

Self-Supervised Consistency Enhanced Disentangled Learning for Neural Decoding Generalization in Brain-Machine Interfaces

Jiyu Wei^{1,2}, Di Hong^{1,2}, Zhanjie Zhang¹, Dazhong Rong^{*,1,2}, Qinming He^{1,2}, Yueming Wang^{*,1,2}

¹College of Computer Science and Technology, Zhejiang University, Hangzhou, China

²Nanhu Brain-Computer Interface Institute, Hangzhou, China

{weijiyu, hongd, cszzj, rdz98, hqm, yumingwang}@zju.edu.cn

Abstract—Brain-Machine Interfaces (BMIs) provide a direct communication pathway between the brain and external devices, enabling humans to control assistive and robotic technologies, with potential applications in rehabilitation, human motor augmentation, and human-centered robotics. However, due to neural drift, the performance of BMIs decreases over time, posing challenges for long-term viability, particularly for invasive BMIs (iBMIs). Existing solutions suffer from two main drawbacks: (i) difficulty in learning robust neural representations, and (ii) neglecting that neural drift varies across motor parameters (e.g., velocity, direction, and speed). To overcome these limitations, we propose Self-Supervised Consistency enhanced Disentangled Learning (SSCDL), a neural decoding generalization framework built on two key innovations. We first design a backbone model named Consistency enhanced Neural Decoder (CND), using a novel teacher-student consistency constraint with simulated neural signal perturbations to learn robust representations invariant to neural drift. Then, we employ three dedicated CNDs under Complementary-Disentangled Generalization (CDG) mechanism, which disentangles motor signals into velocity, direction and speed with the inspiration of neural preference theory. This disentangled learning enables SSCDL to capture invariant neural representations from diverse neural preference perspectives, significantly enhancing cross-day generalization. Extensive experimental results show that SSCDL delivers state-of-the-art decoding performance, exhibiting high robustness and cross-day stability. These capabilities underscore its strong potential for long-term interaction in human-centric robotic and fine-grained assistive applications.

I. INTRODUCTION

Brain-Machine Interfaces (BMIs) translate neural activity into actionable commands for external devices, thereby enabling direct interaction with assistive and robotic technologies [1], [2], [3]. Among them, invasive motor BMIs (iBMIs) leverage high-quality neural recordings from motor cortical areas, providing not only a powerful assistive tool but also a unique experimental platform for investigating the neural mechanisms underlying motor control [4], [5], [6], [7]. Owing to these advantages, iBMIs have demonstrated significant potential in neurorehabilitation, human motor augmentation, and human-centered robotics.

As shown in Fig. 1, a central component of iBMIs is neural decoder, which maps recorded brain activity into control commands through sophisticated decoding algorithms [8], [9], [10], [11], [12], [13]. Although a variety of decoding methods have achieved reliable short-term performance,




Fig. 1. Schematic of a Brain-Machine Interfaces (BMIs): neural activity is recorded from the brain, processed by a decoder, and converted into control signals for external devices such as robots, computers, and assistive systems.

their accuracy typically degrades over days due to neural drift [14], [15]. This variability stems from factors such as electrode drift, neuronal turnover, and synaptic plasticity [16], [17], substantially hindering the widespread adoption and practical application of iBMIs.

Current solutions to address neural drift can be broadly categorized into recalibration-based and generalization-based methods. Recalibration-based methods [18], [19], [20], [15], [21], [22] maintain long-term decoding performance by retraining decoders with newly collected labeled or unlabeled data. However, this reliance on fresh data collection is often impractical in real-world applications, where acquiring such data is both costly and time-consuming. In contrast, generalization-based methods improve cross-day robustness without requiring decoder retraining, which is also the focus of this work. Representative generalization-based methods include: KRSRL [10] and EvoEnsemble [12] model short-term neural signal changes by siamese learning and ensemble strategies; SABLE [23] and CAPTIVATE [24] employ adversarial feature alignment and neuron-based latent synchronization; MRNN [25] models long-term nonlinear neural-kinematic mappings; and LFDA [26] leverages data augmentation with contrastive learning to simulate temporal variations. Despite these advances, learning long-term stable neural representations from only single-day data remains highly challenging.

One key limitation of existing generalization-based methods is that they treat velocity as a unified target variable, overlooking the fact that neural drift may manifest differently across motor parameters. Specifically, since individual neurons exhibit distinct tuning preferences for direction and speed [8], [27], the drift affecting neural representations may

*Corresponding authors: Dazhong Rong and Yueming Wang.

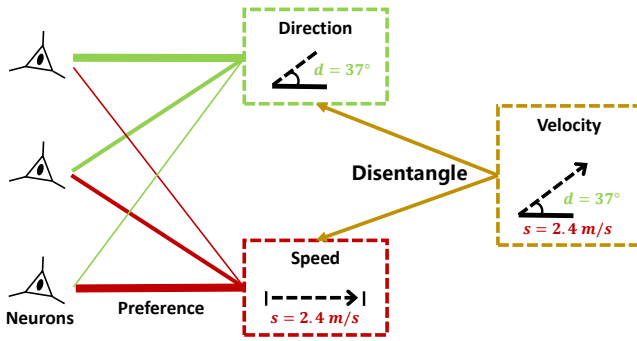


Fig. 2. Neural tuning preference: individual neurons exhibit distinct selectivity to different motor parameters disentangled from velocity (e.g., speed and direction), implying that neural drift may vary across these parameters. Line thickness indicates the strength of preference.

not be uniform across these dimensions (as shown in Fig. 2). Recent studies [22] further suggest that aligning velocity features as a whole can lead to suboptimal alignment in its constituent components, such as speed and direction. Consequently, enforcing uniform generalization across the entire velocity space may compromise the model generalization ability on specific motor parameters.

In this paper, we argue that disentangling neural decoding into velocity and its constituent components (direction and speed) provides a more effective strategy for capturing invariant neural representations and mitigating neural drift, thereby improving cross-day generalization. Rather than treating velocity as a unified target, this parameter-aware decomposition explicitly accounts for the heterogeneous tuning preferences of neurons. Our hypothesis is conceptually aligned with recent advances in disentangled representation learning [28], [29], which have demonstrated strong potential for extracting robust and transferable features under distribution shifts.

To address the aforementioned limitations, we propose Self-Supervised Consistency enhanced Disentangled Learning (SSCDL), a robust neural decoding framework for long-term brain-machine interface in human-centered robotic systems. Building on recent progress in self-supervised learning (SSL) [30], [31], [32], we first develop a Consistency enhanced Neural Decoder (CND) based on a teacher-student paradigm as the backbone model. CND enforces a representation-level consistency constraint, encouraging the teacher and student networks to produce aligned outputs under simulated neural signal perturbations. This design improves the robustness of learned representations against neural variability across days, even when training is limited to single-day recordings. Motivated by the well-established principle that individual neurons exhibit distinct preferences for motor parameters [27], [8], [22], [33], we further introduce Complementary-Disentangled Generalization (CDG) and employ our CNDs under it. Specifically, CDG decomposes conventional neural decoding into three interconnected components—overall velocity, direction, and speed—and assigns a dedicated SSD to each component. The resulting predictions are integrated through a complementary ensemble mechanism, enabling the model to capture invariant repre-

sentations across multiple motor dimensions. This structured design facilitates more precise and stable control, which is particularly critical for long-term, fine-grained interaction with assistive and robotic devices. Together, CND and CDG constitute a scalable and flexible generalization framework that mitigates cross-day performance degradation while remaining readily extensible to other neural decoding scenarios. Our main contributions are summarized as follows:

- We propose Self-Supervised Consistency enhanced Disentangled Learning (SSCDL), a novel neural decoding generalization framework that mitigates cross-day decoding performance degradation, enabling long-term stable BMI control in robotic systems.
- We develop a backbone model, Consistency enhanced Neural Decoder (CND), enforcing representation-level consistency under neural signal perturbations to learn invariant features from single-day data. We apply CND to Complementary-Disentangled Generalization (CDG), which decomposes neural decoding into velocity, direction, and speed components, addressing motor-parameter-dependent neural drift by performing motor-parameter-aware consistency learning.
- Extensive experiments on multiple iBMI datasets show that SSCDL achieves superior cross-day decoding performance and improved fine-grained prediction accuracy in direction and speed, highlighting its effectiveness for long-term stable BMI control.

II. RELATED WORK

Due to neural data drift, BMIs face a significant challenge in maintaining decoding accuracy across days. Existing solutions can be broadly categorized into recalibration-based methods and generalization-based methods.

A. Recalibration-based Methods

Traditional approaches typically rely on periodic decoder retraining using newly collected labeled data [18]. However, the collection of labeled data is often costly, severely limiting the practicality of BMIs. To address this, more advanced methods adopt Unsupervised Domain Adaptation (UDA), treating day₀ data as the source domain and subsequent day_k data as the target domain. Among them, ADAN [20] and UAN [15] adapt the target domain to the source domain by adversarial learning; DA-DCF [34] and WDGRL [21] align the source and target domains from the perspective of feature alignment; SeSA [22] introduce semi-supervised domain adaptation by integrating semi-supervised learning with UDA to conditionally align speed features between the source and target domains. Despite their effectiveness, these methods still require access to unlabeled or partially labeled data from each new session. This dependency poses practical challenges in real-world BMI deployments.

B. Generalization-based Methods

Generalization-based methods aim to maintain cross-day performance without further model updates. To achieve the goal, KRSRL [10] employs a siamese learning framework;

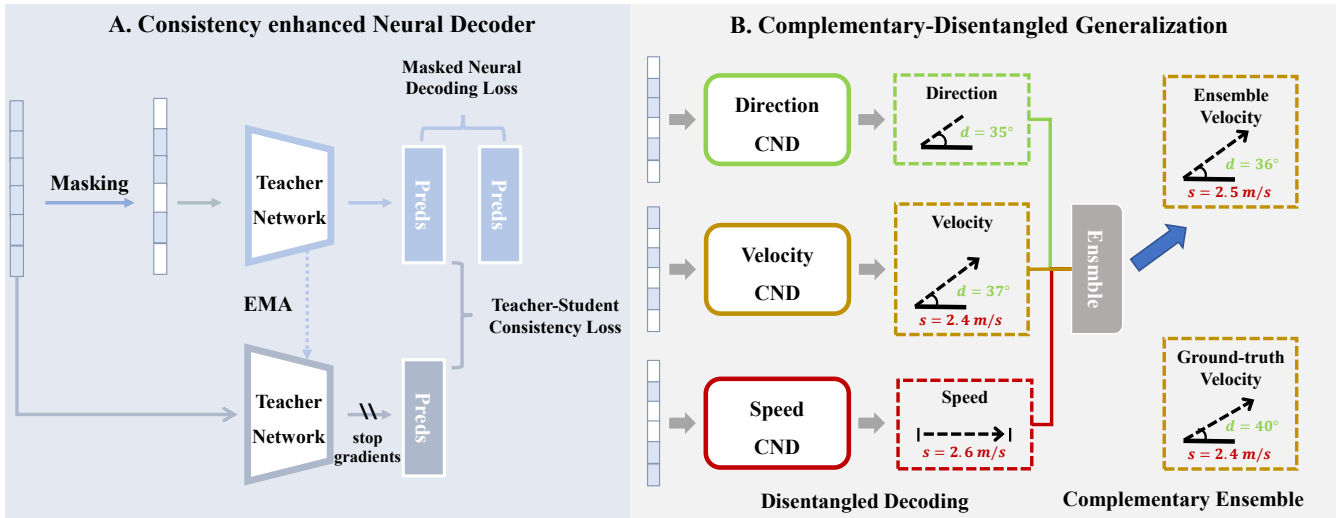


Fig. 3. The overall framework of Self-Supervised Consistency enhanced Disentangled Learning.

DyEnsemble and EvoEnsemble [12] use ensemble strategies to adapt to short-term variations; CAPTIVATE [24] incorporates a neuron locator for feature synchronization; SABLE [23] leverages adversarial alignment to learn invariant latent dynamics; LFDA [26] strengthens generalization by data augmentation and contrastive learning; MRNN [25] models multiple nonlinear mappings between neural signals and kinematics, leveraging historical data to capture long-term variability. However, a significant gap still remains: these methods typically ignore the heterogeneity of neural drift, which is often coupled with motor parameters due to the diverse tuning properties of individual neurons.

III. METHODS

In this section, we first formally define the problem of cross-day neural decoding. Then, inspired by recent advances in Self-Supervised Learning (SSL), we propose a strong neural decoding backbone termed Consistency enhanced Neural Decoder (CND), which leverages data augmentation to simulate cross-day neural drift and learn robust neural representations. Building upon neural preference theory and the success of Disentangled Representation Learning (DRL) in domain generalization, we further introduce Complementary-Disentangled Generalization (CDG). Lastly, by integrating CND and CDG, we present the overall framework of our Self-Supervised Consistency enhanced Disentangled Learning (SSCDL), as shown in Fig. 3.

A. Problem Definition

This paper addresses the challenge of maintaining stable BMI decoding performance across days using only single-session training data. We formulate this problem as a single-domain generalization task. Specifically, data collected on day_0 serve as the source domain $\mathcal{D}_s = \{(x_i, y_i)\}_{i=1}^{N_s}$, while data from a subsequent day day_k ($k > 0$) constitute an unseen target domain \mathcal{D}_t . Here, x_i denotes the recorded neural signal and y_i represents the corresponding kinematic variable. Our goal is to train a decoder F using only \mathcal{D}_s

such that it generalizes effectively to \mathcal{D}_t , thereby minimizing performance degradation caused by neural drift.

B. Consistency enhanced Neural Decoder

To learn neural representations robust to cross-day variability, we propose Consistency enhanced Neural Decoder (CND), a strong backbone model built upon SSL. CND adopts a teacher–student architecture, where the student network is optimized using both a supervised decoding loss and a self-supervised consistency loss, while the teacher network provides a stable reference for representation learning.

Let f_θ and $f_{\hat{\theta}}$ denote the student and teacher networks, respectively. At initialization, the parameters of the two networks are identical and randomly initialized. The training of the student network involves the following two losses:

- 1) **Masked Neural Decoding Loss.** To simulate realistic perturbations such as electrode instability and channel dropout, we introduce a data augmentation strategy termed Neural Signal Masking (NSM). During training, we randomly mask c channels of neural signals x_i in \mathcal{D}_s to produce an augmented input \hat{x}_i . The loss for predicting the corresponding kinematic variable y_i from the masked neural signal \hat{x}_i is defined as:

$$\mathcal{L}_{\text{MND}}(Z) = \frac{1}{N_s} \sum_{(x_i, y_i) \in \mathcal{D}_s} Z(f_\theta(\hat{x}_i), y_i), \quad (1)$$

where Z denotes a function that measures the prediction error.

- 2) **Teacher–Student Consistency Loss.** In addition to the decoding objective, we enforce a teacher–student consistency constraint to further enhance representation robustness. The consistency loss is defined as:

$$\mathcal{L}_{\text{TSC}}(Z) = \frac{1}{N_s} \sum_{(x_i, y_i) \in \mathcal{D}_s} Z(f_\theta(\hat{x}_i), f_{\hat{\theta}}(x_i)). \quad (2)$$

The key idea is that the student network, when given perturbed inputs, should produce predictions consistent

with those of the teacher network operating on the clean neural signal. This constraint encourages the student network to focus on stable neural structures rather than overfitting to specific neurons.

The student network is optimized using the combined loss:

$$\mathcal{L}_{\text{CND}}(Z) = \mathcal{L}_{\text{MND}}(Z) + \lambda \cdot \mathcal{L}_{\text{TSC}}(Z), \quad (3)$$

where λ balances the two loss terms.

As for the teacher network, its parameters are not directly optimized by gradient descent. Instead, its parameters are updated as an Exponential Moving Average (EMA) of the student parameters. Let k denote the training iteration, and \tilde{k} denote a warm-up threshold before which the teacher parameters directly follow the student parameters. The update rule is defined as follows:

$$\hat{\theta}^k = \begin{cases} \theta^k, & \text{if } k \leq \tilde{k}, \\ \frac{k-\tilde{k}}{k-\tilde{k}+1} \cdot \hat{\theta}^{k-1} + \frac{1}{k-\tilde{k}+1} \cdot \theta^k, & \text{otherwise.} \end{cases} \quad (4)$$

This EMA update produces a temporally smoothed teacher model that provides a more stable reference for consistency learning, which is particularly beneficial for noisy neural recordings.

C. Complementary-Disentangled Generalization

Neuroscience studies indicate that different neural populations encode distinct motor parameters, such as preferred direction and preferred speed [8], [27]. This suggests that neural drift may affect these motor parameters differently. Motivated by this observation, we propose Complementary-Disentangled Generalization (CDG), which improves cross-day generalization ability by explicitly modeling different motor parameters rather than treating velocity as a single unified prediction target. CDG involves two steps as follows:

- 1) **Disentangled Learning.** In CDG, velocity decoding is decomposed into three complementary components: overall velocity, direction, and speed. Instead of directly learning velocity as a unified output, we train separate models for each of the three components. This disentanglement allows the decoder to capture parameter-specific neural representations that may generalize differently across days. To measure the prediction discrepancy for different motor parameters, we define parameter-specific difference functions:

$$Z_{\text{direction}}(\hat{y}_i, y_i) = \|y_i\| \cdot \left(1 - \frac{\hat{y}_i \cdot y_i}{\|\hat{y}_i\| \cdot \|y_i\|}\right), \quad (5)$$

$$Z_{\text{speed}}(\hat{y}_i, y_i) = (\|\hat{y}_i\| - \|y_i\|)^2, \quad (6)$$

$$Z_{\text{velocity}}(\hat{y}_i, y_i) = \|\hat{y}_i - y_i\|^2. \quad (7)$$

These parameter-specific functions allow each model to focus on learning the neural representations most relevant to its corresponding motor parameter.

- 2) **Complementary Ensemble.** After obtaining predictions for direction, speed, and velocity, we combine them using a complementary ensemble strategy to produce the final decoding result. Let $\hat{y}_i^{\text{direction}}$, \hat{y}_i^{speed} ,

and $\hat{y}_i^{\text{velocity}}$ denote outputs of the three models, respectively. The final velocity prediction is computed as:

$$\tilde{y}_i = \frac{1}{2} \left(\|\hat{y}_i^{\text{speed}}\| \cdot \frac{\hat{y}_i^{\text{direction}}}{\|\hat{y}_i^{\text{direction}}\|} + \hat{y}_i^{\text{velocity}} \right). \quad (8)$$

The first term reconstructs velocity from the disentangled direction and speed predictions, while the second term corresponds to the holistic velocity prediction. By averaging these two complementary estimates, the ensemble leverages both parameter-specific specialization and holistic decoding, resulting in more accurate and robust predictions.

D. Overall Framework

To realize the introduced CDG mechanism, we employ Consistency enhanced Neural Decoder (CND) as the backbone model for each disentangled branch. Specifically, we construct three independent CND models to decode direction, speed, and overall velocity, respectively. Let f_{θ_d} , f_{θ_s} and f_{θ_v} denote the corresponding student networks. Let $f_{\hat{\theta}_d}$, $f_{\hat{\theta}_s}$ and $f_{\hat{\theta}_v}$ denote their teacher networks. Each branch is trained following the CND learning strategy described previously, where the loss functions of the three student networks are $\mathcal{L}_{\text{CND}}(Z_{\text{direction}})$, $\mathcal{L}_{\text{CND}}(Z_{\text{speed}})$ and $\mathcal{L}_{\text{CND}}(Z_{\text{velocity}})$, respectively. For a given neural signal x_i , the teacher networks produce predictions:

$$\hat{y}_i^{\text{direction}} = f_{\hat{\theta}_d}(x_i), \quad (9)$$

$$\hat{y}_i^{\text{speed}} = f_{\hat{\theta}_s}(x_i), \quad (10)$$

$$\hat{y}_i^{\text{velocity}} = f_{\hat{\theta}_v}(x_i). \quad (11)$$

Since the three branches are fully decoupled, they can be trained in parallel without introducing additional computational overhead. This modular design enables the framework to learn parameter-specific neural representations while maintaining training efficiency. After training, the teacher networks are used for inference. Their predictions are integrated using the complementary ensemble strategy defined before, yielding the final decoding output \tilde{y}_i .

IV. EXPERIMENTS

In this section, we first describe the experimental setup and then present a series of experiments designed to answer the following research questions (RQs):

- **RQ1:** How does our SSCDL compare with other state-of-the-art approaches?
- **RQ2:** How do individual modules in SSCDL contribute to the overall performance?
- **RQ3:** How robust is SSCDL under different experimental conditions?

A. Experimental Settings

1) *Datasets:* We evaluate our method on three non-human primate (NHP) datasets from UAN [15], namely Chewie, Mihili, and Jango. These datasets contain neural recordings collected during various upper-limb motor tasks, including center-out reaching (CO), random-target reaching

(RT), and isometric wrist movements (ISO). Neural signals were recorded from the primary motor cortex (M1) using a 96-channel Utah electrode array implanted in the hand or arm representation area of M1. The signals were digitized and bandpass filtered between 250 Hz and 5000 Hz. Multiunit threshold crossings from the go cue time to the trial end time were extracted. To obtain smoothed firing rates, spike counts were binned using 50 ms non-overlapping windows and convolved with a Gaussian kernel with a standard deviation of 100 ms. The resulting firing rates were then z-score normalized. For cross-day generalization experiments, trials from the first recording day are used as the source domain, while trials from subsequent days are treated as target domains. This setting allows us to evaluate the robustness of neural decoders under neural drift across days.

2) *Implementation Details*: The student and teacher networks of the proposed CND share the same architecture, implemented as a three-layer multilayer perceptron (MLP) with ReLU activations and dropout layers ($p = 0.25$). The MLP serves as the feature extractor, followed by a linear regression head that outputs the predicted kinematic variables. During training, the batch size is set to 128. We use the Adam optimizer with a learning rate of 5×10^{-3} , $\text{betas} = (0.9, 0.999)$, and $\text{weight_decay} = 5 \times 10^{-4}$. For neural signal masking, the number of masked channels c is set to 32 for the Chewie dataset and 16 for the Mihili and Jango datasets. The balance weight λ is set to 1. All experiments are repeated with five random seeds, and the reported results are averaged across these runs to ensure reliability and reduce the effect of stochastic variations.

3) *Metrics*: We evaluate decoding performance using two widely adopted metrics: Correlation Coefficient (CC) and Coefficient of Determination (R^2). For clarity, both metrics are reported in percentage form. In addition, we introduce two auxiliary metrics to evaluate the accuracy of the disentangled motor parameters. For speed prediction, we use Speed Error (SE), defined as the mean squared difference between the predicted and ground-truth speeds. For direction prediction, we use Angle Error (AE), defined as the mean angular difference (in degrees) between the predicted and ground-truth directions. The angular difference is computed using the minimal circular distance to account for the periodic nature of angles (e.g., 0° and 360° represent the same direction). These metrics provide further insight into the model ability to accurately decode different motor components.

B. Comparison with Other Methods (RQ1)

To evaluate the generalization capability of SSCDL, we conduct cross-day experiments using four recording days from each dataset. The earliest day (e.g., C0, J0, M0) is used as the source domain, while the subsequent days (C1–C3, J1–J3, M1–M3) serve as target domains spanning short (about one day), medium (about one month), and long (about three months) temporal intervals. Detailed descriptions of these sessions are provided in Table I. For comparison, we include a MLP baseline (denoted as Vanilla in Table II), several

TABLE I
DESCRIPTORS OF USED RECORDING SESSIONS

Dataset	Task	Session ID	Recording Date
Chewie	CO	C0	2016-09-27
		C1	2016-09-28
		C2	2016-10-12
		C3	2016-11-03
Jango	ISO	J0	2015-07-30
		J1	2015-07-31
		J2	2015-08-20
		J3	2015-10-29
Mihili	RT	M0	2013-12-07
		M1	2013-12-08
		M2	2014-01-15
		M3	2014-02-24

state-of-the-art generalization-based methods (KRSRL [10], MRNN [25], and LFDA [26]), and a recalibration-based method (WDGRL [21]). Note that the WDGRL results are obtained with retraining using target-domain data.

1) *Quantitative Results*: Table II reports the cross-day decoding performance of different methods across three datasets and three time spans. Overall, SSCDL consistently achieves the best or highly competitive results. In particular, it attains the highest CC and R^2 values in most settings, indicating strong robustness to neural drift. Notably, SSCDL even surpasses WDGRL, which utilizes unlabeled target-domain data for recalibration. This result highlights the effectiveness of our single-domain generalization framework.

2) *Qualitative Results*: Trajectory visualizations in Fig. 4 further demonstrate the advantages of CND and SSCDL in decoding stability. The vanilla MLP baseline produces highly erratic trajectories, while LFDA partially alleviates this issue but still exhibits noticeable overshooting. The CND backbone already generates smooth and compact trajectories, demonstrating the effectiveness of the proposed decoder architecture. SSCDL further maintains similarly stable and accurate trajectories that closely match the ground-truth movements. Such stable and precise decoding is essential for reliable and intuitive BMI control.

3) *Fine-grained Analysis*: Beyond overall velocity prediction, SSCDL explicitly disentangles velocity into direction and speed components, enabling a more detailed evaluation of decoding performance. We therefore report the Speed Error (SE) and the Angle Error (AE), and these results validate the effectiveness of the CDG module in capturing neural population shifts related to specific kinematic components.

C. Ablation Study (RQ2)

To analyze the contribution of each module in SSCDL, we conduct ablation experiments on session C0–C1. The results are summarized in Table III. From these results, we draw three key observations:

- 1) **CND provides a strong backbone**. Directly applying a decoder trained on the source domain leads to a significant performance drop on the target domain, indicating severe neural drift. By introducing neural sig-

TABLE II
COMPARISON OF CROSS-DAY GENERALIZATION PERFORMANCE

Dataset	Method	Time Span											
		Short				Medium				Long			
		CC \uparrow	R ² \uparrow	SE \downarrow	AE \downarrow	CC \uparrow	R ² \uparrow	SE \downarrow	AE \downarrow	CC \uparrow	R ² \uparrow	SE \downarrow	AE \downarrow
Chewie	Vanilla	79.39	59.29	0.0572	49.32	61.38	36.14	0.1010	62.11	59.57	32.82	0.1143	64.95
	KRSRL	80.54	60.61	0.0537	49.32	63.32	37.51	0.0996	60.31	59.02	32.15	0.1157	66.33
	MRNN	85.83	71.21	0.0422	47.15	70.01	48.77	0.0910	54.86	63.26	36.45	0.1089	64.41
	LFDA	87.91	72.35	0.0354	49.54	74.58	55.12	0.0768	51.59	68.14	41.21	0.1028	58.58
	CDNG	88.02	75.59	0.0300	42.15	76.26	57.91	0.0693	42.15	68.91	43.35	0.0906	57.77
	WDGRL	87.35	69.79	0.0397	45.89	78.82	65.92	0.0489	47.14	80.09	60.03	0.0619	49.90
	CND	87.07	75.58	0.0348	47.94	77.78	56.54	0.0819	48.83	66.31	43.31	0.1060	56.05
	SSCDL	88.21	77.56	0.0297	39.50	77.53	59.66	0.0639	49.54	67.63	45.40	0.0880	56.73
Mihili	Vanilla	68.07	45.84	30.05	53.07	67.66	49.17	33.05	54.07	48.49	22.48	54.36	61.33
	KRSRL	77.36	58.93	30.55	36.37	68.18	45.60	29.29	53.38	49.43	23.66	55.72	60.96
	MRNN	78.19	61.24	28.79	35.43	70.63	51.37	28.72	50.39	58.42	38.19	52.93	57.38
	LFDA	82.22	64.77	25.09	34.04	74.10	53.89	24.53	48.62	65.24	41.30	48.35	50.77
	CDNG	82.31	65.81	20.69	33.81	74.26	54.91	22.07	48.05	65.43	42.98	40.91	50.11
	WDGRL	82.52	65.49	27.06	35.82	71.28	55.58	23.89	48.69	55.21	40.56	50.22	51.89
	CND	81.67	65.76	21.64	34.40	74.24	54.60	22.80	48.45	64.88	42.78	42.48	50.67
	SSCDL	82.34	66.67	18.20	33.12	74.43	54.95	20.79	47.27	65.53	43.44	36.54	49.57
Jango	Vanilla	77.91	57.53	14.11	68.00	74.37	52.45	14.29	66.76	71.93	48.64	15.88	72.89
	KRSRL	77.08	56.61	14.22	68.70	74.23	51.55	15.02	65.94	72.36	49.36	15.79	70.50
	MRNN	81.03	65.79	10.19	68.16	77.46	65.30	12.76	67.84	74.35	53.79	14.16	70.05
	LFDA	84.39	70.66	8.576	69.90	82.28	65.30	9.877	70.55	80.86	58.43	13.48	72.57
	CDNG	84.71	71.63	7.117	67.91	81.51	66.01	8.870	65.76	65.43	42.98	11.42	69.57
	WDGRL	83.26	68.05	6.214	62.96	78.32	64.98	7.870	60.14	78.15	62.33	13.91	74.11
	CND	84.95	71.81	7.556	69.50	81.34	65.93	8.711	70.30	80.15	62.93	11.30	72.29
	SSCDL	85.04	72.26	6.798	65.70	81.72	67.58	8.526	62.96	80.48	63.57	9.772	66.03

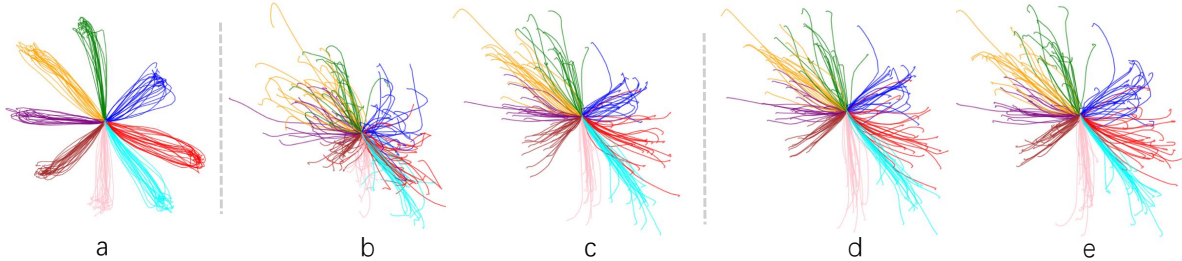


Fig. 4. Comparison of motor trajectories decoded by different models. Trajectories are color-coded by trial directions. (a) Ground truth; (b) Vanilla; (c) LFDA; (d) CND; (e) SSCDL.

nal masking (NSM), teacher-student consistency constraint, and the EMA-based teacher update, CND substantially improves cross-day decoding performance. These improvements demonstrate that CND effectively learns more stable neural representations. Combined with the results in Table II, this confirms SSD as a robust backbone for cross-day neural decoding.

- 2) **CDG improves generalization through disentangled learning.** CDG enhances generalization by decomposing velocity into complementary components: direction, speed, and overall velocity. The velocity branch provides balanced predictions, while the direction and speed branches specialize in their respective motor parameters. Even when using the baseline backbone, CDG improves performance compared with the single-velocity decoder. When combined with CND, the improvements become more pronounced, indicating that

CDG is compatible with robust representation learning.

- 3) **Both modules are essential and beneficial.** While CND and CDG individually improve cross-day generalization, their integration in SSCDL achieves the best overall performance. SSD contributes stable neural representations that mitigate neural drift, whereas CDG introduces parameter-specific specialization through disentangled learning. These two modules together provide a robust framework for long-term neural decoding, enabling more stable and accurate BMI control.

D. Robustness Analysis (RQ3)

1) Robustness Across Tasks, Subjects, and Timescales:

As shown in Table II, SSCDL demonstrates consistent generalization across different neural decoding tasks (CO, ISO, and RT), multiple subjects, and temporal spans of up to three months without any recalibration. These results indicate that

TABLE III
ABLATION RESULTS OF SSCDL ON SESSION C0–C1

Method	CC \uparrow	R^2 \uparrow	AE \downarrow	SE \downarrow
Baseline	79.39	59.29	49.32	0.0572
Baseline + EMA	80.01	60.53	48.37	0.0555
Baseline + NSM	84.35	71.79	48.24	0.0457
Baseline + NSM + EMA	86.35	73.79	48.04	0.0421
SSD	87.07	75.58	47.94	0.0348
Baseline (Velocity)	79.39	59.29	49.32	0.0572
Baseline (Direction)	N/A	N/A	44.37	0.1040
Baseline (Speed)	N/A	N/A	80.54	0.0385
CoDG	84.20	68.52	43.95	0.0374
SSD (Velocity)	87.07	75.58	47.94	0.0348
SSD (Direction)	N/A	N/A	39.69	0.1687
SSD (Speed)	N/A	N/A	98.81	0.0312
SSCDL	88.21	77.56	39.50	0.0297




Fig. 5. Robustness Analysis on Number of Masked Channels.

the proposed framework maintains stable decoding performance under diverse experimental conditions, highlighting its scalability and suitability for long-term BMI applications.

2) *Robustness to Masked Channels*: To further evaluate robustness to neural signal corruption, we analyze the effect of channel masking on session C1 (Fig. 5). The best performance is achieved when 32 channels are masked, suggesting that moderate masking can improve representation learning by encouraging the model to rely on distributed neural information. When the number of masked channels increases to 48, performance slightly decreases but remains stable. However, masking 64 channels leads to a substantial degradation in performance, indicating the practical limit of tolerable channel occlusion.

3) *Robustness on Real-World Human Dataset*: We further evaluate the robustness of SSCDL on a real-world **human robotic effector dataset** from the FALCON benchmark [35]. In this dataset, the source domain consists of recordings from 01-01, where the participant performs a virtual arm reaching task. The results in Table IV demonstrate that SSCDL consistently outperforms existing methods across multiple target domains. These findings further validate the robustness and practical applicability of the proposed framework for real-world human–robot interaction scenarios.

V. DISCUSSION

1) *Computational Cost*: While SSCDL introduces a three-model ensemble, the computational overhead is moderate:

TABLE IV
ROBUSTNESS ANALYSIS ON THE HUMAN DATASET

Target Domain	Vanilla	MRNN	LFDA	SSCDL
01-08	10.43	11.42	12.35	13.54
01-20	9.14	10.53	11.37	12.79
02-09	8.76	9.56	11.10	12.03

TABLE V
COMPUTATIONAL COST ANALYSIS

Metric	Static	MRNN	LFDA	Ours (Parallelized)
Parameters	11k	42k	28k	31k
Inference Time	0.17s	0.41s	0.27s	0.24s

inference for 1000 samples takes about 0.24s, with 31k parameters, comparable to or smaller than existing methods (Table V). These characteristics make SSCDL practically feasible for current BMI devices.

2) *Limitations*: Currently, SSCDL operates offline, limiting immediate application to real-time control. Future work will focus on enabling low-latency online decoding, reducing model complexity via pruning or distillation, and integrating multi-modal sensory feedback to further improve generalization across tasks, subjects, and sessions.

3) *BMIs for Human-Centered Robotics and Human-Robot Interaction*: BMIs enable long-term, human-centered robotic control by translating neural activity into actionable commands for prostheses, exoskeletons, and teleoperation systems. By mitigating neural drift and parameter-specific variability, SSCDL enables stable and fine-grained neural decoding. The learned disentangled representations further support long-term cross-day and cross-subject control, facilitating adaptive and continuous human–robot interaction.

VI. CONCLUSION

In this work, we present SSCDL, a neural decoding framework designed to improve cross-day robustness using only single-day training data. The proposed framework integrates CND with CDG to learn stable and transferable neural representations. CND leverages a self-supervised teacher–student paradigm to enforce prediction consistency under neural perturbations, while CDG decomposes velocity into complementary components (velocity, direction, and speed) to better capture parameter-specific neural dynamics. By mitigating neural drift and parameter-specific variability, SSCDL enables stable and fine-grained neural decoding while learning disentangled neural representations. This capability supports long-term cross-day and cross-subject BMI control, facilitating adaptive and continuous human–robot interaction in applications such as neurorehabilitation, human motor augmentation, and human-centered robotics.

VII. ACKNOWLEDGMENT

This work was supported in part by the National Natural Science Foundation of China under Grant (U25D9015, 62336007), in part by the Starry Night Science Fund of the

Zhejiang University Shanghai Institute for Advanced Study under Grant SN-ZJU-SIAS-002, in part by the Fundamental Research Funds for the Central Universities. The authors acknowledge the use of Gemini (Google) for text polishing to improve the readability of this manuscript.

REFERENCES

- [1] G. Lisi, M. Hamaya, T. Noda, and J. Morimoto, "Dry-wireless eeg and asynchronous adaptive feature extraction towards a plug-and-play co-adaptive brain robot interface," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 959–966.
- [2] K. Shi, R. Huang, F. Mu, Z. Peng, K. Huang, Y. Qin, X. Yang, and H. Cheng, "A novel multimodal human-exoskeleton interface based on EEG and semg activity for rehabilitation training," in *2022 International Conference on Robotics and Automation, ICRA 2022, Philadelphia, PA, USA, May 23-27, 2022*. IEEE, 2022, pp. 8076–8082.
- [3] H. J. Choi, S. Das, S. Peng, R. Bajcsy, and N. Figueroa, "On the feasibility of eeg-based motor intention detection for real-time robot assistive control," in *IEEE International Conference on Robotics and Automation, ICRA 2024, Yokohama, Japan, May 13-17, 2024*. IEEE, 2024, pp. 5592–5599.
- [4] L. R. Hochberg, M. D. Serruya, G. M. Fries, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue, "Neuronal ensemble control of prosthetic devices by a human with tetraplegia," *Nature*, vol. 442, no. 7099, pp. 164–171, 2006.
- [5] L. R. Hochberg, D. Bacher, B. Jarosiewicz, N. Y. Masse, J. D. Simeral, J. Vogel, S. Haddadin, J. Liu, S. S. Cash, P. Van Der Smagt, et al., "Reach and grasp by people with tetraplegia using a neurally controlled robotic arm," *Nature*, vol. 485, no. 7398, pp. 372–375, 2012.
- [6] C. Pandarinath, P. Nuyujukian, C. H. Blabe, B. L. Sorice, J. Saab, F. R. Willett, L. R. Hochberg, K. V. Shenoy, and J. M. Henderson, "High performance communication by people with paralysis using an intracortical brain-computer interface," *elife*, vol. 6, p. e18554, 2017.
- [7] C. E. Bouton, A. Shaikhoui, N. V. Annetta, M. A. Bockbrader, D. A. Friedenber, D. M. Nielson, G. Sharma, P. B. Sederberg, B. C. Glenn, W. J. Mysiw, et al., "Restoring cortical control of functional movement in a human with quadriplegia," *Nature*, vol. 533, no. 7602, pp. 247–250, 2016.
- [8] A. P. Georgopoulos, A. B. Schwartz, and R. E. Kettner, "Neuronal population coding of movement direction," *Science*, vol. 233, no. 4771, pp. 1416–1419, 1986.
- [9] N. Kadmon Harpaz, D. Ungarish, N. G. Hatsopoulos, and T. Flash, "Movement decomposition in the primary motor cortex," *Cerebral cortex*, vol. 29, no. 4, pp. 1619–1633, 2019.
- [10] Y. Li, Y. Qi, Y. Wang, Y. Wang, K. Xu, and G. Pan, "Robust neural decoding by kernel regression with siamese representation learning," *Journal of Neural Engineering*, vol. 18, no. 5, p. 056062, 2021.
- [11] Y. Qi, X. Zhu, K. Xu, F. Ren, H. Jiang, J. Zhu, J. Zhang, G. Pan, and Y. Wang, "Dynamic ensemble bayesian filter for robust control of a human brain-machine interface," *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 12, pp. 3825–3835, 2022.
- [12] X. Zhu, Y. Qi, G. Pan, and Y. Wang, "Tracking functional changes in nonstationary signals with evolutionary ensemble bayesian model for robust neural decoding," *Advances in Neural Information Processing Systems*, vol. 35, pp. 22 576–22 588, 2022.
- [13] X. Gao, D. Rong, and Q. He, "Cross-modal alignment between visual stimuli and neural responses in the visual cortex," in *Journal of Physics: Conference Series*, vol. 3147, no. 1. IOP Publishing, 2025, p. 012026.
- [14] 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," *Nature neuroscience*, vol. 23, no. 2, pp. 260–270, 2020.
- [15] X. Ma, F. Rizzoglio, K. L. Bodkin, E. Perreault, L. E. Miller, and A. Kennedy, "Using adversarial networks to extend brain computer interface decoding accuracy over time," *elife*, vol. 12, p. e84296, 2023.
- [16] H. Duffau, "Brain plasticity and reorganization before, during, and after glioma resection," *Glioblastoma*, pp. 225–236, 2016.
- [17] D. Sussillo, S. D. Stavisky, J. C. Kao, S. I. Ryu, and K. V. Shenoy, "Making brain-machine interfaces robust to future neural variability," *Nature communications*, vol. 7, no. 1, p. 13749, 2016.
- [18] A. B. Ajiboye, F. R. Willett, D. R. Young, W. D. Memberg, B. A. Murphy, J. P. Miller, B. L. Walter, J. A. Sweet, H. A. Hoyen, M. W. Keith, et al., "Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration," *The Lancet*, vol. 389, no. 10081, pp. 1821–1830, 2017.
- [19] S. Wen, A. Yin, T. Furlanello, M. G. Perich, L. E. Miller, and L. Itti, "Rapid adaptation of brain-computer interfaces to new neuronal ensembles or participants via generative modelling," *Nature biomedical engineering*, vol. 7, no. 4, pp. 546–558, 2023.
- [20] A. Farshchian, J. A. Gallego, J. P. Cohen, Y. Bengio, L. E. Miller, and S. A. Solla, "Adversarial domain adaptation for stable brain-machine interfaces," in *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- [21] Y. Qi, L. Ding, Y. Wang, and G. Pan, "Learning robust features from nonstationary brain signals by multiscale domain adaptation networks for seizure prediction," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 14, no. 3, pp. 1208–1216, 2021.
- [22] J. Wei, D. Rong, X. Zhu, Q. He, and Y. Wang, "Speed-enhanced subdomain alignment for long-term stable neural decoding in brain-computer interfaces," in *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2024, pp. 3832–3837.
- [23] J. Jude, M. Perich, L. Miller, and M. Hennig, "Robust alignment of cross-session recordings of neural population activity by behaviour via unsupervised domain adaptation," in *International Conference on Machine Learning*. PMLR, 2022, pp. 10462–10475.
- [24] J. Jude, M. G. Perich, L. E. Miller, and M. H. Hennig, "Capturing cross-session neural population variability through self-supervised identification of consistent neuron ensembles," in *NeurIPS Workshop on Symmetry and Geometry in Neural Representations*. PMLR, 2023, pp. 234–257.
- [25] D. Sussillo, S. D. Stavisky, J. C. Kao, S. I. Ryu, and K. V. Shenoy, "Making brain-machine interfaces robust to future neural variability," *Nature communications*, vol. 7, p. 13749, 2016. [Online]. Available: <https://doi.org/10.1038/ncomms13749>
- [26] S.-H. Yang, C.-J. Huang, and J.-S. Huang, "Increasing robustness of intracortical brain-computer interfaces for recording condition changes via data augmentation," *Computer Methods and Programs in Biomedicine*, vol. 251, p. 108208, 2024.
- [27] Y. Inoue, H. Mao, S. B. Suway, J. Orellana, and A. B. Schwartz, "Decoding arm speed during reaching," *Nature communications*, vol. 9, no. 1, p. 5243, 2018.
- [28] L. Tran, X. Yin, and X. Liu, "Disentangled representation learning gan for pose-invariant face recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1415–1424.
- [29] H. Zhang, Y.-F. Zhang, W. Liu, A. Weller, B. Schölkopf, and E. P. Xing, "Towards principled disentanglement for domain generalization," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 8024–8034.
- [30] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 16 000–16 009.
- [31] A. Jaiswal, A. R. Babu, M. Z. Zadeh, D. Banerjee, and F. Makedon, "A survey on contrastive self-supervised learning," *Technologies*, vol. 9, no. 1, p. 2, 2020.
- [32] D. Rong, H. Dong, X. Gao, J. Wei, D. Hong, Y. Hao, Q. He, and Y. Wang, "Improving unsupervised task-driven models of ventral visual stream via relative position predictivity," in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 47, 2025.
- [33] J. Wei, D. Rong, D. Hong, Z. Zhang, X. Zhu, Q. He, and Y. Wang, "Complementary-disentangled neural generalization: a robust framework for stable brain-computer interfaces," in *2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2025, pp. 4230–4235.
- [34] Y. Dong, D. Hu, S. Wang, and J. He, "Decoder calibration framework for intracortical brain-computer interface system via domain adaptation," *Biomedical Signal Processing and Control*, vol. 81, p. 104453, 2023.
- [35] B. M. Karpowicz, J. Ye, C. Fan, P. Tostado-Marcos, F. Rizzoglio, C. Washington, T. Scodeler, D. de Lucena, S. R. Nason-Tomaszewski, M. J. Mender, et al., "Few-shot algorithms for consistent neural decoding (falcon) benchmark," *Advances in Neural Information Processing Systems*, vol. 37, pp. 76 578–76 615, 2024.