

# Multi-dimensional Neural Decoding with Orthogonal Representations for Brain-Computer Interfaces

Kaixi Tian<sup>1,2,3</sup>, Shengjia Zhao<sup>4</sup>, Yuhan Zhang<sup>1,5</sup>, Shan Yu<sup>1,2,3\*</sup>

<sup>1</sup>Laboratory of Brain Atlas and Brain-inspired Intelligence, Institute of Automation, Chinese Academy of Sciences

<sup>2</sup>School of Future Technology, University of Chinese Academy of Sciences

<sup>3</sup>State Key Laboratory of Brain Cognition and Brain-inspired Intelligence, CAS Center for Excellence in Brain Science and Intelligence Technology

<sup>4</sup>Department of Physics, University of Oxford

<sup>5</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences

tiankaixi2020@ia.ac.cn, shengjia.zhao@gmail.com, zhangyuhan2022@ia.ac.cn, shan.yu@nlpr.ia.ac.cn

## Abstract

Current brain-computer interfaces primarily decode single motor variables, limiting natural control requiring simultaneous multi-dimensional extraction. We introduce Multi-dimensional Neural Decoding (MND), a task that simultaneously extracts multiple motor variables (direction, position, velocity, acceleration) from single neural population recordings. MND faces two key challenges: cross-task interference when decoding correlated motor dimensions from shared cortical representations, and generalization issues across sessions, subjects, and paradigms. To address these challenges, we propose **OrthoSchema**, a multi-task framework inspired by cortical orthogonal subspace organization and cognitive schema reuse. OrthoSchema enforces representation orthogonality to eliminate cross-task interference and employs selective feature reuse transfer for few-shot cross-session, subject and paradigm adaptation. Experiments on macaque motor cortex datasets demonstrate that OrthoSchema significantly improves decoding accuracy in cross-session, subject and paradigm generalization tasks, with larger performance improvements when fine-tuning samples are limited. Ablation studies confirm the synergistic effects of all components are crucial, with OrthoSchema effectively modeling cross-task features and capturing session relationships for robust transfer. Our results provide new insights into scalable and robust neural decoding for real-world BCI applications.

## Introduction

Brain-computer interfaces (BCIs) restore motor function in neural injury patients by decoding neural signals into actionable commands. However, traditional BCIs focused on single-task decoding, extracting individual motor variables from motor cortex (Hochberg et al. 2012; Flesher et al. 2021). Motor cortex simultaneously encodes multiple correlated motor variables within shared neural populations (Cunningham and Yu 2014; Li 2014). Consequently, current BCIs struggle to support natural, fluent interactions that require information-rich, high-dimensional neural representations (Collinger et al. 2013; Wolpaw 2013), resulting in slow decoding responses and poor generalization in real-world

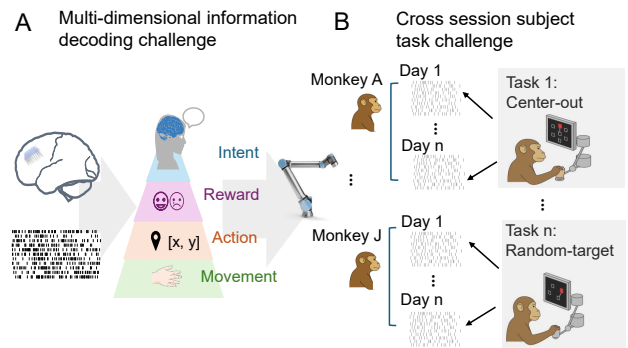


Figure 1: Key challenges in multi-dimensional neural decoding. (A) Growing demand for parallel decoding of diverse neural signals. (B) Limited generalization under distribution shifts across sessions, subjects, and paradigms.

scenarios (Shih, Krusienski, and Wolpaw 2012). We propose Multi-dimensional Neural Decoding (MND) as a novel task formulation that simultaneously extracts multiple correlated motor variables from shared neural population activity. Unlike traditional single-task approaches, MND formulation encounters the complex interactions that arise when decoding correlated motor dimensions from shared cortical representations. This task formulation not only captures the biological reality of cortical motor encoding but also meets the growing demand for high-bandwidth BCI control (Gallego, Makin, and McDougle 2022).

MND faces two key challenges, as illustrated in Figure 1. First, it remains unclear whether sufficient multi-dimensional information across different hierarchical levels can be extracted from limited implanted brain regions. When extending to multi-dimensional scenarios, cross-task interference exists among different dimensions, where extraction of one type of information adversely affects the decoding accuracy of other dimensions (Li et al. 2024). Second, since neural signals naturally drift over time (Karpowicz et al. 2025; Wimalasena, Miller, and Pandarinath 2020), existing methods (Lim et al. 2025) exhibit poor adaptability across sessions, subjects, and paradigm variations, often requiring

\*Corresponding author

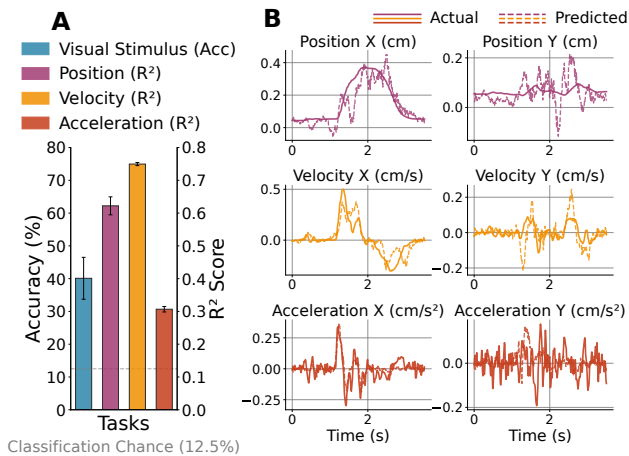


Figure 2: Feasibility of decoding distinct motor features from cortical activity. (A) Independent decoding of multiple motor features (direction, position, velocity, and acceleration) from motor cortex recordings within the same session. (B) Sample predictions for hand motor variables.

extensive retraining when facing new conditions.

Neural coding research provides insights for addressing these challenges. Studies in rodents (Zhou et al. 2021) and non-human primates (Bernardi et al. 2020) demonstrate that the brain can form schema-like representational patterns that remain consistent across sessions and subjects. Meanwhile, non-human primate research reveals that different functional representations within the same cortical area exhibit orthogonal properties (Tian et al. 2024; Flesch et al. 2022), and this orthogonal subspace organization facilitates interference-free representation of multi-dimensional information. These findings inspire two insights for algorithm design: orthogonal subspace organization may offer an approach to mitigate feature coupling problems in multi-dimensional decoding, while schema-like stable representation could provide a theoretical foundation for more efficient transfer learning strategies.

Drawing on these insights, we propose **OrthoSchema**—a structured multi-task neural decoding framework that aims to extract independent latent representations from neural activity and achieve improved generalization across subjects, sessions, and paradigms. The framework comprises three main components: (1) orthogonality constraints in the latent space to enhance separation among task-relevant features and reduce interference across decoding objectives; (2) an explicit session modeling mechanism that captures variations across sessions and subjects, with these representations excluded during transfer to remove drift-related information; (3) a parameter transfer strategy guided by feature relevance, enabling efficient few-shot adaptation to new sessions, subjects, or paradigms.

To validate the effectiveness of our approach, we employed macaque motor cortex neural datasets from two different paradigms: center-out reaching and random target tasks (Perich et al. 2025). We first experimentally demon-

strated the practical feasibility of independently decoding multiple motor features from motor cortex (M1), as shown in Figure 2, confirming the potential for multi-task neural decoding. Subsequently, we evaluated the few-shot capabilities of the OrthoSchema architecture across different backbone networks in cross-subject, cross-session, and cross-paradigm scenarios. Experimental results demonstrate that our method achieves significant performance improvements across all testing scenarios, with particularly outstanding performance in the extremely few-shot regime ( $K=2$ ). Ablation studies and visualization analyses suggest that orthogonal constraints promote the formation of task-specific feature patterns, while the schema-based transfer mechanism successfully decouples task-relevant representations from session-specific variations, providing mechanistic insights into the model’s superior generalization capability. The main contributions of this work include:

- We propose MND as a BCI task formulation that meets the growing demand in BCI systems for simultaneously decoding multiple correlated motor variables from the same neural population.
- We develop the **OrthoSchema** framework, combining orthogonal constraints and schema-based transfer mechanisms to effectively address cross-task interference and cross-session/subject/paradigm generalization challenges.
- We validate the effectiveness of our framework on macaque motor cortex datasets, demonstrating superior few-shot transfer performance across cross-session, subject, and paradigm scenarios.

## Related Work

### Neural Decoding

Neural decoding refers to the computational process of extracting behavioral intentions and cognitive states from recorded neural signals (Musallam et al. 2004). Traditional BCI decoding methods primarily rely on linear models, such as Kalman filters (Wu et al. 2006) and population vector algorithms (Georgopoulos, Schwartz, and Kettner 1986), for motor intention prediction. Deep learning techniques have significantly improved decoding capabilities (Glaser et al. 2020). LSTM networks have been applied to motion trajectory decoding (Liu et al. 2022), while convolutional neural networks have enhanced spatial feature extraction (Xie, Schwartz, and Prasad 2018). Recent approaches focus on modeling population neural dynamics. LFADS (Pandarinath et al. 2018) uses variational autoencoders for latent neural state inference. NOMAD (Karpowicz et al. 2025) improves motor decoding through latent variable alignment. However, these methods typically address single decoding objectives and face significant challenges when simultaneously extracting multiple behavioral variables from the same neural population.

### Multi-dimensional Information Decoding

With the demand for richer interaction capabilities in BCI applications, multi-dimensional information decoding has

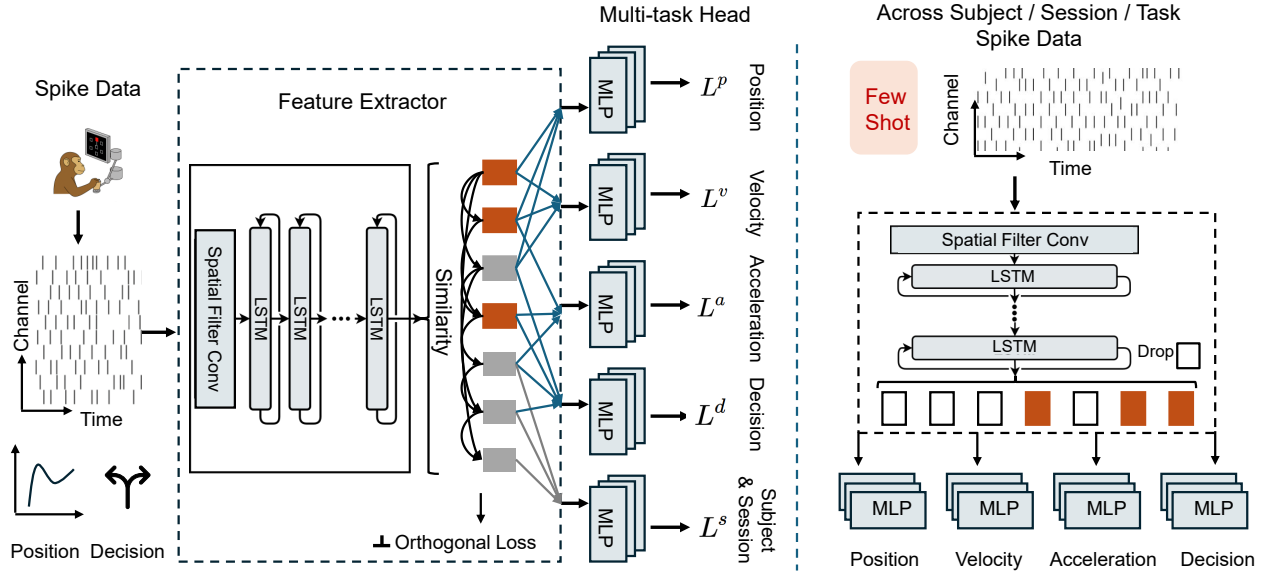


Figure 3: Overview of the OrthoSchema framework. Spike data are processed by a global convolution layer followed by LSTM blocks (replaceable) to obtain latent representations. Orthogonality constraints are applied in the latent space. The model includes multiple decoding heads for direction classification and the regression of hand position, velocity, and acceleration. A session/subject classification head is employed during training to model distributional shifts. During inference, the session classification components are removed, and parameters are selectively reused for few-shot fine-tuning.

emerged as a central challenge in current neural decoding methods (Wolpaw 2013). MND has been applied to EEG-based decoding, including Bayesian multi-task models (Alamgir, Grosse-Wentrup, and Altun 2010; Oikonomou, Nikolopoulos, and Kompatsiaris 2022) and deep neural networks with multi-head outputs (Autthasan et al. 2021), demonstrating cross-dimensional interference challenges.

Existing multi-task learning methods address task conflicts through shared/private subspace decomposition (Dong et al. 2022), adaptive loss weighting via gradient magnitudes (Chen et al. 2018), or uncertainty-based task weighting (Kendall, Gal, and Cipolla 2018). However, neuroscience research suggests that neural population can form mutually orthogonal encodings in output subspaces (Kaufman et al. 2014; Tang et al. 2020; Elsayed et al. 2016), encouraging features to be approximately orthogonal without interfering with the shared encoder’s general representation capability.

### Cross-session/subject/paradigm Adaptation

Neural decoding stability across sessions, subjects, and paradigms poses major challenges for practical BCI deployment, as neural activity changes significantly due to electrode displacement and tissue reactions (Perge et al. 2013; Chestek et al. 2011; Willett et al. 2021). Farshchian et al. (2018) proposed adversarial learning-based domain adaptation methods that reduce cross-session variability by learning domain-invariant representations. Degenhart et al. (2020) developed alignment strategies based on shared neural manifolds, achieving cross-session generalization. Recent foundation models such as Neural Data Transformer (Ye et al. 2023) and POYO (Azabou et al. 2023) achieve

cross-subjects generalization through large-scale pretraining. However, these methods require extensive neural data for pretraining while ignoring the differential stability of schema-like neural representations (Goudar et al. 2023) across sessions, subjects, and paradigms revealed in neuroscience research.

Unlike existing single-output decoders and data-intensive foundation models, OrthoSchema is designed for MND tasks, integrating orthogonality constraints in latent space with selective schema reuse to enhance multi-dimensional decoding stability across sessions, subjects, and paradigms.

## Method

MND is defined as the simultaneous extraction of multiple correlated motor variables  $\{y_1, y_2, \dots, y_n\}$  from shared neural population recordings  $X$ , where variables exhibit hierarchical dependencies (e.g., position, velocity, acceleration). We propose OrthoSchema to address MND’s core challenges: cross-task interference from correlated dimensions and poor generalization across sessions, subjects, and paradigms through orthogonal feature disentanglement and schema-based transfer learning.

Figure 3 illustrates OrthoSchema neural decoding framework comprising four modules: (a) orthogonally-constrained feature extraction; (b) multi-decoding heads; (c) transfer with selective feature reuse. The pipeline processes neural spike data through spatial filtering and hierarchical networks, applies orthogonal constraints for feature diversity, employs multi-task learning for BCI decoding and session classification, and achieves few-shot generalization

through selective parameter transfer.

### Orthogonally-Constrained Feature Extractor

The orthogonally-constrained feature extractor extracts decoupled features from neural spike data to reduce cross-dimensional interference. This module employs a model-agnostic design with a scalable shared encoder that can be adapted to different backbone architectures. Neural spike data  $X \in \mathbb{R}^{B \times 1 \times C \times T}$  (where  $B$  is batch size,  $C$  is channels,  $T$  is time steps) is first processed through a spatial filter to extract coordinated neural activity patterns:

$$h_0 = \text{ELU}(\text{BatchNorm2d}(\text{Conv2d}(X, \theta_0))) \quad (1)$$

where  $\text{Conv2d}(\cdot; \theta_0)$  denotes 2D convolution with learnable parameters  $\theta_0$ ,  $\text{BatchNorm2d}(\cdot)$  represents batch normalization for stable training, and  $\text{ELU}(\cdot)$  is the Exponential Linear Unit activation function. After dimension adjustment, the data is processed through bidirectional LSTM layers:

$$h_{k+1}, (c_{k+1}, m_{k+1}) = \text{LSTM}_k(h_k, (c_k, m_k)) \quad (2)$$

$$h_{k+1} = \text{ELU}(\text{BatchNorm1d}(h_{k+1}^T))^T \quad (3)$$

where  $k = 1, 2, \dots, 6$ ,  $h_k$  represents the hidden state,  $(c_k, m_k)$  denote the cell state and memory state of the LSTM at layer  $k$ .

**Orthogonal Feature Constraint** We impose orthogonal constraints on the final RNN layer features  $h_6$  to enhance feature diversity and promote independent representations across-dimensions. The orthogonality loss is computed as:

$$L_{\text{orthogonal}} = \|\text{mean}(x_{\text{norm}} \cdot x_{\text{norm}}^T) - I\|_F^2 \quad (4)$$

where  $x_{\text{norm}}$  represents the normalized features  $h_6$ ,  $I$  is the identity matrix. This constraint incurs minimal computational overhead, adding approximately 3% to training time, while the selective transfer mechanism introduces less than 1% additional cost during few-shot adaptation.

### Multi-Task Head

All task heads connect to the final RNN features  $h_L$  processed with orthogonal constraints. The architecture includes two types of task heads: (1) BCI decoding heads for motor-related variables and (2) a session classification head to model subject/session-specific features and promote cross-session generalization.

Each BCI decoding head is implemented as a multi-layer perceptron (MLP), targeting a specific aspect of motor behavior. For direction classification, we use temporally averaged features:

$$\bar{x} = \frac{1}{T} \sum_{t=1}^T x^{(t)} \in \mathbb{R}^{B \times 128} \quad (5)$$

$$f_{\text{direction}} = \text{MLP}_{\text{direction}}(\bar{x}) \in \mathbb{R}^{B \times 8} \quad (6)$$

where  $x^{(t)}$  represents the feature at the  $t$ -th time step, and  $\bar{x}$  denotes the temporally averaged features. For the three motor regression tasks:

$$f_{\text{position}} = \text{MLP}_{\text{position}}(x), \quad (7)$$

$$f_{\text{velocity}} = \text{MLP}_{\text{velocity}}(x) \quad (8)$$

$$f_{\text{acceleration}} = \text{MLP}_{\text{acceleration}}(x) \quad (9)$$

where  $x \in \mathbb{R}^{B \times T \times 128}$  represents the features for time-step-wise processing.

To identify and model session/subject drift information, we design a classification task head to recognize neural signal characteristics across different subjects and sessions:

$$f_{\text{session}} = \text{MLP}_{\text{session}}(x_{\text{avg}}) \in \mathbb{R}^{B \times N_{\text{sessions}}} \quad (10)$$

where  $f_{\text{session}}$  represents the output probabilities for session classification,  $x_{\text{avg}}$  denotes the temporally averaged features, and  $N_{\text{sessions}}$  is the total number of sessions in the training data. This task head helps capture session/subject-related features and enables identification of stable representations by filtering out session/subject-specific components during few-shot transfer.

**Loss Function** The total loss function integrates losses from all tasks along with the orthogonal constraint loss:

$$L_{\text{total}} = \alpha_1 L_{\text{dir}} + \alpha_2 (L_{\text{pos}} + L_{\text{vel}} + L_{\text{acc}}) + \alpha_3 L_{\text{orthogonal}} + \alpha_4 L_{\text{session}} \quad (11)$$

where classification tasks employ cross-entropy loss and regression tasks use mean squared error loss.  $\alpha_1, \alpha_2, \alpha_3, \alpha_4$  are weighting factors that balance different objectives.

### Transfer with Selective Feature Reuse

Inspired by the brain's ability to reuse stable cognitive schemas for cross-session, subject and paradigm generalization, we develop a selective parameter transfer strategy for efficient cross-session/subject/paradigm adaptation. Our approach focuses on the final layer of RNN features and computes task-relevant and session/subject-specific relevance separately:

We first compute the average weight importance across all task heads to identify shared representations for multiple motor decoding tasks:

$$I_{\text{task}}(\theta_i) = \frac{1}{|H|} \sum_{h \in H} |\text{weight}_{h,i}| \quad (12)$$

where  $H$  is the set of BCI task heads and  $\text{weight}_{h,i}$  is the weight connecting the  $i$ -th feature to task head  $h$ . Simultaneously, we compute the parameter importance for the session classification task to capture session/subject-specific drift features:

$$I_{\text{session}}(\theta_i) = |\text{weight}_{\text{session},i}| \quad (13)$$

where  $\text{weight}_{\text{session},i}$  is the weight connecting the  $i$ -th feature to the session classification head.

We extract task-relevant yet stable representations through importance difference:

$$I_{\text{transfer}}(\theta_i) = I_{\text{task}}(\theta_i) - \beta \cdot I_{\text{session}}(\theta_i) \quad (14)$$

$$\Theta_{\text{transfer}} = \{\theta_i | I_{\text{transfer}}(\theta_i) \geq \text{percentile}(I_{\text{transfer}}, 100 - \gamma)\} \quad (15)$$

where  $\beta$  controls the degree of down-weighting session-specific features, and  $\gamma$  is the percentage of top-ranked features to select. This strategy ensures that selected features are important for BCI tasks but contribute little to session classification. During  $k$ -shot fine-tuning, we only update the selected parameters in  $\Theta_{\text{transfer}}$ , ensuring that only task-relevant stable features receive gradient updates for effective few-shot generalization.

Dataset	Methods	Across Session							
		Direction (%)		Hand Position ( $R^2$ )		Hand Velocity ( $R^2$ )		Hand Acceleration ( $R^2$ )	
		K=2	K=5	K=2	K=5	K=2	K=5	K=2	K=5
CO	CNN	33.97 ± 3.93	73.55 ± 6.11	0.533 ± 0.080	0.724 ± 0.038	0.650 ± 0.022	0.783 ± 0.024	0.299 ± 0.147	0.562 ± 0.049
	CNN+OS	<b>37.99 ± 1.35</b>	<b>79.42 ± 2.24</b>	<b>0.644 ± 0.033</b>	<b>0.769 ± 0.017</b>	<b>0.720 ± 0.017</b>	<b>0.801 ± 0.019</b>	<b>0.440 ± 0.020</b>	0.550 ± 0.014
	RNN	72.60 ± 4.64	88.32 ± 0.53	0.886 ± 0.015	0.950 ± 0.002	0.838 ± 0.013	0.916 ± 0.003	0.576 ± 0.015	0.715 ± 0.017
	RNN+OS	<b>80.23 ± 0.51</b>	<b>89.68 ± 0.41</b>	<b>0.903 ± 0.005</b>	<b>0.952 ± 0.004</b>	<b>0.862 ± 0.014</b>	<b>0.922 ± 0.002</b>	<b>0.611 ± 0.004</b>	<b>0.735 ± 0.010</b>
	MLP	/	94.02 ± 2.15	/	0.860 ± 0.009	/	0.529 ± 0.055	/	/
	MLP+OS	/	<b>94.39 ± 0.86</b>	/	<b>0.863 ± 0.006</b>	/	<b>0.554 ± 0.006</b>	/	/
RT	CNN	—	—	0.566 ± 0.034	0.708 ± 0.022	0.698 ± 0.017	0.774 ± 0.022	—	—
	CNN+OS	—	—	<b>0.601 ± 0.022</b>	<b>0.727 ± 0.032</b>	<b>0.719 ± 0.007</b>	<b>0.798 ± 0.012</b>	—	—
	RNN	—	—	0.635 ± 0.017	0.712 ± 0.057	0.694 ± 0.008	0.782 ± 0.008	—	—
	RNN+OS	—	—	<b>0.681 ± 0.019</b>	<b>0.768 ± 0.050</b>	<b>0.717 ± 0.014</b>	<b>0.794 ± 0.020</b>	—	—

Table 1: Cross-session performance under 2-shot and 5-shot settings (mean±std). Direction is reported as accuracy (%), while position, velocity, and acceleration are reported as  $R^2$  values. Bold values represent the best results within each k-shot group. “/” indicates results at chance level. “—” indicates the task is not applicable. +OS denotes enhancement with OrthoSchema.

Methods	k-shot	Across Subject		
		Direction (%)	Position ( $R^2$ )	Velocity ( $R^2$ )
CNN	K=2	38.93 ± 6.53	0.504 ± 0.049	0.538 ± 0.026
CNN+OS	K=2	<b>44.76 ± 2.84</b>	<b>0.546 ± 0.039</b>	<b>0.543 ± 0.012</b>
CNN	K=5	65.52 ± 4.31	0.620 ± 0.044	0.622 ± 0.015
CNN+OS	K=5	<b>66.50 ± 5.41</b>	<b>0.663 ± 0.040</b>	<b>0.644 ± 0.018</b>
RNN	K=2	45.53 ± 4.45	0.685 ± 0.039	0.592 ± 0.017
RNN+OS	K=2	<b>60.01 ± 3.69</b>	<b>0.783 ± 0.022</b>	<b>0.653 ± 0.021</b>
RNN	K=5	75.65 ± 4.43	0.824 ± 0.017	0.697 ± 0.022
RNN+OS	K=5	<b>77.76 ± 1.97</b>	<b>0.872 ± 0.009</b>	<b>0.743 ± 0.011</b>
MLP	K=5	90.25 ± 1.46	0.854 ± 0.008	0.589 ± 0.020
MLP+OS	K=5	<b>91.02 ± 0.71</b>	<b>0.856 ± 0.008</b>	<b>0.607 ± 0.004</b>

Table 2: Cross-subject performance on CO task under 2-shot and 5-shot settings (mean±std). Direction is reported as accuracy (%), while position, velocity are reported as  $R^2$  values. Bolded values are the best results within each k-shot group. +OS indicates enhanced with OrthoSchema.

Methods	k-shot	Across Paradigm (CO→RT)	
		Position ( $R^2$ )	Velocity ( $R^2$ )
CNN	K=2	0.233 ± 0.061	0.400 ± 0.046
CNN+OS	K=2	<b>0.240 ± 0.060</b>	<b>0.427 ± 0.026</b>
CNN	K=5	0.359 ± 0.038	0.531 ± 0.031
CNN+OS	K=5	<b>0.425 ± 0.024</b>	<b>0.546 ± 0.011</b>
RNN	K=2	0.579 ± 0.040	0.688 ± 0.023
RNN+OS	K=2	<b>0.622 ± 0.013</b>	<b>0.697 ± 0.000</b>
RNN	K=5	0.707 ± 0.013	0.762 ± 0.006
RNN+OS	K=5	<b>0.728 ± 0.009</b>	<b>0.782 ± 0.003</b>
MLP	K=5	0.618 ± 0.018	0.250 ± 0.015
MLP+OS	K=5	<b>0.624 ± 0.002</b>	<b>0.257 ± 0.010</b>

Table 3: Cross-paradigm performance (CO→RT) under 2-shot and 5-shot settings (mean±std). Position and Velocity are reported as  $R^2$  values. Bolded values are the best results. +OS indicates enhanced with OrthoSchema.

## Experiments

### CO and RT Dataset

We use motor cortex neural datasets from three rhesus macaques (Perich et al. 2025). The dataset contains two motor paradigms: center-out (CO) tasks require macaques to perform reaching movements from the center position to the targets based on one of directional cues displayed on the screen; random target (RT) tasks require macaques to reach targets that appear at random locations on the screen. Subject C provided 14 CO task sessions (recorded September-October 2016), Subject M provided 20 CO task sessions and 6 RT task sessions (recorded January 2014-June 2015), and Subject T provided 6 CO task sessions and 6 RT task sessions (recorded August-September 2013). These data pro-

vide a testing foundation for validating the model’s cross-session, cross-subject, and cross-paradigm generalization capabilities in MND. CO tasks contain neural spike activity, hand position, velocity, acceleration, and directional choice labels; RT tasks contain neural spike activity, hand position, and velocity. For neural data processing, we extract  $T = 3.5$  seconds of neural activity starting from target onset with sampling rate  $f_s = 100$  Hz. Additional implementation details are provided in the appendix.

### Experimental Results

We conduct generalization experiments using CO and RT datasets on three representative neural decoding backbone architectures (CNN, RNN, and MLP) (Liu et al. 2022), evaluating performance across three scenarios: cross-session,

Methods	Direction (%)	Position ( $R^2$ )	Velocity ( $R^2$ )	Acceleration ( $R^2$ )
RRR	37.60 ± 4.20	0.456 ± 0.066	0.102 ± 0.037	-0.260 ± 0.062
LFADS	64.88 ± 4.35	0.841 ± 0.040	0.818 ± 0.019	0.555 ± 0.021
NoMAD	67.74 ± 2.21	0.814 ± 0.058	0.78 ± 0.006	0.524 ± 0.018
Cycle-GAN	36.19 ± 3.61	0.343 ± 0.064	0.180 ± 0.026	0.012 ± 0.049
STNDT	32.12 ± 1.90	0.738 ± 0.108	0.513 ± 0.094	0.081 ± 0.012
UW	52.03 ± 1.70	0.790 ± 0.007	0.763 ± 0.020	0.455 ± 0.018
GradNorm	44.35 ± 1.12	0.807 ± 0.022	0.802 ± 0.003	0.499 ± 0.006
<b>RNN+OS</b>	<b>80.23 ± 0.51</b>	<b>0.902 ± 0.005</b>	<b>0.862 ± 0.001</b>	<b>0.611 ± 0.004</b>

Table 4: Performance comparison with multi-task and neural decoding baselines on cross-session generalization under 2-shot settings (mean±std). Direction is reported as accuracy (%), regression tasks (position, velocity, acceleration) reported as  $R^2$  values. Bold indicates best performance per k-shot group. Results for k=5 are provided in Appendix.

cross-subject, and cross-paradigm. We use direction classification accuracy to evaluate motor intention decoding performance and  $R^2$  values to assess motor regression for position, velocity, and acceleration under few-shot settings with  $k \in \{2, 5\}$ . Bolded values are the best results within each k-shot group. “/” represents results at chance level due to insufficient feature extraction capability. “-” indicates the task is not applicable. +OS indicates enhanced with OrthoSchema.

Table 1 presents cross-session generalization performance on CO and RT paradigms. Results demonstrate that our method (+OS) achieves significant improvements across all metrics. Under the more restrictive 2-shot setting, performance improvements are more pronounced, validating the method’s effectiveness under extremely limited sample conditions. Table 2 shows cross-subject generalization results. Due to differences in neuronal recording, and motor control strategies across different individuals, cross-subject transfer is more challenging than cross-session. All backbone networks showed sustained improvements, with MLP results omitted under k=2 conditions due to chance-level performance. Similarly, improvement effects are more pronounced under k=2 conditions across all architectures. Table 3 presents cross-paradigm transfer results from CO to RT paradigms. Results for RT to CO are provided in Appendix. Cross-paradigm transfer is the most challenging scenario due to different motor patterns, dynamics, and neural representations across paradigms. Nevertheless, our method still achieves clear improvements on most metrics, with more significant improvements under k=2 conditions.

We also compare OrthoSchema with classic multi-task and neural decoding baseline methods to validate its effectiveness in MND tasks. Table 4 shows cross-session performance comparisons on CO task with neuroscience decoding method Reduced Rank Regression (RRR) (Borgognon et al. 2025), LFADS (Pandarinath et al. 2018), NoMAD (Karpowicz et al. 2025), Cycle-GAN (Ma et al. 2023), STNDT (Le and Shlizerman 2022), classic multi-task learning uncertainty weighting (UW) (Kendall, Gal, and Cipolla 2018), and gradient normalization (GradNorm) (Chen et al. 2018)

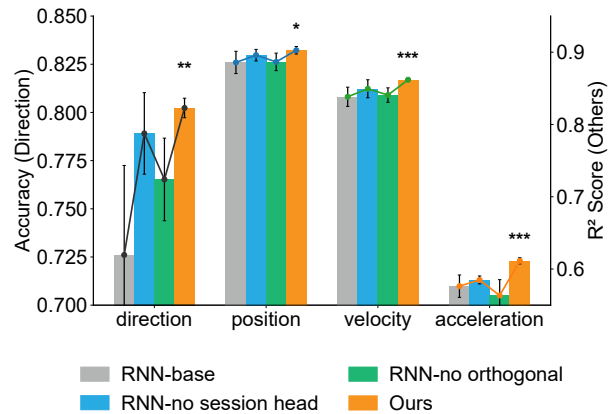


Figure 4: Performance comparison of different module combinations on cross-session generalization using RNN backbone under 2-shot setting. \*, \*\*, \*\*\* indicate statistical significance at  $p < 0.05$ ,  $p < 0.01$ ,  $p < 0.001$ , respectively.

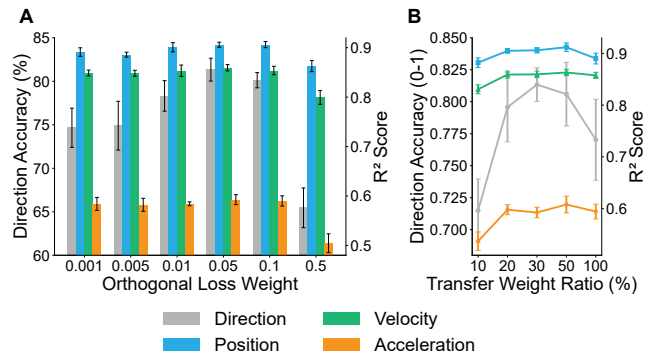


Figure 5: Hyperparameter analysis on cross-session generalization under 2-shot setting. (A) Impact of different orthogonal loss weights. (B) Impact of different feature retention ratios for transfer.

methods. Our results outperform all methods, especially under k=2 conditions. OrthoSchema also maintains similar results for other cross-subject and cross-paradigm scenarios, with details provided in the appendix.

Overall, our method achieves significant improvements under both cross-session and cross-subject settings, and also demonstrates good enhancement effects in challenging cross-paradigm condition. The k-shot fine-tuning mechanism provides more significant performance improvements under extreme sample constraints (k=2), further demonstrating the effectiveness of our method.

## Ablation Studies

**Impact of Module Combinations** We first evaluate the impact of different module combinations on performance. The variant models are defined as follows: (i) Base multi-task RNN, containing only direction and motor regression tasks; (ii) Multi-task RNN with orthogonal loss; (iii) Multi-task RNN with session classification head; (iv) Ours, com-

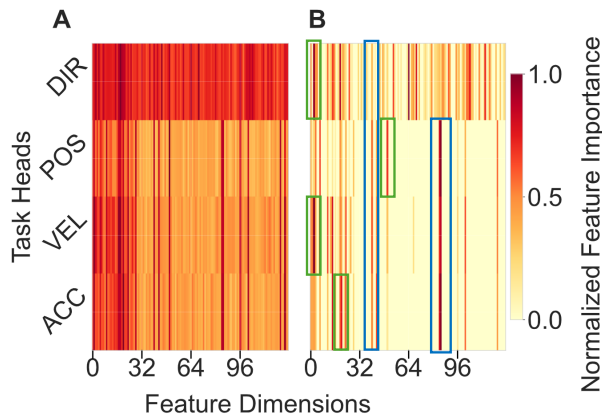


Figure 6: Comparison heatmap of feature importance extracted by orthogonality loss versus transferred feature importance. (A) Feature importance distribution of each head extracted by the trained network. (B) Adjusted feature importance transferred to new data after removing the influence of session/subject heads. Blue boxes indicate cross-task shared features, while green boxes represent task-specific features.

binning all components. As shown in Figure 4, all components positively impact performance. Compared to the base model, adding orthogonal loss improves direction accuracy by approximately 6%, adding session classification head brings about 2% improvement, while the complete method achieves a significant 8% enhancement. Statistical significance is also observed in regression tasks compared to other methods. This validates the effectiveness of each component and the importance of their synergistic effects.

**Impact of Hyperparameter Settings** We analyze the impact of key hyperparameters to determine optimal configurations and understand model behavior. Figure 5(A) shows the impact of orthogonal loss weights. Small weights (0.001-0.005) provide insufficient orthogonal constraints to decouple task features, while large weights (0.1-0.5) impose overly strong constraints that harm shared feature learning. A weight of 0.05 achieves optimal balance, ensuring feature decoupling while maintaining task performance. Figure 5(B) analyzes feature retention ratios. Low ratios (10-20%) cause loss of critical task information and performance degradation, while high ratios (50-100%) retain excessive session-specific noise that affects generalization. A 30% retention ratio achieves optimal performance by balancing information preservation and noise filtering.

**Extended Validation Experiments** To further validate the robustness of our method, we conducted two extended experiments. First, the impact of training days (Appendix Figure 3) shows that more training days lead to better performance, especially pronounced in direction tasks, demonstrating that the model can effectively utilize more training data to learn robust neural-behavioral mapping relationships. Second, single-task decoding validation (Appendix Figure 4) shows that our model framework is also effective




Figure 7: T-SNE projection of the learned session embeddings. Different colors represent different sessions.

in single-task decoding, confirming the method’s broad applicability. Detailed results are provided in the appendix.

### Visualization Analysis

To provide insights into our method’s internal mechanisms, we conduct feature importance visualization analysis. Figure 6 demonstrates the model’s effectiveness in feature decomposition and transfer. Figure 6 (A) shows that different task heads learn distinct feature importance patterns, proving successful cross-task feature disentanglement. Figure 6 (B) displays the adjusted feature importance after removing day/subject head influences. The transferred patterns are more concentrated and stable, with shared feature regions (blue boxes) showing stronger consistency while preserving task-specific features (green boxes). The transferred feature distribution exhibits less noise and clearer structure, directly explaining the method’s significant improvements in cross-scenario generalization. We further visualize the learned representations to observe how the model learns inter-session relationships (Figure 7). Analysis shows that adjacent sessions map to similar regions in low-dimensional space, indicating that the model not only learns to identify task-specific and shared patterns, but also captures macro-structural relationships across multiple sessions.

### Conclusion

This work defines the MND task and proposes the **OrthoSchema** as the solution, eliminating cross-task interference through orthogonal constraints and achieving cross-session, subjects and paradigms few-shot adaptation through schema-like selective feature transfer. Experiments on macaque motor cortex data validate the effectiveness of OrthoSchema, particularly under limited data conditions (K=2), with visualization revealing effective orthogonal disentanglement and schema-based transfer. Unlike data-intensive foundation models, OrthoSchema provides a data-efficient solution for scenarios with limited training samples. This work provides insights into multi-dimensional neural information extraction and generalization, informing the development of high-bandwidth BCIs.

## References

- Alamgir, M.; Grosse-Wentrup, M.; and Altun, Y. 2010. Multitask Learning for Brain-Computer Interfaces. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 17–24.
- Autthasan, P.; Chaisaen, R.; Sudhawiyangkul, T.; Rangpong, P.; Kiatthaveephong, S.; Dilokthanakul, N.; Bhakdisongkhram, G.; Phan, H.; Guan, C.; and Wilaiprasitporn, T. 2021. MIN2Net: End-to-End Multi-Task Learning for Subject-Independent Motor Imagery EEG Classification. *IEEE Transactions on Biomedical Engineering*, 69(6): 2105–2118.
- Azabou, M.; Arora, V.; Ganesh, V.; Mao, X.; Nachimuthu, S.; Mendelson, M.; Richards, B.; Perich, M.; Lajoie, G.; and Dyer, E. 2023. A Unified, Scalable Framework for Neural Population Decoding. In *Advances in Neural Information Processing Systems 36*, 44937–44956.
- Bernardi, S.; Benna, M. K.; Rigotti, M.; Munuera, J.; Fusi, S.; and Salzman, C. D. 2020. The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell*, 183(4): 954–967.
- Borgognon, S.; Macellari, N.; Hickey, A. M.; Perich, M. G.; Javaheri, H.; Ornelas-Kobayashi, R.; Delacombaz, M.; Hitz, C.; Fallegger, F.; Lacour, S. P.; et al. 2025. Regional Specialization of Movement Encoding across the Primate Sensorimotor Cortex. *Nature Communications*, 16: 5729.
- Chen, Z.; Badrinarayanan, V.; Lee, C.-Y.; and Rabinovich, A. 2018. GradNorm: Gradient Normalization for Adaptive Loss Balancing in Deep Multitask Networks. In *Proceedings of the 35th International Conference on Machine Learning*, 794–803.
- Chestek, C. A.; Batista, A. P.; Santhanam, G.; Yu, B. M.; Afshar, A.; Cunningham, J. P.; Gilja, V.; Ryu, S. I.; Churchland, M. M.; and Shenoy, K. V. 2011. Single-Neuron Stability during Repeated Reaching in Macaque Premotor Cortex. *Journal of Neuroscience*, 31(46): 17235–17243.
- Collinger, J. L.; Wodlinger, B.; Downey, J. E.; Wang, W.; Tyler-Kabara, E. C.; Weber, D. J.; McMorland, A. J. C.; Velliste, M.; Boninger, M. L.; and Schwartz, A. B. 2013. High-Performance Neuroprosthetic Control by an Individual with Tetraplegia. *The Lancet*, 381(9866): 557–564.
- Cunningham, J. P.; and Yu, B. M. 2014. Dimensionality Reduction for Large-Scale Neural Recordings. *Nature Neuroscience*, 17(11): 1500–1509.
- 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. *Nature Biomedical Engineering*, 4(7): 672–685.
- Dong, X.; Wu, R.; Xiong, C.; Li, H.; Cheng, L.; He, Y.; Qian, S.; Cao, J.; and Mo, L. 2022. GDOD: Effective Gradient Descent Using Orthogonal Decomposition for Multi-Task Learning. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management*, 386–395.
- Elsayed, G. F.; Lara, A. H.; Kaufman, M. T.; Churchland, M. M.; and Cunningham, J. P. 2016. Reorganization between Preparatory and Movement Population Responses in Motor Cortex. *Nature Communications*, 7: 13239.
- Farshchian, A.; Gallego, J. A.; Cohen, J. P.; Bengio, Y.; Miller, L. E.; and Solla, S. A. 2018. Adversarial Domain Adaptation for Stable Brain-Machine Interfaces. *arXiv preprint arXiv:1810.00045*.
- Flesch, T.; Juechems, K.; Dumbalska, T.; Saxe, A.; and Summerfield, C. 2022. Orthogonal Representations for Robust Context-Dependent Task Performance in Brains and Neural Networks. *Neuron*, 110(7): 1258–1270.
- Flesher, S. N.; Downey, J. E.; Weiss, J. M.; Hughes, C. L.; Herrera, A. J.; Tyler-Kabara, E. C.; Boninger, M. L.; Collinger, J. L.; and Gaunt, R. A. 2021. A Brain-Computer Interface That Evokes Tactile Sensations Improves Robotic Arm Control. *Science*, 372(6544): 831–836.
- Gallego, J. A.; Makin, T. R.; and McDougale, S. D. 2022. Going beyond Primary Motor Cortex to Improve Brain-Computer Interfaces. *Trends in Neurosciences*, 45(3): 176–183.
- Georgopoulos, A. P.; Schwartz, A. B.; and Kettner, R. E. 1986. Neuronal Population Coding of Movement Direction. *Science*, 233(4771): 1416–1419.
- Glaser, J. I.; Benjamin, A. S.; Chowdhury, R. H.; Perich, M. G.; Miller, L. E.; and Kording, K. P. 2020. Machine Learning for Neural Decoding. *eNeuro*, 7(4).
- Goudar, V.; Peysakhovich, B.; Freedman, D. J.; Buffalo, E. A.; and Wang, X.-J. 2023. Schema Formation in a Neural Population Subspace Underlies Learning-to-Learn in Flexible Sensorimotor Problem-Solving. *Nature Neuroscience*, 26(5): 879–890.
- Hochberg, L. R.; Bacher, D.; Jarosiewicz, B.; Masse, N. Y.; Simeral, J. D.; Vogel, J.; Haddadin, S.; Liu, J.; Cash, S. S.; Van Der Smagt, P.; et al. 2012. Reach and Grasp by People with Tetraplegia Using a Neurally Controlled Robotic Arm. *Nature*, 485(7398): 372–375.
- Karpowicz, B. M.; Ali, Y. H.; Wimalasena, L. N.; Sedler, A. R.; Keshtkaran, M. R.; Bodkin, K.; Ma, X.; Rubin, D. B.; Williams, Z. M.; Cash, S. S.; et al. 2025. Stabilizing Brain-Computer Interfaces through Alignment of Latent Dynamics. *Nature Communications*, 16: 4662.
- Kaufman, M. T.; Churchland, M. M.; Ryu, S. I.; and Shenoy, K. V. 2014. Cortical Activity in the Null Space: Permitting Preparation without Movement. *Nature Neuroscience*, 17(3): 440–448.
- Kendall, A.; Gal, Y.; and Cipolla, R. 2018. Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 7482–7491.
- Le, T.; and Shlizerman, E. 2022. STNDT: Modeling Neural Population Activity with Spatiotemporal Transformers. In *Advances in Neural Information Processing Systems 35*, 17926–17939.

- Li, Y.; Zhu, X.; Qi, Y.; and Wang, Y. 2024. Revealing Unexpected Complex Encoding but Simple Decoding Mechanisms in Motor Cortex via Separating Behaviorally Relevant Neural Signals. *eLife*, 12: RP87881.
- Li, Z. 2014. Decoding Methods for Neural Prostheses: Where Have We Reached? *Frontiers in Systems Neuroscience*, 8: 129.
- Lim, M. J. R.; Lo, J. Y. T.; Tan, Y. Y.; Lin, H.-Y.; Wang, Y. H.; Tan, D.; Wang, E.; Naing, M. Y. Y.; Ng, J. J. W.; Jefree, R. A. B.; et al. 2025. The State-of-the-Art of Invasive Brain-Computer Interfaces in Humans: A Systematic Review and Individual Patient Meta-Analysis. *Journal of Neural Engineering*, 22(2): 026013.
- Liu, F.; Meamardoost, S.; Gunawan, R.; Komiyama, T.; Mewes, C.; Zhang, Y.; Hwang, E.; and Wang, L. 2022. Deep Learning for Neural Decoding in Motor Cortex. *Journal of Neural Engineering*, 19(5): 056021.
- Ma, X.; Rizzoglio, F.; Bodkin, K. L.; Perreault, E.; Miller, L. E.; and Kennedy, A. 2023. Using Adversarial Networks to Extend Brain Computer Interface Decoding Accuracy over Time. *eLife*, 12: e84296.
- Musallam, S.; Corneil, B. D.; Greger, B.; Scherberger, H.; and Andersen, R. A. 2004. Cognitive Control Signals for Neural Prosthetics. *Science*, 305(5681): 258–262.
- Oikonomou, V. P.; Nikolopoulos, S.; and Kompatsiaris, I. 2022. A Multitask Bayesian Framework for the Analysis of Motor Imagery EEG Data. In *Proceedings of the 30th European Signal Processing Conference (EUSIPCO)*, 1308–1312.
- Pandarinath, C.; O’Shea, D. J.; Collins, J.; Jozefowicz, R.; Stavisky, S. D.; Kao, J. C.; Trautmann, E. M.; Kaufman, M. T.; Ryu, S. I.; Hochberg, L. R.; et al. 2018. Inferring Single-Trial Neural Population Dynamics Using Sequential Auto-Encoders. *Nature Methods*, 15(10): 805–815.
- Perge, J. A.; Homer, M. L.; Malik, W. Q.; Cash, S.; Eskandar, E.; Friehs, G.; Donoghue, J. P.; and Hochberg, L. R. 2013. Intra-Day Signal Instabilities Affect Decoding Performance in an Intracortical Neural Interface System. *Journal of Neural Engineering*, 10(3): 036004.
- Perich, M. G.; Miller, L. E.; Azabou, M.; and Dyer, E. L. 2025. Long-Term Recordings of Motor and Premotor Cortical Spiking Activity during Reaching in Monkeys. DANDI Archive. <https://doi.org/10.48324/dandi.000688/0.250122.1735>.
- Shih, J. J.; Krusienski, D. J.; and Wolpaw, J. R. 2012. Brain-Computer Interfaces in Medicine. *Mayo Clinic Proceedings*, 87(3): 268–279.
- Tang, C.; Herikstad, R.; Parthasarathy, A.; Libedinsky, C.; and Yen, S.-C. 2020. Minimally Dependent Activity Subspaces for Working Memory and Motor Preparation in the Lateral Prefrontal Cortex. *eLife*, 9: e58154.
- Tian, K.; Zhao, Z.; Chen, Y.; Ge, N.; Cao, S.; Han, X.; Gu, J.; and Yu, S. 2024. Domain-Specific Schema Reuse Supports Flexible Learning to Learn in Primate Brain. *bioRxiv*.
- Willett, F. R.; Avansino, D. T.; Hochberg, L. R.; Henderson, J. M.; and Shenoy, K. V. 2021. High-Performance Brain-to-Text Communication via Handwriting. *Nature*, 593(7858): 249–254.
- Wimalasena, L. N.; Miller, L. E.; and Pandarinath, C. 2020. From Unstable Input to Robust Output. *Nature Biomedical Engineering*, 4(7): 665–667.
- Wolpaw, J. R. 2013. Brain-Computer Interfaces. *Handbook of Clinical Neurology*, 110: 67–74.
- Wu, W.; Gao, Y.; Bienenstock, E.; Donoghue, J. P.; and Black, M. J. 2006. Bayesian Population Decoding of Motor Cortical Activity Using a Kalman Filter. *Neural Computation*, 18(1): 80–118.
- Xie, Z.; Schwartz, O.; and Prasad, A. 2018. Decoding of Finger Trajectory from ECoG Using Deep Learning. *Journal of Neural Engineering*, 15(3): 036009.
- Ye, J.; Collinger, J.; Wehbe, L.; and Gaunt, R. 2023. Neural Data Transformer 2: Multi-Context Pretraining for Neural Spiking Activity. In *Advances in Neural Information Processing Systems 36*, 80352–80374.
- Zhou, J.; Jia, C.; Montesinos-Cartagena, M.; Gardner, M. P. H.; Zong, W.; and Schoenbaum, G. 2021. Evolving Schema Representations in Orbitofrontal Ensembles during Learning. *Nature*, 590(7847): 606–611.