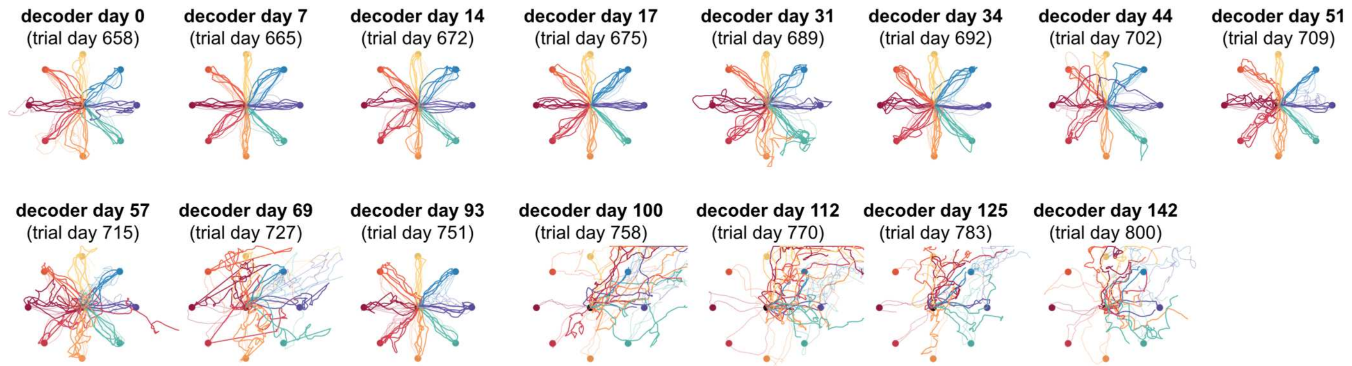

Fig. 2. Cursor trajectories of the first 5-min block for each trial day. Colors correspond to different target locations. A solid line indicates the cursor trajectory of a trial going from center to a peripheral target, and a dotted line indicates cursor trajectory of a trial of return to the center.

F. *Outlier Events*

Changes in the acquired neural data can degrade the performance of a static BCI decoder. Such perturbations may result from slow drifts in neuronal tuning or changes in task context, but also from rapid within trial-changes due to noise or dropped wireless neural data packets. To examine the potential influence of acute perturbations on task performance, we identified outlier trials as those having more than 5% loss of wireless neural data packets or large neural responses greater than 8 standard deviations from the mean [12]. Across the 15 sessions, 34 of the 1840 trials were labeled as outlier trials.

G. *Performance Metrics*

The overall task performance on each day was quantified by the percent of successful trials in the assessment task. Cursor control was further assessed using the following performance metrics [13]: *angle error (AE)*, defined as the angle between the cursor-to-target vector and the directional vector estimated by the decoder (best = 0°; max = 180°); *time to target*, which is the time it takes to acquire the target before timeout*; path efficiency*, which is the ratio of the actual trajectory length to the ideal straight-line path to the target (best = 1); and the number of *orthogonal direction changes (ODC)* where the cursor reversed away from or toward the target (best = 0), which quantifies the path consistency towards the target. These performance metrics were computed for each trial on all LSTM performance assessment days.

To quantify how the LSTM's performance over days related to changes in the underlying neural activity, we first applied demixed principal component analysis (dPCA) [14] to the population neural activity from the initial LSTM assessment session to find the top two task-relevant neural dimensions. Then, for each assessment day, we projected the neural data from the initial second of each trial (excluding outlier trials) onto these day 0 neural components and visualized the resulting average trajectories. We further quantified the amount of task-related neural activity in each assessment session by comparing the variance accounted for (VAF) by these initial task-related neural components.

III. RESULTS

Over the course of 142 days, we assessed a fixed LSTM decoder's ability to enable control of an iBCI cursor without recalibration. Cursor trajectories in the closed-loop center-out task each day reflected the quality of control provided by the fixed LSTM decoder over more than 4.5 months of assessment (Fig. 2). The LSTM decoder achieved accurate online cursor control in all but one session during the first three months. The decoder achieved an average of 93.8% success in the first 11 sessions (93 decoder days). However, there were signs of decreasing performance on decoder day 69. We found consistently poor performance from day 100 to 142 (average of 33.1% targets acquired). Interestingly, as the LSTM performance worsened, its errors were directionally biased consistently to the upper right direction (Fig. 2, decoder day 100-142) as opposed to displaying errors in all directions.

The computed cursor performance metrics over time (Fig. 3) were consistent with the observed cursor trajectories. These performance metrics revealed significant changes between early and later assessment sessions, including *angle error* (early: 26.8° ± 22.6°; later: 88.4° ± 46.1°; $p < 0.001$; Wilcoxon rank sum), *time to target* (early: 4.02 ± 1.98 seconds; later: 8.01

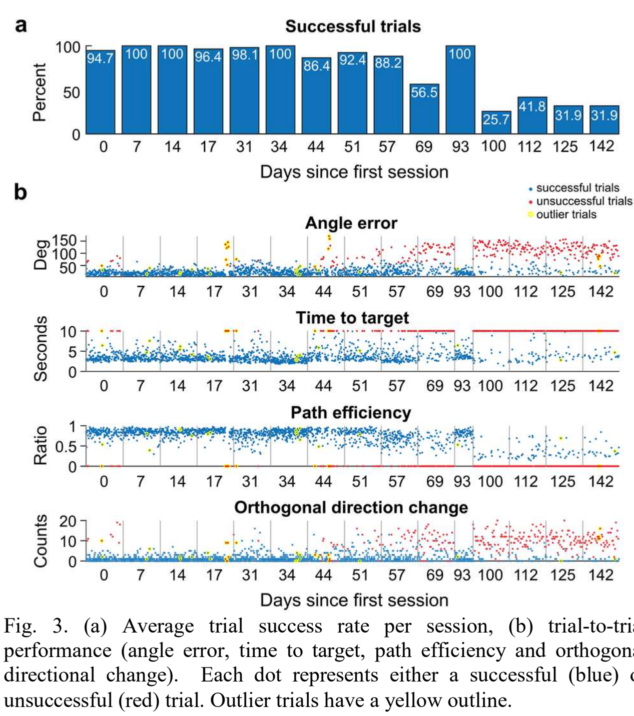

Fig. 3. (a) Average trial success rate per session, (b) trial-to-trial performance (angle error, time to target, path efficiency and orthogonal directional change). Each dot represents either a successful (blue) or unsuccessful (red) trial. Outlier trials have a yellow outline.

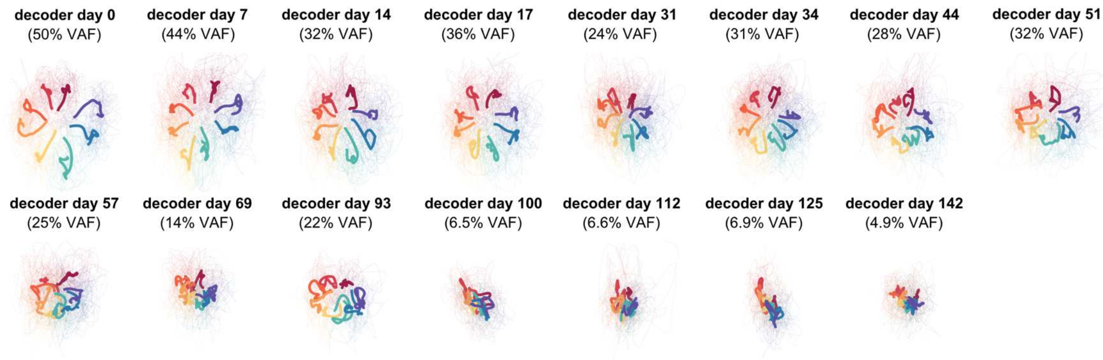

Fig. 4. Projection of neural features of assessment sessions onto the top two task-dependent demixed principal components of decoder day 0. Color represents movement intention directions. Solid lines are trial averages per direction.

± 3.04 seconds; $p < 0.001$), *path efficiency* (early: 0.79 ± 0.12; later: 0.38 ± 0.17; $p < 0.001$; excluding unsuccessful trials) and *orthogonal directional change* (early: 1.82 ± 2.82; later: 7.40 ± 5.35; $p < 0.001$).

How does the LSTM's performance over days compared to linear estimates of change in the neural data? We examined the top two task-relevant dPCs computed from the neural data in the first assessment session to compare changes in task-relevant neural activity across sessions. We visualized this change over days by projecting the neural data from the beginning of each trial to observe the average neural trajectory per target direction for each session (Fig. 4). We found that the average trajectories became less distinct as the session dates progressed, reflecting changes in the underlying population activity over time consistent with the decline in task performance. We then computed the VAF from the top two task-relevant neural dPCs for each session. We found that VAF was initially 50% on decoder day 0 and remained above 20% on days in which clear target trajectories were observed. As the task performance declined, the VAF decreased, ending with a value of 4.2% on decoder day 142. We found a negative correlation between the neural VAF and the average angle error per session (-0.88, Pearson correlation; similar correlation found when using the top two principal components, -0.92).

IV. Discussion & Conclusion

An LSTM decoder trained on previously collected human intracortical recordings offered stable and accurate 2D closed-loop iBCI cursor control for a person with tetraplegia for up to three months. The stable long-term performance reported may be attributed to the nonlinear recurrent LSTM decoder. Previous work has shown that a nonlinear, recurrent neural network decoder can enable robust 2D cursor decoding in NHPs over several days [8]. Furthermore, stable low dimensional neural manifolds have been demonstrated to persist for hundreds of days [15]. Training the LSTM on a large number of historical sessions may implicitly regularize its coefficients toward a stable, low dimensional manifold, enabling better across-session neural-to-kinematic mappings despite neural and task variability. In other words, it avoids overfitting to a given day of data, unlike traditional daily recalibrated decoders, thereby becoming more robust to across-day instabilities. Including various types of 2D closed-loop cursor tasks in the training data also encompasses a wider range of cursor control behavior, which allows the LSTM to adapt to different control contexts. We previously reported that peak offline LSTM performance can be achieved by training with human neural data spanning 73 days on average [10], which is consistent with the training data set we used for the LSTM in this study.

We assessed a fixed LSTM over 15 sessions across 142 days and found robust high performance decoding up to 93 days after the initial assessment. However, the LSTM had low performance that did not improve 100 days and beyond. Comparing angle error with the VAF of the top two task-relevant dPCs revealed a strong negative correlation between LSTM performance and variance explained. This suggests that the neural subspace maintained a strong task relationship over the first 93 days of LSTM use (except day 69). Future work will investigate if this relationship holds true for linear decoders.

The decline of performance during decoder day 69 followed by a recovery of performance on decoder day 93 suggests that, although there may be a period of several days in which the decoder's performance is variable, the task-relevant neural subspace was recoverable. Follow-up studies will sample performance more often to better understand the transition between high performance and lower performance decoding, perhaps providing additional insight into the nature of the underlying nonstationarity.

Our findings validate the use of a nonlinear LSTM decoder as a strategy towards long-term BCI decoding with a greatly reduced need for recalibration. Large volumes of previously collected data can be useful for training decoders that can accommodate neural variability. It is likely that other pre-processing steps such as nonlinear data alignment may help the LSTM decoder to better generalize and adapt to larger nonlinear shifts in neural representation.

Acknowledgment

The authors would like to thank participant T11, his family and caregivers, and Beth Travers, Dave Rosler, and Maryam Masood for their contributions to this research.

References

[1] D. M. Brandman *et al.*, "Rapid calibration of an intracortical brain–computer interface for people with tetraplegia," *J. Neural Eng.*, vol. 15, no. 2, p. 026007, Jan. 2018, doi: 10.1088/1741-2552/AA9EE7.

[2] C. Pandarinath *et al.*, "High performance communication by people with paralysis using an intracortical brain-computer interface," *Elife*, vol. 6, Feb. 2017, doi: 10.7554/eLife.18554.

[3] F. R. Willett, D. T. Avansino, L. R. Hochberg, J. M. Henderson, and

K. V Shenoy, “High-performance brain-to-text communication via handwriting,” *Nature*, vol. 593, no. 7858, pp. 249–254, 2021, doi: 10.1038/s41586-021-03506-2.

[4] J. A. Perge *et al.*, “Intra-day signal instabilities affect decoding performance in an intracortical neural interface system.,” *J. Neural Eng.*, vol. 10, no. 3, p. 036004, Jun. 2013, doi: 10.1088/1741-2560/10/3/036004.

[5] B. Jarosiewicz *et al.*, “Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface,” *Sci. Transl. Med.*, vol. 7, no. 313, pp. 313ra179-313ra179, Nov. 2015, doi: 10.1126/scitranslmed.aac7328.

[6] P. H. Tseng, N. A. Urpi, M. Lebedev, and M. Nicolelis, “Decoding movements from cortical ensemble activity using a long short-term memory recurrent network,” *Neural Comput.*, vol. 31, no. 6, pp. 1085–1113, Jun. 2019, doi: 10.1162/neco_a_01189.

[7] B. Premchand *et al.*, “Decoding movement direction from cortical microelectrode recordings using an LSTM-based neural network,” in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2020, pp. 3007–3010. doi: 10.1109/EMBC44109.2020.9175593.

[8] D. Sussillo, S. D. Stavisky, J. C. Kao, S. I. Ryu, and K. V. Shenoy, “Making brain–machine interfaces robust to future neural variability,” *Nat. Commun.*, vol. 7, p. 13749, Dec. 2016, doi: 10.1038/ncomms13749.

[9] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, doi: 10.1162/neco.1997.9.8.1735.

[10] T. Hosman *et al.*, “BCI decoder performance comparison of an LSTM recurrent neural network and a Kalman filter in retrospective simulation,” *Int. IEEE/EMBS Conf. Neural Eng. NER*, vol. 2019-March, pp. 1066–1071, May 2019, doi: 10.1109/NER.2019.8717140.

[11] L. R. Hochberg *et al.*, “Neuronal ensemble control of prosthetic devices by a human with tetraplegia,” *Nature*, vol. 442, no. 7099, pp. 164–171, Jul. 2006, doi: 10.1038/nature04970.

[12] J. D. Simeral *et al.*, “Home use of a percutaneous wireless intracortical brain-computer interface by individuals with tetraplegia,” *IEEE Trans. Biomed. Eng.*, vol. 68, no. 7, pp. 2313–2325, Jul. 2021, doi: 10.1109/TBME.2021.3069119.

[13] S.-P. Kim, J. D. Simeral, L. R. Hochberg, J. P. Donoghue, G. M. Friehs, and M. J. Black, “Point-and-click cursor control with an intracortical neural interface system by humans with tetraplegia.,” *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 19, no. 2, pp. 193–203, Apr. 2011, doi: 10.1109/TNSRE.2011.2107750.

[14] D. Kobak *et al.*, “Demixed principal component analysis of neural population data,” *Elife*, vol. 5, Apr. 2016, doi: 10.7554/eLife.10989.

[15] J. A. Gallego, M. G. Perich, R. H. Chowdhury, S. A. Solla, and L. E. Miller, “Long-term stability of cortical population dynamics underlying consistent behavior,” *Nat. Neurosci.*, vol. 23, no. 2, pp. 260–270, Feb. 2020, doi: 10.1038/s41593-019-0555-4.