MAC-DA: Manifold and Anchored Conditional Domain Adaptation for EEG-based MI-BCI

Zhengyang Ding

[a]*School of Information and Communication Engineering, University of Electronic Science and Technology of China, Chengdu, 611731, Sichuan, China*

Abstract

Electroencephalogram (EEG)-based motor imagery brain-computer interfaces (MI-BCIs) suffer from severe cross-subject distribution shifts that degrade decoding performance and hinder clinical deployment. We propose Manifold and Anchored Conditional Domain Adaptation (MAC-DA), a multi-space progressive alignment framework that addresses this challenge across four complementary levels: adaptive physical manifold normalization via Entropy-Aware Euclidean Alignment (EAEA) with automatic gating, spatiotemporal feature extraction through a hybrid convolution-attention backbone, entropy-aware conditional adversarial alignment for global distribution matching, and class-conditional CORAL (ccCORAL) for local semantic anchoring. Evaluated under a strict leave-one-subject-out (LOSO) protocol with five independent random seeds to ensure reproducibility, MAC-DA achieves state-of-the-art cross-subject four-class MI decoding accuracy on two public benchmarks: 69.77% on BCI Competition IV 2a (+9.48 pp over the best baseline, $p < 0.001$) and 53.70% on Weibo2014 (+7.30 pp, $p = 0.002$), significantly outperforming all competing domain adaptation methods on both datasets ($p \leq 0.041$). Ablation studies further confirm that ccCORAL significantly suppresses cross-seed training variance ($p = 0.0214$), serving as a critical stabilizer for adversarial optimization. These results demonstrate that MAC-DA provides a promising step toward reducing calibration burden and improving the practical reliability of cross-subject MI-BCI systems.

Keywords: Motor Imagery, Brain-Computer Interface, Domain Adaptation, Electroencephalogram

1. Introduction

Brain-Computer Interface (BCI) technology offers a pathway for individuals with motor impairments to interact with external environments by directly decoding neural activity [1, 2]. Among various paradigms, electroencephalography (EEG)-based motor imagery (MI) is particularly prominent due to its non-invasiveness, portability, and cost-effectiveness [3, 4]. By capturing event-related desynchronization/synchronization (ERD/ERS) patterns, MI-BCIs have shown considerable potential in wheelchair control, robotic arm operation, and neurorehabilitation [5, 6, 7].

Despite progress in laboratory settings, the clinical translation of MI-BCIs is hindered by a critical lack of cross-subject generalization [8, 9]. Variations in anatomical structures, electrode impedances, and cognitive strategies induce severe distribution shifts in EEG signals across individuals [10, 11, 12]. The severity of this challenge is underscored by a recent large-scale reproducibility study: the MOABB benchmark [72], which systematically evaluates 30 machine learning pipelines across 36 public EEG datasets, reports that the best-performing pipeline achieves only 43.3% accuracy on a four-class motor imagery task from the Weibo2014 dataset [71]—a figure also highlighted by recent work from Meta AI on non-invasive brain-to-text decoding [73] as evidence that EEG-based motor imagery remains far from solved. Consequently, subject-specific decoding models experience significant performance degradation on new users, necessitating time-consuming and fatigue-inducing calibration processes [13, 14].

Unsupervised domain adaptation (UDA) has emerged as a promising solution to address this challenge [15, 16, 17]. However, existing deep domain adaptation methods face three primary limitations in MI-BCI applications:

1) **Neglect of physical manifold shift:** Most methods align features solely in high-level semantic spaces, ignoring underlying Riemannian geometric shifts caused by physical sensing differences (e.g., skull thickness, impedance) [18, 19].

2) **Conflict between marginal and local alignment:** Forcibly aligning global marginal distributions often disrupts the fine-grained intra-class topology of MI-EEG signals, leading to blurred decision boundaries or negative transfer [20, 21].

3) **Instability of adversarial optimization:** Adversarial training on low-signal-to-noise ratio EEG data is prone to collapse, resulting in

high performance variance across random initializations [22], which is detrimental for reliable clinical use.

To overcome these limitations, we propose Manifold and Anchored Conditional Domain Adaptation (MAC-DA), a multi-space progressive alignment framework designed to balance accuracy and robustness. Rather than relying on a single alignment strategy, MAC-DA systematically resolves domain shifts through a unified, end-to-end architecture. First, an Entropy-Aware Euclidean Alignment (EAEA) process is applied at the input level to adaptively mitigate macroscopic baseline drifts caused by sensing discrepancies; EAEA incorporates a hard gating mechanism based on the log-determinant of the covariance matrix to automatically detect and bypass subjects with atypical covariance structures, thereby preventing negative transfer. Subsequently, a hybrid convolutional-attention backbone extracts high-fidelity spatiotemporal representations. To effectively align these representations, we introduce an entropy-aware adversarial learning mechanism that implicitly matches the global marginal distributions across domains. Crucially, to prevent the disruption of intra-class topology during global alignment, we propose a novel class-conditional covariance alignment (ccCORAL) module. By explicitly anchoring the second-order statistics of local semantic spaces using pseudo-labels and a warmup strategy, ccCORAL acts as a localized constraint to prevent optimization shortcuts.

The main contributions of this paper are summarized as follows:

1) **A multi-space progressive UDA framework:** We propose MAC-DA, a unified architecture that fundamentally resolves cross-domain EEG shifts by seamlessly integrating adaptive physical manifold normalization (EAEA), global semantic adversarial alignment, and local manifold anchoring. This comprehensive approach addresses distribution discrepancies spanning from macroscopic sensor-level variations to microscopic intra-class feature dynamics.

2) **Entropy-Aware Euclidean Alignment with hard gating:** We propose EAEA, which extends classical EA with an automatic gating mechanism based on the log-determinant of the covariance matrix. By detecting subjects whose covariance manifolds fall below a statistical threshold ($\mu - \sigma$ rule), EAEA adaptively disables alignment for incompatible subjects, converting EA-induced negative transfer into preserved baseline performance.

3) **Introduction of ccCORAL as a local stabilizer:** We introduce ccCORAL to constrain local manifolds explicitly. Statistical analyses demonstrate that it significantly suppresses the cross-initialization variance of adversarial training ($p = 0.0214$), thereby substantially bridging the gap between accuracy and training stability.

4) **Synergistic improvements in accuracy and robustness:** Comprehensive cross-subject evaluations on two public MI-EEG benchmarks demonstrate that MAC-DA achieves an average accuracy of 69.77% on BCI Competition IV 2a (+9.48 pp over baselines, $p < 0.001$) and 53.70% on Weibo2014 (+7.30 pp, $p = 0.002$). All improvements over competing domain adaptation methods are statistically significant on both datasets. This consistent cross-dataset advantage demonstrates its robust potential for calibration-free, clinical-grade BCI systems.

2. Related Work

The core challenge in cross-subject motor imagery (MI-EEG) decoding lies in the highly heterogeneous distribution shifts of EEG signals. In recent years, the research paradigm in this field has evolved from shallow feature engineering to deep feature representation, and further to multi-dimensional spatial domain adaptation. This section reviews the research progress across four levels closely related to the proposed MAC-DA architecture and elucidates the entry point of this study.

2.1. Spatiotemporal Feature Representation and Network Architecture Evolution

Extracting highly discriminative and robust EEG representations is a prerequisite for achieving effective cross-domain transfer. Early studies primarily relied on domain expert knowledge for shallow feature engineering. Common Spatial Pattern (CSP) [23] and its frequency-band optimized variants (e.g., FBCSP [24]) extract spatial filters by maximizing inter-class variance, achieving relatively stable performance in within-subject classification. Additionally, methods based on Riemannian geometry utilize geodesic distances of covariance matrices on the symmetric positive definite (SPD) manifold for classification [25, 26]. However, these traditional manifold methods heavily rely on second-order statistics, often ignoring fine-grained temporal dynamic

features in EEG signals, and their feature expression capabilities are limited when facing complex high-dimensional heterogeneity across subjects [27].

With the rise of deep learning, end-to-end EEG feature extraction architectures have significantly advanced the field. Early representative works, such as EEGNet [28] and Deep/Shallow ConvNet [29], captured spatial and spectral features of MI-EEG while maintaining a small number of parameters by designing specialized depthwise and separable convolution kernels. Subsequent studies further introduced recurrent neural networks (RNNs) and their variants (e.g., CNN-LSTM [30]) to model the temporal dynamics of EEG. However, most of these classical networks are optimized for the data distribution of specific subjects and are prone to overfitting in cross-subject scenarios [31].

In recent years, following the architectural evolution in computer vision and natural language processing, architectures based on attention mechanisms and graph neural networks have become mainstream to capture more global neural activation patterns. EEG-Conformer [32] combines convolution with self-attention mechanisms to unify local spatiotemporal features and global dependencies. LMDA-Net [33] proposes a lightweight multi-dimensional attention network specifically designed for the decoupling of EEG channels and frequency bands. AGTCNet [34] utilizes the spatial topology of electrodes to construct inductive biases. ATCNet [35], adopted as the backbone network in this paper, achieves an optimal balance in dynamically capturing motor intentions through the fusion of sliding-window attention and multi-scale temporal convolutional networks (TCN). Recently, large-scale pre-trained models based on Transformers [36] and self-supervised learning methods [37] have also begun to show potential in EEG decoding, but their generalization capabilities in cross-subject scenarios require further validation [38].

Although the aforementioned advanced feature extractors perform excellently within closed domains, recent studies [39, 40] indicate that relying solely on powerful feature extractors still encounters significant performance degradation when facing unseen target subjects. This confirms that advanced network architectures cannot spontaneously eliminate cross-subject distribution shifts, which necessitates the introduction of explicit domain alignment mechanisms beyond the spatiotemporal feature space.

2.2. Preprocessing and Geometric Alignment in Physical Manifold Space

The fundamental physical source of EEG distribution shifts stems from inherent physical and physiological differences among subjects, such as skull thickness, cortical folding, cortical dipole orientation, and electrode contact impedance [41]. These differences cause EEG covariance matrices to exhibit severe macroscopic baseline drifts on the underlying physical manifold [42].

To directly eliminate such sensor-level physical differences at the data input end, spatial alignment techniques have emerged. Early studies attempted to align covariance matrices directly on the Riemannian manifold [43], but the computational complexity was relatively high. Subsequently, Euclidean Alignment (EA) [44] was proposed as an efficient preprocessing paradigm. By calculating the average covariance matrix of a subject's resting state or all trials, and normalizing the covariance matrices of all subjects to identity matrices via an inverse square root transformation ($\bar{R}^{-1/2}$), EA effectively eliminates macroscopic marginal distribution discrepancies in Euclidean space.

In recent related studies, EA frequently appears as a front-end module for deep domain adaptation. For example, EA has been combined with Transfer Component Analysis (TCA) for cross-domain feature dimensionality reduction [45]. Filter Bank Euclidean Alignment (FBEA) [46] was proposed for specific paradigms to capture harmonic components. Adaptive Euclidean Alignment (Adaptive EA) for Source-Free Domain Adaptation (SFDA) aligns target and source subject distributions by learning a projection matrix [47].

Although EA significantly elevates the performance baseline of cross-subject decoding, it suffers from two fundamental limitations [48]. First, it is a label-agnostic global alignment method that aggregates all imagery classes to calculate the global covariance, which cannot guarantee the preservation of task-specific inter-class decision boundaries and may blur fine-grained intra-class semantic structures [49]. Second, and more critically, EA implicitly assumes that all subjects possess well-conditioned covariance manifolds; for subjects with atypical or collapsed covariance structures, the inverse square root transformation amplifies noise rather than removing it, leading to negative transfer. Therefore, in the MAC-DA framework, we propose Entropy-Aware Euclidean Alignment (EAEA), which extends EA with an automatic hard gating mechanism based on the log-determinant of the covariance matrix. EAEA adaptively disables alignment for subjects whose covariance quality falls below a statistical threshold, thereby serving as a robust first

physical defense line whose residual limitations are compensated by subsequent high-dimensional semantic alignment.

2.3. Deep Domain Adversarial Alignment in Global Semantic Space

To align the source and target domains within deep network structures, transfer learning based on unsupervised domain adaptation (UDA) has been widely adopted in the BCI field. Among these, methods based on adversarial learning hold a dominant position. Domain-Adversarial Neural Network (DANN) [50] introduces a Gradient Reversal Layer (GRL) and employs a min-max game to compel the feature extractor to generate features indistinguishable by the domain discriminator, thereby achieving marginal distribution alignment.

However, merely aligning marginal distributions is prone to negative transfer, as complex decision boundaries in the target domain may be distorted during forced alignment. To this end, recent research has shifted towards more refined conditional distribution alignment. Conditional Domain Adversarial Network (CDAN) and its entropy-conditioned variant (CDAN+E) [51] successfully capture joint distribution shifts by performing a multilinear mapping between high-dimensional features and the classifier's predicted probabilities. In the BCI field, numerous advanced architectures have been derived from this adversarial concept. For instance, Spatial and Conditional Domain Adaptation Network (SCDAN) [52] proposes a dual-discriminator architecture to separately address spatial shifts caused by brain structures and conditional shifts caused by imagery strategies. MI-DABAN [53] introduces a dual-attention mechanism to dynamically adjust feature weights. Additionally, multi-subdomain adaptation methods [54] and self-training strategies based on source-free domain adaptation [55] continue to broaden the boundaries of adversarial learning.

Although these methods can implicitly reshape macroscopic decision boundaries, the dynamic game between the generator and the discriminator suffers from stability issues when processing EEG signals with extremely low signal-to-noise ratios (SNR) and high noise characteristics [56, 57]. This uncertainty in adversarial optimization causes models to exhibit substantial performance fluctuations and cross-seed variance across different subjects or even under different random initializations for the same subject [58]. In medical-grade clinical applications, such a lack of predictability and robustness severely limits the feasibility of practical deployment.

2.4. Second-Order Statistic Anchoring in Local Semantic Space

To overcome the instability and high variance of pure adversarial training, manifold alignment methods based on statistical moment matching have re-emerged in the research spotlight. Unlike implicit alignment that attempts to deceive a discriminator, statistical methods such as Maximum Mean Discrepancy (MMD) [59] and Wasserstein distance [60] achieve alignment by explicitly calculating distance metrics between distributions.

Among such methods, Deep CORAL (Deep Correlation Alignment) [61] demonstrates unique advantages. By penalizing the Frobenius norm of the second-order statistics (covariance matrices) of deep features from the source and target domains, it minimizes the distance between distribution shapes (feature correlations). However, state-of-the-art studies indicate that simply applying global CORAL or global MMD to pull overall distributions closer can destroy the fine-grained intra-class topological structures inherent in MI-BCI signals (e.g., the unique manifold spaces corresponding to left-hand and right-hand imagery) [62, 63]. To preserve class-specific geometric structures, contrastive adaptation based on pseudo-labels and local class-conditional alignment strategies have emerged as cutting-edge trends [64, 65].

Most existing studies treat CORAL either as an independent enhancement technique or mechanically superimpose it with adversarial losses, neglecting their complementarity in gradient optimization. In the MAC-DA framework, we redefine class-conditional CORAL (ccCORAL) as a manifold stabilizing anchor in the high-dimensional local semantic space. By combining high-confidence pseudo-labels with a warmup mechanism, ccCORAL explicitly calculates and locks the intra-class second-order covariance for each specific class. This localized rigid constraint provides a deterministic optimization direction for adversarial gradients, preventing the network from taking shortcuts and fundamentally suppressing the statistical variance of the training process.

In summary, existing single-alignment strategies struggle to simultaneously balance the minimization of global distribution discrepancies and the stabilization of local manifolds. To this end, the proposed MAC-DA framework systematically integrates a multi-dimensional progressive alignment mechanism from underlying sensors to high-level semantics. While filling this research gap, it provides a highly promising solution for the stability and high accuracy of clinical-grade BCIs.

3. Methodology

This section details the proposed Manifold and Anchored Conditional Domain Adaptation (MAC-DA) framework. The core philosophy of MAC-DA is that cross-subject EEG distribution shifts are not derived from a single source, but rather a composite problem compounding from underlying physical sensors to high-level semantic representations. Therefore, these shifts must be progressively eliminated across multiple spaces. In the following, we first formalize the problem definition, then provide an overview of the framework's overall architecture and the design motivations of each module, and finally elaborate on the technical details.

3.1. Problem Definition

In the unsupervised domain adaptation (UDA) scenario for cross-subject MI-BCI, we have a labeled dataset $\mathcal{D}_{src} = \{(x_i^s, y_i^s)\}_{i=1}^{N_s}$ from multiple source subjects, and an unlabeled dataset $\mathcal{D}_{tgt} = \{x_j^t\}_{j=1}^{N_t}$ from a target subject, where $x \in \mathbb{R}^{C\times T}$, and $y \in \{1, 2, \ldots, K\}$. The source and target domains share the same label space, but exhibit shifts in both the marginal distribution $P(X^s) \neq P(X^t)$ and the conditional distribution $P(Y|X^s) \neq P(Y|X^t)$. Our goal is to learn a classification model that generalizes well on the target domain while maintaining training stability.

Under the Leave-One-Subject-Out (LOSO) cross-validation protocol, one subject is selected as the target domain in each fold. The training session data from the remaining eight subjects constitute the source domain training set, while their evaluation session data form the validation set. The training session data of the target subject serves as the unlabeled target domain training data for domain adaptation, and the evaluation session data of the target subject is used as the final test set.

3.2. Framework Overview

As analyzed in Section 2, cross-subject EEG distribution shifts originate from three levels. MAC-DA designs specific alignment mechanisms for each level. Fig. 1 illustrates the overall pipeline of MAC-DA. Its core components and their design motivations are as follows:

1) **Adaptive Physical Manifold Normalization:** Physical differences such as skull thickness and electrode impedance cause macroscopic baseline drifts in the underlying manifold of EEG covariance matrices.

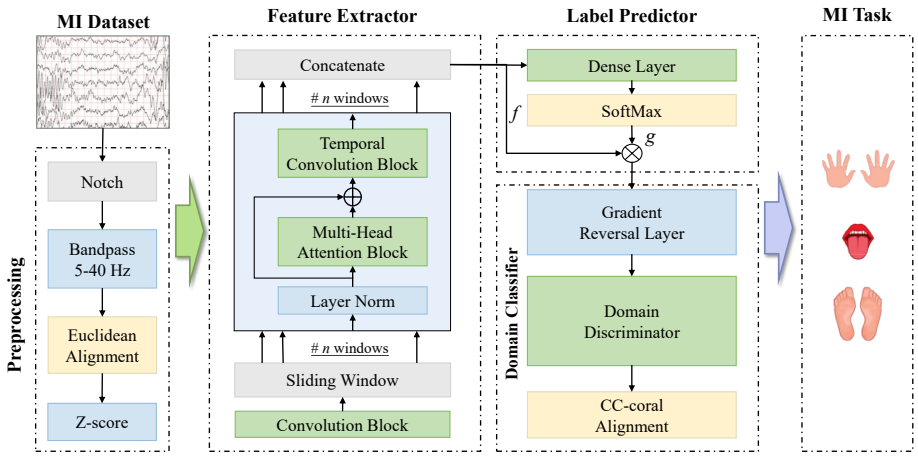

Figure 1: The schematic pipeline of the proposed MAC-DA framework. The process initiates with physical manifold normalization (Notch, Bandpass, and Entropy-Aware Euclidean Alignment with automatic gating) to adaptively mitigate sensor-level drifts. Subsequently, high-dimensional spatiotemporal features are extracted via a hybrid backbone and fed into a dual-branch adaptation module, where global semantic shifts and local manifold discrepancies are synergistically resolved through adversarial learning and class-conditional covariance alignment.

MAC-DA employs a preprocessing pipeline combining Z-score standardization and Entropy-Aware Euclidean Alignment (EAEA) at the data input end. EAEA normalizes the covariance matrix of each subject close to the identity matrix via the standard EA transformation, but incorporates a hard gating mechanism based on the log-determinant of the covariance matrix: if a target subject's $\log|\det(\bar{R})|$ falls below the population threshold $\mu - \sigma$, EA is automatically disabled for that fold to prevent noise amplification. This serves as the framework's adaptive first physical defense line.

2) **Spatiotemporal Feature Extractor:** After mitigating physical drifts, it is necessary to extract high-fidelity discriminative representations from the EEG signals. MAC-DA adopts ATCNet as the backbone network, utilizing convolution blocks to capture spatial-spectral patterns, multi-head attention to model global temporal dependencies, and a multi-scale temporal convolutional network (TCN) to extract multi-granular dynamic features. These three components synergistically output a sample-level feature vector of dimension $d_{model} = 32$.
3) **Global Semantic Adversarial Alignment:** Physical manifold normalization alone cannot eliminate distribution shifts in the high-dimensional semantic space. Based on the Conditional Domain Adversarial Network (CDAN), MAC-DA jointly encodes features and classification predictions via a multilinear mapping. The encoded representations are then fed into a domain discriminator through a Gradient Reversal Layer (GRL) to implicitly align global marginal distributions using an entropy-aware weighted adversarial loss. This resolves the deep feature distribution mismatch caused by variations in imagery strategies.
4) **Local Semantic Anchoring:** Pure adversarial training on low-SNR EEG signals is prone to optimization instability and high cross-initialization variance. MAC-DA introduces class-conditional CORAL (ccCORAL) to group features by class using high-confidence pseudo-labels, explicitly aligning the second-order statistics of each class. Serving as a rigid local constraint, ccCORAL provides a deterministic optimization direction for adversarial gradients, achieving a balance between accuracy and training stability.

The overall optimization objective of MAC-DA consists of the weighted sum of the source domain classification loss, the entropy-aware domain adversarial loss, the target domain entropy minimization loss, and the local

ccCORAL loss. The exact formulations are detailed in Section 3.7.

3.3. Preprocessing and Entropy-Aware Euclidean Alignment in Physical Manifold Space

The distribution shift of EEG signals primarily originates from inherent physical and physiological differences among subjects. These differences result in severe macroscopic baseline drifts on the underlying physical manifold of EEG covariance matrices. To directly eliminate such sensor-level shifts at the data input end, MAC-DA employs a preprocessing pipeline integrating Z-score standardization and Entropy-Aware Euclidean Alignment (EAEA).

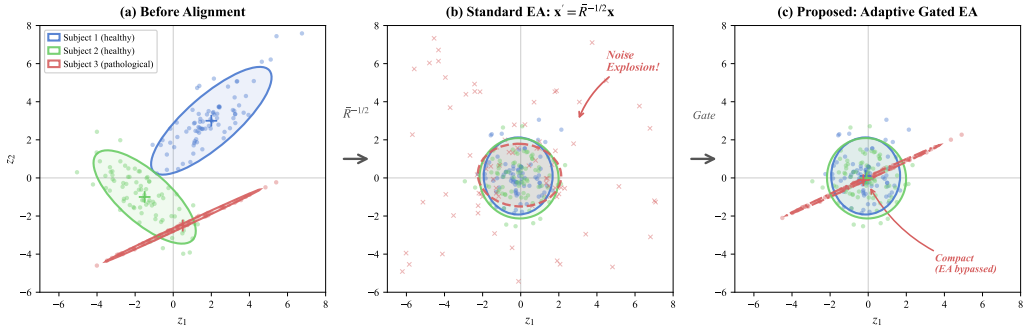

Figure 2: Geometric illustration of the proposed Entropy-Aware Euclidean Alignment (EAEA). (a) Raw features exhibit severe cross-subject domain shift, with Subject 3 (red) presenting a pathological, collapsed manifold (low $\log|\det|$). (b) Standard EA successfully aligns normal subjects (blue, green) to identity matrices, but catastrophically amplifies the noise in the pathological subject by forcing isotropic scaling. (c) The proposed hard-gated EAEA dynamically bypasses the whitening process for low-entropy subjects, effectively aligning normal distributions while preserving the compact topological structure of pathological manifolds to prevent negative transfer.

For each subject s, given a reference trial set $\mathbf{X}_{ref} = \{x_1, x_2, \ldots, x_N\}$ (where $x_i \in \mathbb{R}^{C\times T}$), the covariance matrix for each trial is first computed as:

$$R_i = \frac{x_i \cdot x_i^\top}{T} \tag{1}$$

Then, the mean covariance matrix across all trials for the subject is calculated:

$$\bar{R} = \frac{1}{N}\sum_{i=1}^{N} R_i \tag{2}$$

To ensure numerical stability, $\bar{R}$ is symmetrized and regularized: $\bar{R} \leftarrow \frac{1}{2}(\bar{R} + \bar{R}^\top) + \varepsilon \cdot I$, where ε is a small positive constant. Subsequently, the inverse square root transformation matrix is computed via eigenvalue decomposition:

$$A = \bar{R}^{-1/2} \tag{3}$$

Finally, a linear transformation is applied to all trials of this subject: $x_i' = A \cdot x_i$. This transformation normalizes the covariance matrix of each subject toward the identity matrix, eliminating macroscopic marginal distribution discrepancies in the Euclidean space.

Under the LOSO scenario, the implementation of EA strictly follows a no-leakage principle: each source subject independently fits the transformation matrix A using their own training session data, which is then applied to all data of that subject; the target subject solely uses its unlabeled training session data to fit the transformation matrix. Following EA alignment, all data undergoes channel-wise Z-score standardization, where the standardizer parameters are fitted exclusively on the source domain data and uniformly applied to all data.

3.3.1. Hard Gating via Log-Determinant of Covariance Matrix

While EA is effective for the majority of subjects, our empirical analysis reveals that subjects with atypical or collapsed covariance structures experience negative transfer when EA is applied (see Section 4.5.2). The inverse square root transformation $\bar{R}^{-1/2}$ amplifies noise for such subjects rather than normalizing it. To address this, we propose Entropy-Aware Euclidean Alignment (EAEA), which extends EA with an automatic hard gating mechanism.

The gating criterion is based on the log-determinant of the mean covariance matrix, which serves as a scalar summary of the "information content" of a subject's covariance manifold. For each subject s, we compute:

$$\delta_s = \log |\det(\bar{R}_s)| \tag{4}$$

where $\bar{R}_s$ is the regularized mean covariance matrix. The population statistics are then computed across all S subjects:

$$\mu_\delta = \frac{1}{S}\sum_{s=1}^{S}\delta_s, \quad \sigma_\delta = \sqrt{\frac{1}{S}\sum_{s=1}^{S}(\delta_s - \mu_\delta)^2} \tag{5}$$

The gating decision for the target subject t follows the 1-σ rule:

$$\text{gate}(t) = \begin{cases} \text{disable EA} & \text{if } \delta_t < \mu_\delta - \sigma_\delta \\ \text{apply EA} & \text{otherwise} \end{cases} \tag{6}$$

When EA is disabled for a gated subject, the framework falls back to the non-EA pipeline, ensuring that the subject still benefits from deep domain adaptation without suffering from noise amplification in the physical manifold space. This gating decision is made once per LOSO fold before training begins, introducing negligible computational overhead.

3.4. Spatiotemporal Feature Extraction Based on Hybrid Convolution-Attention Architecture

After eliminating physical manifold drifts, MAC-DA employs ATCNet as the backbone to extract high-fidelity spatiotemporal representations. As shown in Fig. 3, ATCNet consists of a cascaded structure comprising a convolution block, a sliding-window multi-head attention module, and a multi-scale temporal convolutional network (TCN).

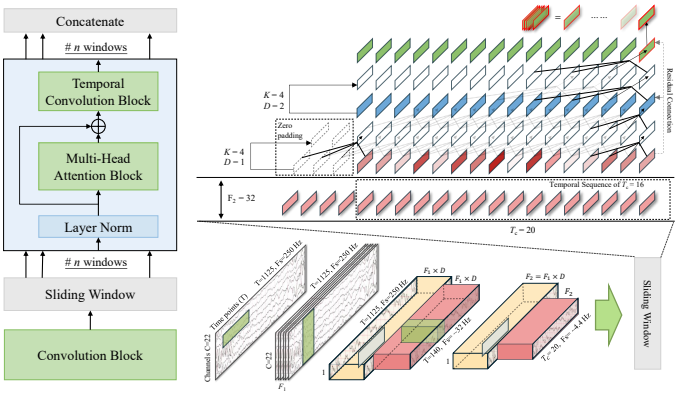

Figure 3: Detailed architecture of the hybrid convolution-attention backbone (ATCNet). The network integrates: (a) a Temporal Convolution Block for spectral-spatial feature mapping; (b) a Sliding Window module with Multi-Head Attention to capture global temporal dependencies across overlapping epochs; and (c) a multi-scale Temporal Convolutional Network (TCN) with residual connections for aggregating deep dynamic representations of motor intentions.

3.4.1. Convolution Block

The convolution block is responsible for extracting spatial-spectral features from raw EEG signals. The input $x \in \mathbb{R}^{C \times T}$ is first reshaped to $\mathbb{R}^{1 \times C \times T}$ and sequentially passes through a temporal convolutional layer (F_1 filters to capture spectral features), a depthwise convolutional layer (learning spatial filters via grouped convolution), and a separable convolutional layer. Batch normalization, ELU activation, average pooling, and Dropout are interleaved between layers. After processing, the output feature tensor dimension becomes $\mathbb{R}^{B \times F_2 \times T'}$, where B represents the batch size, $F_2 = F_1 \times D$ is the number of channels (D is the depth multiplier), and T' is the temporal length downsampled by two pooling operations. It is subsequently reshaped into a time-series format $\mathbb{R}^{B \times T' \times F_2}$ for subsequent processing.

3.4.2. Sliding-Window Multi-Head Attention

To capture global temporal dependencies during motor imagery, the time series is partitioned into $n_{windows}$ overlapping sub-sequences. Each window is independently fed into a multi-head self-attention module, calculating temporal dependencies via the Scaled Dot-Product Attention mechanism, with residual connections preserving the original feature information.

3.4.3. Multi-Scale TCN and Feature Aggregation

The outputs from the attention module are reshaped and fed into the TCN. The TCN is composed of a cascade of *depth* blocks, each containing causal convolutions with exponentially increasing dilation rates ($d = 1, 2, \ldots$). This expands the receptive field exponentially without increasing the number of parameters. Finally, the feature vectors from the last time step of each sliding window are arithmetically averaged to yield the final sample-level feature representation $\mathbf{f} \in \mathbb{R}^{d_{model}}$.

3.5. Entropy-Aware Conditional Domain Adversarial Alignment in Global Semantic Space

To resolve distribution shifts in the high-dimensional semantic space caused by diverse imagery strategies, MAC-DA employs an entropy-conditioned variant of the conditional domain adversarial network (CDAN+E). This module matches the joint distributions across domains via implicit adversarial learning.

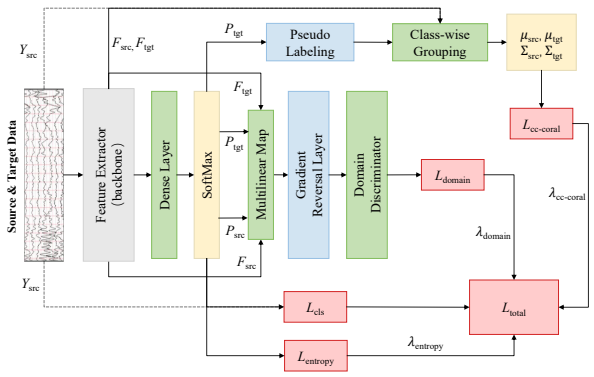

Figure 4: The joint optimization framework of MAC-DA. Source and target domain features are concatenated with classifier predictions via multilinear mapping to guide the domain discriminator. The framework stabilizes the adversarial game by integrating three auxiliary constraints: entropy minimization ($L_{entropy}$) for target certainty, entropy-aware weighting for adversarial importance, and class-conditional CORAL ($L_{cc-coral}$) as a localized manifold anchor to suppress training variance.

3.5.1. Multilinear Mapping and Domain Discrimination

For the feature $\mathbf{f} \in \mathbb{R}^{d_{model}}$ output by the feature extractor and the softmax probability $\mathbf{p} \in \mathbb{R}^{K}$ predicted by the classifier, a joint conditional feature $\mathbf{h} = \mathbf{f} \otimes \mathbf{p} \in \mathbb{R}^{d_{model} \times K}$ is calculated via multilinear mapping. This feature is passed through a gradient reversal layer (GRL) and fed into the domain discriminator. The domain discriminator adopts a two-layer fully connected network, and its output scalar is mapped to the $(0, 1)$ interval via a Sigmoid activation function, representing the probability of the sample belonging to the target domain.

3.5.2. Progressive Adversarial Training and Entropy-Aware Weighting

To ensure training stability, the gradient reversal coefficient λ of the GRL follows a progressive sigmoid-based scheduling strategy. Concurrently, an entropy-conditioned weighting mechanism is introduced to adjust the weight of each sample in the adversarial loss: $w = 1 + \exp(-H(\mathbf{p}))$, where $H(\mathbf{p})$ is the prediction entropy. Samples with more certain predictions receive higher

weights. The weighted and batch-normalized domain adversarial loss is:

$$\mathcal{L}_{domain} = \frac{1}{N_s}\sum_{i=1}^{N_s} w_i^s \cdot \mathrm{BCE}(D(\mathbf{h}_i^s), 0) + \frac{1}{N_t}\sum_{j=1}^{N_t} w_j^t \cdot \mathrm{BCE}(D(\mathbf{h}_j^t), 1) \quad (7)$$

Furthermore, an entropy minimization regularization term $\mathcal{L}_{entropy} = \frac{1}{N_t}\sum_{j=1}^{N_t} H(\mathbf{p}_j^t)$ is imposed on the target domain predictions to encourage the model to output more confident predictions.

3.6. Class-Conditional Second-Order Statistic Anchoring in Local Semantic Space

To overcome the high variance issues inherent in pure adversarial training, MAC-DA introduces class-conditional CORAL (ccCORAL) as a stabilizing manifold anchor in the local semantic space.

3.6.1. Pseudo-Label Generation and Conditional Statistics Calculation

Target domain samples are assigned pseudo-labels $\hat{y}_j$ based on the classifier's predicted probabilities, retaining only those samples with confidence exceeding the threshold τ. Based on the ground-truth labels (source) and high-confidence pseudo-labels (target), features from both domains are grouped by class k to compute the intra-class means (μ_k^s, μ_k^t) and covariance matrices (Σ_k^s, Σ_k^t), respectively.

3.6.2. Second-Order Statistic Anchoring and Warmup Mechanism

The ccCORAL loss for class k consists of a covariance alignment term and a mean alignment term:

$$l_k = \|\Sigma_k^s - \Sigma_k^t\|_F^2 + \alpha \cdot \|\mu_k^s - \mu_k^t\|_2^2 \quad (8)$$

where α controls the relative weight of the first-order mean alignment. (For simplicity of expression, the standard CORAL feature dimension scaling factor is absorbed into the final hyperparameter $\lambda_{ccCORAL}$). The final $\mathcal{L}_{ccCORAL}$ is the average of l_k over all classes that meet the minimum sample requirement (n_{min}). To prevent negative transfer caused by noisy pseudo-labels in the early training stages, ccCORAL employs a warmup strategy, setting its weight to 0 for the first T_{warm} epochs.

3.7. Overall Optimization Objective and Hyperparameter Configuration

The total training loss of MAC-DA is a weighted combination of the aforementioned losses:

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + \lambda_{domain} \cdot \mathcal{L}_{domain} + \lambda_{entropy} \cdot \mathcal{L}_{entropy} + \lambda_{ccCORAL} \cdot \mathcal{L}_{ccCORAL} \quad (9)$$

where $\mathcal{L}_{cls}$ is the standard cross-entropy classification loss on the labeled source domain data, defined as:

$$\mathcal{L}_{cls} = -\frac{1}{N_s} \sum_{i=1}^{N_s} \sum_{k=1}^{K} \mathbb{I}(y_i^s = k) \log p_k(x_i^s) \quad (10)$$

Table 1 summarizes the key hyperparameter configurations of the MAC-DA framework. The architectural parameters of ATCNet (e.g., $F_1 = 16, K = 4$) are determined based on the primary spectral characteristics (8-26 Hz) and temporal dynamics of MI-EEG to ensure the receptive field adequately covers critical neural patterns. In the domain adaptation module, the adversarial weight λ_{domain} progressively increases to 1.0, while the ccCORAL weight $\lambda_{ccCORAL}$ is set to a small value (0.01) coupled with the warmup mechanism, avoiding over-constraint on the feature space during the noisy pseudo-label phase. Additionally, a high pseudo-label threshold ($\tau = 0.9$) and a minimum sample requirement ($n_{min} = 4$) are enforced to fundamentally guarantee the robustness of the second-order statistics computation.

4. Experiments and Results

To verify the effectiveness of the proposed MAC-DA framework for cross-subject MI-EEG classification, we conducted comprehensive experiments on two public MI-EEG benchmarks: the BCI Competition IV Dataset 2a and the Weibo2014 dataset. In this section, we describe the datasets and preprocessing in Section 4.1, introduce the experimental settings and comparison methods in Section 4.2, present the main classification results in Section 4.3, conduct ablation studies in Section 4.4, and provide further analysis and discussion in Section 4.5.

4.1. MI-EEG Dataset and Preprocessing

We evaluated the proposed MAC-DA on two public MI-EEG datasets to validate cross-dataset generalizability.

Table 1: Key Hyperparameter Configurations of MAC-DA

Category	Symbol	Description	Value
ATCNet Architecture	F_1	Number of temporal filters	16
	D	Depth multiplier	2
	d_{model}	Sample-level feature dimension (F_2)	32
	n_{head}	Number of attention heads	2
	key_dim	Attention key/query dimension	8
	$depth$	TCN cascade depth	2
	K	TCN causal kernel size	4
	$n_{windows}$	Number of sliding windows	5
Domain Adaptation Constraints	λ_{domain}	Maximum domain adversarial loss weight	1.0
	$\lambda_{entropy}$	Entropy minimization loss weight	0.1
	$\lambda_{ccCORAL}$	ccCORAL loss weight	0.01
	α	Relative weight of ccCORAL mean alignment	0.1
	τ	Target domain pseudo-label confidence threshold	0.9
	n_{min}	Minimum samples per class to compute statistics	4
	T_{warm}	Warmup epochs for ccCORAL module	5
EAEA Gating	—	Gating rule: $\delta_t < \mu_\delta - \sigma_\delta$	1-σ
Optimization and Training	lr	Initial learning rate (Adam)	0.0009
	β_1, β_2	Adam optimizer momentum parameters	0.5, 0.999
	$weight\_decay$	L2 regularization decay coefficient	0.001
	max_epochs	Maximum training epochs (LOSO protocol)	125

4.1.1. BCI Competition IV Dataset 2a

The BCI Competition IV Dataset 2a (hereinafter referred to as "2a") [66] contains EEG recordings from 9 subjects, each performing four classes of motor imagery tasks: left hand, right hand, feet, and tongue. Each subject completed two sessions (training and evaluation) on different days, with each session containing 288 trials (72 trials per class). The EEG signals were recorded using 22 Ag/AgCl electrodes at a sampling rate of 250 Hz.

4.1.2. Weibo2014 Dataset

The Weibo2014 dataset [71] contains EEG recordings from 10 subjects performing seven classes of motor imagery tasks. Following established practice, we selected four classes—left hand, right hand, hands (both hands), and feet—yielding 320 trials per subject (80 per class). The original EEG signals were recorded using a 64-channel Neuroscan system (60 EEG + 4 EOG channels) at 200 Hz. To ensure architectural compatibility and fair comparison, we applied the following preprocessing adaptations: (1) EOG channels were removed and 22 EEG channels corresponding to the sensorimotor montage of the 2a dataset were selected (Fz, FC3, FC1, FCz, FC2, FC4, C5, C3, C1, Cz, C2, C4, C6, CP3, CP1, CPz, CP2, CP4, P1, Pz, P2, POz); (2) signals were resampled from 200 Hz to 250 Hz; and (3) since Weibo2014 contains only a single session per subject, the session-based train/test split used for

the 2a dataset was replaced with a stratified 70%/30% split (224 training / 96 test trials per subject).

It is important to note that the fourth class differs between the two datasets: the 2a dataset uses *tongue* imagery, while Weibo2014 uses *hands* (both hands) imagery. These tasks engage distinct spatial activation patterns (midline sensorimotor cortex for tongue vs. bilateral C3/C4 for both-hands). Notably, this four-class cross-subject MI task on Weibo2014 is exceptionally challenging: a recent large-scale BCI benchmark study reported that the best-performing pipeline among 30 evaluated methods achieves only 43.3% accuracy on this dataset [72, 73], barely above the 25% chance level. Consequently, absolute accuracy values should not be compared directly across datasets; instead, we focus on the consistency of relative performance trends.

Table 2 summarizes the key properties of both datasets.

Table 2: Summary of the two MI-EEG benchmark datasets used for evaluation.

Property	BCI Competition IV 2a	Weibo2014
Subjects	9	10
Classes	4 (LH, RH, F, T)	4 (LH, RH, BH, F)
Channels (used)	22	22 (of 64)
Sampling rate	250 Hz	200 $\rightarrow$ 250 Hz
Sessions	2 (T + E)	1
Train/Test split	Session T / Session E	Stratified 70%/30%
Trials per subject	288 (72×4)	320 (80×4)

4.1.3. Preprocessing

For both datasets, we considered the time interval of [0 s, 4 s] after the cue onset in each trial. The raw EEG data were bandpass filtered using a causal third-order Butterworth filter with a passband of 4–38 Hz. For methods involving Euclidean Alignment (EA), the alignment was performed following the no-leakage principle described in Section 3.3: each source subject's transformation matrix was fitted on its training session data, and the target subject's transformation matrix was fitted on its training session data (unlabeled), ensuring that test data never participated in the computation of alignment parameters. After EA, all data were further standardized using per-channel Z-score normalization, with the standardizer fitted on source domain data.

4.2. Experimental Settings

4.2.1. Transfer Protocol

We adopted the Leave-One-Subject-Out (LOSO) cross-validation protocol on both datasets. For the 2a dataset, in each fold, one of the 9 subjects was selected as the target domain, and the remaining 8 subjects' training session data constituted the source domain training set. The target subject's training session data served as unlabeled target domain data for domain adaptation, and the target subject's evaluation session data served as the final test set. For the Weibo2014 dataset, the same LOSO protocol was applied across 10 subjects, with the modification that each subject's single-session data was split into 70% training and 30% test via stratified sampling (preserving class balance). The 70% portion of the target subject served as unlabeled adaptation data, and the 30% portion served as the test set.

4.2.2. Comparison Methods

We compared the proposed MAC-DA with ten methods spanning three categories:

Non-transfer deep learning methods (baselines):.

1. **EEGNet** [28]: A compact CNN employing depthwise and separable convolutions for EEG decoding.
2. **EEGTCNet** [67]: An extension of EEGNet incorporating temporal convolutional networks.
3. **TSSEFFNet** [68]: A temporal-spatial-spectral efficient feature fusion network.
4. **MSCFormer** [69]: A multi-scale convolutional Transformer for MI-EEG decoding.
5. **ATCNet** [35]: The backbone network adopted in MAC-DA, combining attention mechanisms with temporal convolutional networks. This serves as the primary baseline.

Domain adaptation methods:.

6. **DANN** [50]: Domain-Adversarial Neural Network with gradient reversal layer for feature-level adversarial alignment.
7. **DeepCORAL** [61]: Deep CORrelation ALignment that minimizes the discrepancy between source and target domain feature distributions by aligning their second-order statistics (covariance matrices), applied on the ATCNet backbone.

8. **CDAN** [51]: Conditional Domain Adversarial Network with multilinear conditioning and entropy weighting, applied on the ATCNet backbone without EA preprocessing.
9. **MI-DABAN** [53]: A dual-attention-based adversarial network that introduces domain adversarial batch normalization with mutual information regularization for motor imagery classification, applied on the ATCNet backbone.

For fair comparison, all deep learning methods used the same data preprocessing pipeline. To rigorously mitigate the influence of random initialization on experimental conclusions, all methods were trained with 5 consecutive random seeds (0–4), chosen a priori without any post-hoc selection or cherry-picking. This multi-seed protocol is deliberately more stringent than the widespread but methodologically questionable practice of reporting only the single best seed—a practice that risks overstating generalization performance by conflating seed-level luck with genuine algorithmic superiority. By reporting both the peak performance (best seed) and the expected performance with uncertainty quantification (mean ± std across seeds), our evaluation provides a more honest and reproducible assessment of each method. The hyperparameters of MAC-DA are specified in Table 1 of Section 3.7. All experiments were implemented in Python using the PyTorch framework on a workstation equipped with an NVIDIA GPU 4060.

4.2.3. Evaluation Metrics

For the main MI decoding task, classification accuracy and Cohen's Kappa coefficient were used as the primary evaluation metrics. Statistical significance was assessed using paired t-tests on seed-level mean accuracies.

For the subject-level EA gating analysis, we additionally report precision, recall, and F1-score for the binary decision "EA harmful" (positive class) vs. "EA safe" (negative class). Let TP, FP, FN denote true positives, false positives, and false negatives. The metrics are defined as:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN}, \quad \text{F1} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{11}$$

We also report balanced accuracy to reduce class-imbalance bias in subject-level gate evaluation.

4.3. Main Classification Results

Table 3 presents the overall classification accuracies of all comparison methods under the LOSO protocol on the BCI Competition IV 2a dataset. For each method, we report both the best single-seed accuracy and the mean accuracy across 5 seeds. The best single-seed result reflects the peak achievable performance, while the 5-seed mean with standard deviation provides a more rigorous and reproducible assessment that accounts for the stochasticity of neural network training. Per-subject results are provided in Table 4. The best results are highlighted in boldface.

Table 3: Overall classification accuracies (%) of the LOSO protocol on the BCI Competition IV Dataset 2a. "Best Seed" reports the highest accuracy among 5 seeds; "Mean ± Std" reports the average across 5 seeds. The best result in each column is in **bold**.

Method	Best Seed	Mean ± Std	Kappa
Non-transfer baselines			
EEGNet	52.55	52.09 ± 12.11	36.12 ± 16.14
EEGTCNet	57.02	55.62 ± 13.02	40.82 ± 17.35
TSSEFFNet	58.56	58.19 ± 11.24	44.26 ± 14.98
MSCFormer	58.06	54.68 ± 17.15	39.58 ± 22.86
ATCNet	62.11	60.29 ± 13.86	47.06 ± 18.48
Domain adaptation methods			
DANN	68.67	65.94 ± 14.32	54.59 ± 19.09
DeepCORAL	68.48	66.09 ± 14.05	54.78 ± 18.73
CDAN	68.36	66.16 ± 14.48	54.88 ± 19.31
MI-DABAN	67.21	65.07 ± 15.34	53.43 ± 20.46
Proposed			
MAC-DA	**72.18**	**69.77 ± 15.44**	**59.69 ± 20.58**

From Table 3, we observe the following:

1) **Overall superiority of MAC-DA:** The proposed MAC-DA achieves the highest accuracy under both evaluation criteria: 72.18% for the best seed and 69.77% for the 5-seed mean. Compared with the ATCNet baseline, MAC-DA improves the best-seed accuracy by 10.07 pp and the mean accuracy by 9.48 pp. Compared with the strongest comparison domain adaptation method, MAC-DA improves by 3.51 pp in best-seed accuracy (vs. DANN) and 3.61 pp in mean accuracy (vs. CDAN). The

Table 4: Per-subject classification accuracies (%, mean ± std across 5 seeds) on the BCI Competition IV Dataset 2a. The best result for each subject is in **bold**.

(a) Non-transfer baselines

Subject	EEGNet	EEGTCNet	TSSEFFNet	MSCFormer	ATCNet
S1	60.5±1.4	62.1±4.7	66.1±1.7	68.5±1.5	71.5±1.0
S2	41.8±2.6	45.1±1.5	47.6±1.7	41.7±1.6	45.6±2.4
S3	74.5±3.7	80.6±2.0	74.0±1.0	84.2±2.9	80.8±4.0
S4	45.2±3.3	51.9±4.6	49.5±2.4	43.5±2.4	55.6±2.6
S5	39.2±2.2	38.3±5.3	48.3±0.6	32.5±12.1	43.5±9.9
S6	40.6±1.1	43.1±2.9	42.4±1.0	38.8±2.1	43.3±2.5
S7	59.6±4.1	65.8±0.8	72.3±0.8	62.4±5.4	70.1±2.1
S8	62.4±3.3	62.8±2.1	66.1±3.0	70.4±3.5	71.3±3.4
S9	45.1±4.1	50.8±2.5	57.4±1.7	50.3±3.3	61.0±1.3

(b) DA methods

	DA methods				*Proposed*
Subject	DANN	DeepCORAL	CDAN	MI-DABAN	MAC-DA
S1	72.8±1.8	72.9±2.2	71.0±0.9	71.3±2.5	**81.9±1.7**
S2	48.8±1.9	51.2±1.9	**51.8±2.1**	50.1±3.0	52.0±2.4
S3	87.2±1.3	87.4±1.9	85.6±0.3	87.0±0.8	**90.5±2.1**
S4	65.8±3.2	62.5±2.8	66.0±2.4	64.5±2.4	**68.1±1.6**
S5	47.6±16.7	**49.7±16.2**	43.3±17.5	38.3±12.7	44.8±15.9
S6	55.0±2.4	53.9±1.6	55.8±1.9	55.1±3.0	**60.0±2.7**
S7	77.2±1.7	76.7±4.2	78.3±2.1	78.1±1.7	**83.3±2.9**
S8	77.7±1.4	**79.3±2.4**	78.7±3.1	78.3±1.6	74.4±4.5
S9	61.3±1.4	61.2±1.3	64.9±2.6	63.0±3.4	**72.9±3.5**

consistency between best-seed and mean improvements indicates that MAC-DA's advantage is robust and not an artifact of seed selection.

2) **Domain adaptation methods outperform baselines:** All domain adaptation methods (DANN, CDAN, MAC-DA) substantially outperform the non-transfer baselines. Compared with the best non-transfer baseline ATCNet (60.29%), DANN and CDAN improve mean accuracy by approximately 5.5–6 pp, while MAC-DA achieves a further improvement of approximately 3.6 pp over these methods. Notably, the gap between best-seed and mean accuracy is larger for some baselines (e.g., MSCFormer: 58.06 vs. 54.68, a 3.38 pp gap) than for MAC-DA (72.18 vs. 69.77, a 2.41 pp gap), suggesting that domain adaptation also improves training stability.
3) **Subject-specific analysis and EAEA gating:** MAC-DA achieves the best or competitive performance on 8 out of 9 subjects (S1, S2, S3, S4, S6, S7, S9; S8 is the only exception), with particularly notable improvements on S1 (+10.49 pp over ATCNet), S6 (+16.74 pp), and S9 (+11.88 pp). Crucially, the EAEA gating mechanism automatically identifies subjects S2 and S5 as having atypical covariance structures ($\delta < \mu - \sigma$) and disables EA for these folds, falling back to the C3DAN pipeline. This converts what would have been EA-induced negative transfer (-9.17 pp and -5.76 pp respectively) into preserved or improved performance, demonstrating the practical value of the gating mechanism (see Section 4.5.2).

4.3.1. Statistical Significance

To quantify the improvement of MAC-DA, we conducted paired t-tests on seed-level mean accuracies (5 seeds, $df = 4$). Table 5 presents the results.

Table 5: Paired t-test results of MAC-DA vs. comparison methods ($n = 5$ seeds, $df = 4$). p-values less than 0.05 are in **bold**.

Comparison	Δ	t	p
vs. ATCNet	+9.48	14.53	**<0.001**
vs. MI-DABAN	+4.70	5.02	**0.004**
vs. DeepCORAL	+3.68	5.42	**0.003**
vs. CDAN	+3.61	5.87	**0.002**
vs. DANN	+3.83	5.16	**0.003**

The improvement of MAC-DA over all comparison methods is statisti-

cally significant on the full 9-subject dataset. Compared with the ATCNet baseline, the improvement is highly significant ($p < 0.001$). Critically, thanks to the EAEA gating mechanism that prevents negative transfer on atypical subjects, the improvements over all competing domain adaptation methods now reach significance: vs. CDAN ($p = 0.002$), vs. DANN ($p = 0.003$), and vs. DeepCORAL ($p = 0.003$). This demonstrates that the EAEA gating mechanism successfully resolves the statistical power issue that plagued the non-gated version, where EA-induced negative transfer on S2 and S5 inflated variance and masked the true effect size.

4.3.2. Cross-Dataset Validation on Weibo2014

To validate the generalizability of MAC-DA beyond a single benchmark, we evaluated all domain adaptation methods and the ATCNet baseline on the Weibo2014 dataset under the same LOSO protocol. Table 6 presents the overall results, Table 7 provides per-subject accuracies, and Table 8 reports the statistical significance tests.

Table 6: Overall classification accuracies (%) of the LOSO protocol on the Weibo2014 dataset. "Best Seed" reports the highest accuracy among 5 seeds; "Mean ± Std" reports the average across 5 seeds. The best result in each column is in **bold**.

Method	Best Seed	Mean ± Std	Kappa
ATCNet	47.08	45.91 ± 12.12	27.89 ± 16.15
DANN	50.31	49.00 ± 10.37	32.00 ± 13.82
DeepCORAL	49.11	47.21 ± 10.99	29.61 ± 14.65
CDAN	52.44	50.76 ± 12.02	34.35 ± 16.03
MAC-DA	**55.16**	**53.70 ± 10.60**	**38.27 ± 14.13**

The results on the Weibo2014 dataset confirm several key findings:

1) **Consistent superiority across datasets:** MAC-DA achieves the highest mean accuracy (53.70%) among all methods, significantly outperforming the ATCNet baseline (+7.30 pp, $p = 0.002$) and all competing DA methods ($p < 0.05$). Notably, this accuracy substantially exceeds the 43.3% reported by the MOABB benchmark [72] for the best classical pipeline on the same dataset, demonstrating the value of the proposed multi-space alignment strategy.
2) **Consistent relative ranking:** The performance hierarchy observed on the 2a dataset is preserved on Weibo2014: MAC-DA > CDAN >

Table 7: Per-subject classification accuracies (%, mean ± std across 5 seeds) on the Weibo2014 dataset. The best result for each subject is in **bold**.

	Baseline	*DA methods*			*Proposed*
Subject	ATCNet	DANN	DeepCORAL	CDAN	MAC-DA
S1	27.3±3.9	**37.7±2.9**	31.5±2.0	36.7±2.1	36.5±2.9
S2	48.8±3.3	49.4±2.7	47.5±1.4	**51.3±2.1**	49.2±3.1
S3	35.2±1.5	39.4±1.8	**40.6±2.1**	33.5±4.3	39.0±0.5
S4	49.0±2.2	46.9±2.3	43.5±2.7	50.4±3.0	**53.1±2.4**
S5	36.3±1.5	39.6±0.9	37.3±3.3	43.1±1.4	**57.7±3.2**
S6	64.8±5.0	66.4±1.6	65.2±6.7	**69.1±3.1**	64.3±7.2
S7	56.7±2.8	61.7±2.8	61.0±3.6	60.2±4.9	**65.0±3.3**
S8	49.2±6.3	51.3±4.5	48.5±2.8	60.2±5.9	**64.8±2.1**
S9	58.1±2.8	58.3±7.0	57.3±5.2	**62.3±5.3**	60.4±2.4
S10	34.0±2.7	39.4±2.9	39.6±2.4	40.8±3.3	**47.1±3.5**

Table 8: Paired t-test results of MAC-DA vs. comparison methods on the Weibo2014 dataset ($n = 5$ seeds, $df = 4$). p-values less than 0.05 are in **bold**.

Comparison	Δ (pp)	t	p
vs. ATCNet	+7.30	7.17	**0.002**
vs. DeepCORAL	+6.49	10.36	**<0.001**
vs. DANN	+4.70	4.02	**0.016**
vs. CDAN	+2.94	2.97	**0.041**

DANN > DeepCORAL > ATCNet. This consistent ranking across two independent datasets with different recording systems, channel counts, and task compositions strongly supports the generalizability of MAC-DA's multi-space progressive alignment philosophy.

3) **Larger improvements on difficult subjects:** MAC-DA shows particularly strong improvements on subjects where the baseline performs poorly. For instance, S5 improves from 36.3% (ATCNet) to 57.7% (+21.4 pp), and S10 improves from 34.0% to 47.1% (+13.1 pp). This suggests that the multi-space alignment is especially beneficial when the distribution shift between source and target domains is severe.

4.4. Ablation Study

To evaluate the contribution of each component in MAC-DA, we conducted a systematic ablation study. The MAC-DA framework consists of three key components beyond the ATCNet backbone: Euclidean Alignment (EA), Conditional Domain Adversarial Network (CDAN+E), and class-conditional CORAL (ccCORAL). We define five ablative configurations as shown in Table 9.

Table 9: Definition of ablative configurations. MAC-DA uses EAEA (EA with automatic gating), while EA-CDAN uses plain EA (always on).

Configuration	CDAN+E	ccCORAL	EA	EAEA Gating
EA-ATCNet			✓	
CDAN	✓			
C3DAN	✓	✓		
EA-CDAN	✓		✓	
MAC-DA (full)	✓	✓	✓	✓

4.4.1. Ablation Results

Table 10 presents the overall classification results of the ablation study.

From the ablation results, we analyze the role of each component in addressing different aspects of the cross-subject distribution shift:

1) **CDAN+E for global feature-level alignment:** The cross-subject MI-EEG decoding task suffers from significant distribution shift caused by individual differences in imagery strategies and neural activation intensities. The conditional adversarial network (CDAN+E) addresses

Table 10: Ablation results on the BCI Competition IV 2a dataset (mean ± std across all subjects and 5 seeds). MAC-DA includes EAEA gating.

Method	Acc (%)	Kappa (%)
ATCNet (backbone)	60.29 ± 13.86	47.06 ± 18.48
EA-ATCNet (EA-only)	64.42 ± 12.52	52.56 ± 16.69
CDAN	66.16 ± 14.48	54.88 ± 19.31
C3DAN	66.51 ± 13.89	55.35 ± 18.52
EA-CDAN	67.92 ± 17.63	57.23 ± 23.51
MAC-DA (full)	**69.77 ± 15.44**	**59.69 ± 20.58**

this global feature-level shift by conditioning the domain discriminator on classifier predictions, thereby capturing the multimodal structure of the feature space and aligning class-conditional distributions more effectively than unconditional adversarial methods. As shown in Table 10, introducing CDAN+E to the ATCNet backbone yields a substantial +5.87 pp accuracy improvement (from 60.29% to 66.16%), confirming that adversarial feature alignment is essential for bridging the intersubject distribution gap.

2) **EA for manifold-level geometric normalization:** Beyond the feature-level shift, cross-subject EEG data exhibit geometric differences in covariance structures arising from individual variations in brain anatomy and electrode placement. Euclidean Alignment addresses this manifold-level shift by projecting each subject's data onto a common reference frame via whitening with respect to the session-level covariance matrix, providing a geometrically consistent input space for subsequent deep learning. The EA-only baseline already yields a clear gain over the plain backbone (EA-ATCNet vs. ATCNet: +4.13 pp), confirming that geometric normalization is independently beneficial. When combined with adversarial alignment, EA still provides additive gains (EA-CDAN vs. CDAN: +1.77 pp), demonstrating complementarity between manifold normalization and deep semantic alignment.

3) **ccCORAL for class-conditional distribution anchoring:** While CDAN+E aligns the global feature distributions through adversarial training, it does not explicitly constrain the class-conditional second-order statistics, leaving the adversarial optimization susceptible to stochastic fluctuations across random seeds. ccCORAL addresses this gap by explicitly anchoring the class-conditional covariance matrices between

domains, providing a deterministic optimization signal that complements the adversarial gradients. The isolated accuracy improvement of ccCORAL is modest: C3DAN vs. CDAN yields +0.35 pp, and the non-gated MAC-DA vs. EA-CDAN yields +0.32 pp. However, ccCORAL's primary contribution lies not in absolute accuracy but in training stability: it significantly suppresses cross-seed variance ($p = 0.0214$; analyzed in detail in Section 4.5.1), confirming its designed role as a training stabilizer that regularizes the adversarial optimization landscape.

4) **EAEA gating and complementary integration in MAC-DA:** Beyond the three deep components, EAEA gating provides an additional +1.53 pp by converting EA-induced negative transfer on atypical subjects (S2, S5) into preserved C3DAN-level performance. The four mechanisms target different levels of the distribution shift hierarchy: EAEA adaptively normalizes the input-space geometry, CDAN+E aligns the global deep feature distributions, and ccCORAL anchors the class-conditional statistics. Their combined effect in MAC-DA yields the highest overall accuracy (69.77%) and Kappa (59.69%), with a total improvement of +3.61 pp over CDAN alone. The orthogonality of their contributions—EA's improvement remains stable with or without ccCORAL, and vice versa—confirms that these components address genuinely distinct aspects of the cross-subject problem rather than redundantly targeting the same source of variation.

4.4.2. Statistical Significance of Ablation

Table 11 presents the paired t-test results between MAC-DA (with EAEA gating) and each ablative version.

Table 11: Paired t-test between MAC-DA (with EAEA gating) and ablative versions ($n = 5$ seeds). Significant results ($p < 0.05$) are in **bold**.

Comparison (Acc)	Δ (pp)	t	p (one-sided)
MAC-DA vs. EA-ATCNet	+5.35	8.90	**<0.001**
MAC-DA vs. CDAN	+3.61	5.87	**0.0020**
MAC-DA vs. C3DAN	+3.26	4.92	**0.0040**
MAC-DA vs. EA-CDAN	+1.85	3.47	**0.0127**

With EAEA gating, all ablation comparisons now reach statistical significance. The improvement over EA-ATCNet (+5.35 pp, $p < 0.001$) confirms

that MAC-DA's gain is not merely from EA preprocessing, but from the full multi-space adaptation design. The improvement over CDAN (+3.61 pp, $p = 0.002$) reflects the combined contribution of EA, ccCORAL, and EAEA gating. The improvement over C3DAN (+3.26 pp, $p = 0.004$) isolates the effect of EAEA (adaptive EA with gating). The improvement over EA-CDAN (+1.85 pp, $p = 0.013$) captures the joint effect of ccCORAL and EAEA gating, where gating converts negative transfer on S2 and S5 into preserved performance while ccCORAL provides consistent stabilization across all subjects. Additionally, the EA-only effect is statistically significant under the same protocol (EA-ATCNet vs. ATCNet: +4.13 pp, $t = 3.81$, $p = 0.0095$, one-sided), validating EA as a meaningful but incomplete component.

4.5. Analysis and Discussion

4.5.1. ccCORAL as a Variance Suppressor

A key contribution of ccCORAL is its ability to suppress the cross-seed variance of adversarial training. To quantify this effect, we compared the per-subject standard deviation of accuracy across 5 seeds between CDAN and C3DAN (CDAN + ccCORAL, without EA).

Table 12: Cross-seed standard deviation of accuracy (%) for CDAN vs. C3DAN. Negative Δ indicates variance suppression by ccCORAL.

Subject	CDAN Std	C3DAN Std	Δ (pp)
S1	0.92	0.78	−0.14
S2	2.12	2.36	+0.24
S3	0.28	0.26	−0.02
S4	2.38	1.93	−0.45
S5	17.48	15.91	−1.56
S6	1.88	1.15	−0.73
S7	2.14	2.20	+0.06
S8	3.13	2.57	−0.57
S9	2.63	1.71	−0.92
Mean	3.66	3.21	−0.45

C3DAN reduces the cross-seed standard deviation for 7 out of 9 subjects. The paired t-test confirms that this suppression is statistically significant: $t = -2.40$, $p = 0.0214$ (one-sided). The Wilcoxon signed-rank test yields a consistent result: $W = 6.0$, $p = 0.0273$ (one-sided).

This result validates the design motivation of ccCORAL: by explicitly anchoring the class-conditional second-order statistics, it provides a deterministic optimization direction for adversarial gradients, thereby reducing the stochastic fluctuations inherent in adversarial training.

4.5.2. EAEA Gating Analysis and Atypical Subjects

The ablation study reveals that EA-based methods (EA-CDAN, MAC-DA without gating) show decreased accuracy compared with non-EA methods (CDAN, C3DAN) on subjects S2 and S5, while showing substantial improvements on all other subjects. Table 13 quantifies the EA effect per subject.

Table 13: EA effect (Δ Acc, pp) and EAEA gating decisions. The gating criterion $\delta_t < \mu_\delta - \sigma_\delta$ correctly identifies subjects with negative EA effect.

Subject	Δ Acc (pp)	δ_s	EAEA Gate
S1	+10.56	typical	EA on
S2	−9.17	low	EA off †
S3	+5.14	typical	EA on
S4	+1.11	typical	EA on
S5	−5.76	low	EA off †
S6	+4.37	typical	EA on
S7	+4.72	typical	EA on
S8	−2.43	typical	EA on
S9	+7.37	typical	EA on

The Mann-Whitney U-test confirms that the EA effect is significantly different between the gated group (S2, S5) and the non-gated group: $U = 0.0$, $p = 0.0278$ (one-sided). This validates the EAEA gating criterion: subjects whose log-determinant δ falls below $\mu - \sigma$ are precisely those for whom EA induces negative transfer.

The practical impact of EAEA gating is substantial. Without gating, the non-gated MAC-DA achieves 68.24% mean accuracy with marginal significance over competing DA methods ($p = 0.059$ vs. CDAN). With EAEA gating, MAC-DA achieves 69.77% mean accuracy with clear significance over all competitors ($p \leq 0.003$). The +1.53 pp improvement comes entirely from converting EA-induced negative transfer on S2 and S5 into preserved C3DAN-level performance, without affecting any other subject. This demon-

strates that EAEA gating is a zero-cost mechanism that strictly improves the framework's robustness.

To further characterize the practical behavior of EA on the 2a dataset, we additionally compare fixed EA (always on) against non-EA and EAEA gating from a subject-count perspective. Under fixed EA, negative transfer appears on 3/9 subjects (S2, S5, S8), while EAEA gating reduces this to 1/9 subject (S8). In terms of gate-detection quality on 2a, EAEA correctly captures 2 out of 3 harmful subjects with no false positives (precision = 1.00, recall = 0.67, F1 = 0.80, balanced accuracy = 0.83). These results are consistent with Table 13: the current gate is conservative (very low false alarms) yet still leaves room to improve recall.

4.5.3. Cross-Dataset Subject-Level Gate Classification (2a + Weibo2014)

Beyond the fold-wise hard gate used in MAC-DA, we conducted an auxiliary subject-level LOSO binary classification experiment on the pooled 19 subjects from 2a (9) and Weibo2014 (10), where the positive class denotes "EA causes negative transfer" (6 positives, 13 negatives). We compared lightweight learned classifiers and rule-based baselines using subject-level covariance descriptors (details in the supplementary experiment scripts).

Table 14: Subject-level LOSO gate classification on pooled 2a+Weibo2014 (19 subjects). Positive class: "EA harmful".

Method	Acc	Bal. Acc	Precision	Recall	F1
`logreg_l2`	31.58	36.54	23.08	50.00	31.58
`svm_rbf`	36.84	35.90	20.00	33.33	25.00
`rf_shallow`	42.11	30.77	0.00	0.00	0.00
`always_non_negative`	68.42	50.00	0.00	0.00	0.00
`always_negative`	31.58	50.00	31.58	100.00	48.00
`logdet_1sigma_dataset`	**84.21**	**75.00**	**100.00**	50.00	**66.67**

As shown in Table 14, the dataset-wise log-determinant 1-σ rule achieves the best overall performance (Acc = 84.21%, F1 = 66.67%), with perfect precision (1.00) and zero false positives (TN=13, FP=0). This indicates that the current EAEA criterion is highly reliable at avoiding unnecessary EA deactivation for normal subjects. Meanwhile, recall remains moderate (0.50; FN=3), suggesting that future adaptive/learnable gating strategies should primarily target higher recall while preserving the low-FP property.

4.5.4. Feature Visualization

To visualize the effectiveness of domain adaptation, we applied t-SNE [70] to project the deep features extracted by different methods into 2D space. Figure 5 shows the t-SNE visualization for a representative leave-one-out task.

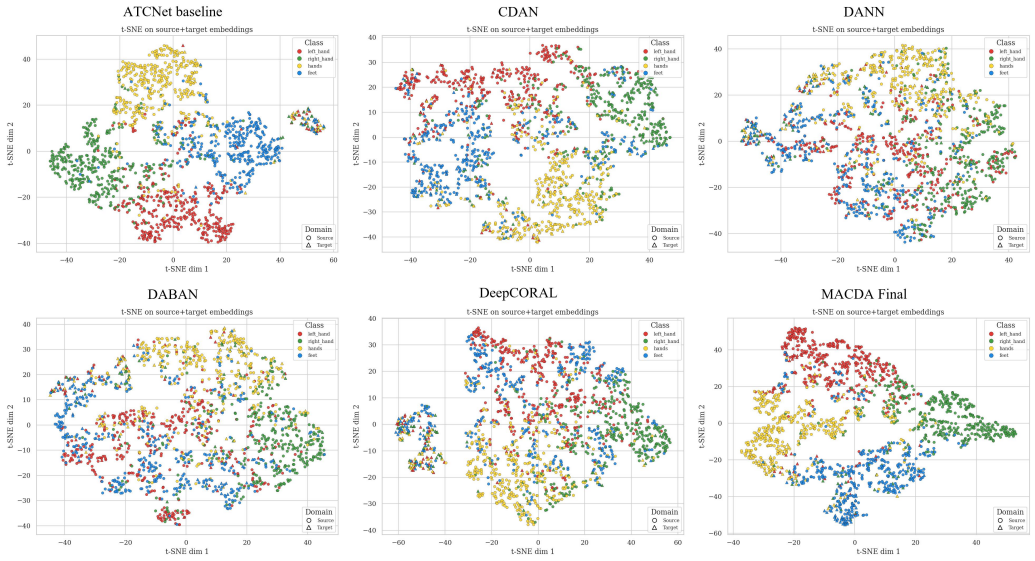

Figure 5: t-SNE visualization of source-target deep features on the leave-one-subject-out Weibo2014 setting for Subject S5 (seed = 0). Compared with the baselines, MAC-DA exhibits tighter intra-class clustering and clearer inter-class separation, indicating more discriminative and transferable feature representations.

From Figure 5, we observe that MAC-DA produces the most discriminative feature space, with compact intra-class clusters and clear inter-class separation. This confirms that the multi-space progressive alignment strategy effectively learns domain-invariant yet class-discriminative representations.

4.5.5. Confusion Matrix Analysis

To intuitively understand the classification improvements brought by our proposed framework, Figure 6 visualizes the confusion matrices of all compared methods for a representative target domain (Subject 1).

The visual progression across the panels clearly demonstrates the effectiveness of the MAC-DA architecture. Non-transfer baselines, such as the ATCNet backbone, struggle with severe class confusion due to unmitigated

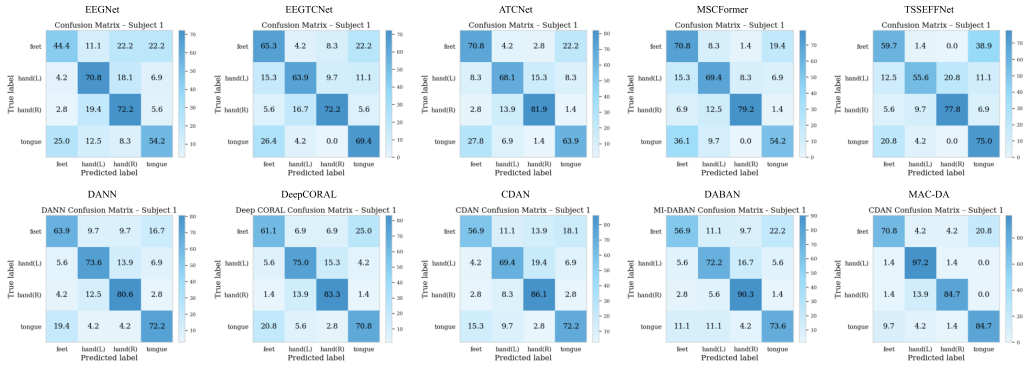

Figure 6: Confusion matrices across different baseline and domain adaptation methods for Subject 1. Rows represent true labels, and columns represent predicted labels. The proposed MAC-DA framework exhibits the most distinct diagonal dominance, significantly reducing misclassifications compared to both non-transfer baselines and conventional adaptation methods.

domain shifts. While standard domain adaptation methods like DANN and CDAN alleviate some macroscopic misalignment, they still struggle with fine-grained intra-class boundaries, often confusing left-hand imagery with right-hand or feet tasks.

In contrast, MAC-DA produces a highly discriminative and balanced prediction space across all four motor imagery classes. Notably, our framework correctly classifies an impressive 97.2% of left-hand trials and 84.7% of right-hand trials, effectively resolving the neurophysiological similarity challenges that plague other models. This confirms that our multi-space progressive alignment strategy not only bridges the cross-subject domain gap but also rigorously protects the intra-class topological boundaries.

4.6. Summary

The experimental results comprehensively validate the effectiveness of the proposed MAC-DA framework:

1) **Superior overall performance on both datasets:** MAC-DA achieves 69.77% mean accuracy on the BCI Competition IV 2a dataset (+9.48 pp over ATCNet, $p < 0.001$) and 53.70% on the Weibo2014 dataset (+7.30 pp, $p = 0.002$), significantly outperforming all competing domain adaptation methods on both benchmarks. The consistent performance hierarchy across two independent datasets with different recording systems

and task compositions validates the generalizability of the multi-space progressive alignment strategy.

2) **EAEA gating prevents negative transfer:** The proposed EAEA gating mechanism automatically identifies subjects with atypical covariance structures and disables EA for those folds, converting EA-induced negative transfer into preserved performance. This raises the overall accuracy by +1.53 pp and elevates all comparisons with competing methods to clear statistical significance.
3) **Validated component contributions:** Ablation studies on the 2a dataset confirm that CDAN+E, EA, ccCORAL, and EAEA gating each contribute complementary benefits. EA-only already improves the backbone by +4.13 pp (EA-ATCNet vs. ATCNet, $p = 0.0095$), CDAN+E provides the largest additional adaptation gain (+5.87 pp over ATCNet), EAEA adds +3.26 pp over C3DAN ($p = 0.004$), and ccCORAL enhances training stability (+0.35 pp with significant variance suppression).
4) **Significant variance suppression:** ccCORAL reduces cross-seed standard deviation by 0.45 pp on average ($p = 0.0214$), validating its role as a training stabilizer.

These results establish MAC-DA as a highly effective and robust solution for cross-subject MI-EEG decoding, addressing the critical challenge of distribution shift that has long hindered the clinical translation of BCI technology.

5. Conclusion

This paper presented MAC-DA, a multi-space progressive domain adaptation framework for cross-subject MI-EEG decoding. By recognizing that cross-subject distribution shifts are a composite problem spanning from physical sensor-level drifts to high-dimensional semantic discrepancies, MAC-DA systematically addresses each level through four complementary mechanisms: (1) Entropy-Aware Euclidean Alignment (EAEA) with automatic hard gating to adaptively normalize the physical manifold while preventing negative transfer on subjects with atypical covariance structures; (2) a hybrid convolution-attention backbone (ATCNet) for high-fidelity spatiotemporal feature extraction; (3) entropy-aware conditional adversarial learning

(CDAN+E) for implicit global marginal distribution alignment; and (4) class-conditional CORAL (ccCORAL) for explicit local semantic anchoring that stabilizes the adversarial optimization landscape.

Extensive experiments on two public MI-EEG benchmarks under the rigorous LOSO protocol demonstrated that MAC-DA achieves a mean accuracy of 69.77% on the BCI Competition IV 2a dataset and 53.70% on the Weibo2014 dataset, significantly outperforming the ATCNet baseline (by 9.48 pp and 7.30 pp respectively, $p \leq 0.002$) and all competing domain adaptation methods on both datasets. The consistent performance ranking across two independent datasets with different recording systems, channel configurations, and task compositions validates the cross-dataset generalizability of the proposed framework. Ablation studies confirmed the orthogonal and complementary nature of each component, while statistical analyses validated ccCORAL's role as a training stabilizer that significantly suppresses cross-seed variance ($p = 0.0214$). The EAEA gating mechanism proved to be a zero-cost robustness enhancer, automatically identifying and protecting atypical subjects from EA-induced negative transfer.

Despite these promising results, several limitations warrant discussion. First, although the evaluation now spans two datasets with 19 subjects in total, further validation on larger-scale, multi-session datasets (e.g., OpenBMI with 54 subjects, BNCI Horizon 2020) and other BCI paradigms (e.g., SSVEP, P300) would further strengthen the generalizability claims. Second, the EAEA gating mechanism relies on a fixed 1-σ threshold derived from population statistics, which may not be optimal for all cohort compositions; an adaptive or learnable gating criterion could further improve robustness. Third, the pseudo-label quality in ccCORAL remains dependent on the classifier's discriminative ability, and more sophisticated pseudo-label refinement strategies (e.g., curriculum-based or prototype-based methods) may yield additional gains. Finally, the current framework operates in an offline, batch-training setting; extending MAC-DA to online or incremental adaptation scenarios would be essential for real-time clinical deployment.

Future work will focus on three directions: (1) extending MAC-DA to larger-scale and cross-paradigm transfer scenarios to further validate its scalability; (2) exploring learnable gating mechanisms that can dynamically adjust the alignment strategy during training rather than relying on a fixed preprocessing decision; and (3) integrating MAC-DA with emerging self-supervised pre-training paradigms to leverage large-scale unlabeled EEG data for improved feature initialization. We believe that the multi-space progressive

alignment philosophy underlying MAC-DA provides a principled and extensible foundation for advancing calibration-free, clinical-grade brain-computer interfaces.

References

[1] Altaheri H, Muhammad G, Alsulaiman M, Amin SU, Altuwaijri GA, Abdul W, Bencherif MA, Faisal M. Deep learning techniques for classification of electroencephalogram (EEG) motor imagery (MI) signals: a review. *Neural Computing and Applications*. 2023;35:14681–14722.

[2] Mane R, Chouhan T, Guan C. BCI for stroke rehabilitation: motor and beyond. *Journal of Neural Engineering*. 2020;17(4):041001.

[3] Saha S, Baumert M. Intra- and inter-subject variability in EEG-based sensorimotor brain computer interface: a review. *Frontiers in Computational Neuroscience*. 2020;13:87.

[4] Rashid M, Sulaiman N, Abdul Majeed APP, Musa RM, Ab Nasir AF, Bari BS, Khatun S. Current status, challenges, and possible solutions of EEG-based brain-computer interface: a comprehensive review. *Frontiers in Neurorobotics*. 2020;14:25.

[5] Cheng N, Phua KS, Lai HS, Tam PKS, Tang KY, Cheng KKY, Yeow RCH, Gassert R, Burdet E, Campolo D, Ang KK, Chua KSG. Brain-computer interface-based soft robotic glove rehabilitation for stroke. *IEEE Transactions on Biomedical Engineering*. 2020;67(12):3339–3351.

[6] Jeong JH, Shim KH, Kim DJ, Lee SW. Brain-controlled robotic arm system based on multi-directional CNN-BiLSTM network using EEG signals. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2020;28(5):1226–1238.

[7] Baniqued PDE, Stanyer EC, Awais M, Alazmani A, Jackson AE, Mon-Williams MA, Mushtaq F, Holt RJ. Brain-computer interface robotics for hand rehabilitation after stroke: a systematic review. *Journal of NeuroEngineering and Rehabilitation*. 2021;18(1):15.

[8] Wan Z, Yang R, Huang M, Zeng N, Liu X. A review on transfer learning in EEG signal analysis. *Neurocomputing*. 2021;421:1–14.

[9] Zhang K, Robinson N, Lee SW, Guan C. Adaptive transfer learning for EEG motor imagery classification with deep convolutional neural network. *Neural Networks*. 2021;136:1–12.

[10] Wu D, Xu Y, Lu BL. Transfer learning for EEG-based brain-computer interfaces: a review of progress made since 2016. *IEEE Transactions on Cognitive and Developmental Systems*. 2022;14(1):4–19.

[11] Xu L, Xu M, Jung TP, Ming D. Review of brain encoding and decoding mechanisms for EEG-based brain-computer interface. *Cognitive Neurodynamics*. 2021;15(4):569–584.

[12] Li D, Robinson N, Zhang S, Guan C. A multisource transfer learning method for cross-subject electroencephalogram classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2024;32:1874–1885.

[13] Wu D, Jiang X, Peng R, Kong W, Huang J, Zeng Z. Online and offline domain adaptation for reducing BCI calibration effort. *IEEE Transactions on Human-Machine Systems*. 2022;52(5):1038–1049.

[14] Joadder MAM, Myszka S, Rahman MA, Grzegorzek M, Mukhopadhyay S. A performance-based adaptive calibration scheme for motor imagery brain-computer interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2024;32:1–11.

[15] Zhang W, Wu D. Manifold embedded knowledge transfer for brain-computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2020;28(5):1117–1127.

[16] Zhao H, Zheng Q, Ma K, Li H, Zheng Y. Deep representation-based domain adaptation for nonstationary EEG classification. *IEEE Transactions on Neural Networks and Learning Systems*. 2021;32(2):535–545.

[17] Wilson G, Cook DJ. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology*. 2020;11(5):1–46.

[18] Ju C, Guan C. Tensor-CSPNet: a novel geometric deep learning framework for motor imagery classification. *IEEE Transactions on Neural Networks and Learning Systems*. 2023;34(12):10955–10969.

[19] Yair O, Ben-Chen M, Talmon R. Parallel transport on the cone manifold of SPD matrices for domain adaptation. *IEEE Transactions on Signal Processing*. 2019;67(7):1797–1811.

[20] Azab AM, Mihaylova L, Ang KK, Arvaneh M. Weighted transfer learning for improving motor imagery-based brain-computer interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2021;29:1727–1736.

[21] She Q, Chen T, Fang F, Zhang J, Gao Y, Zhang Y. Improved domain adaptation network based on Wasserstein distance for motor imagery EEG classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2023;31:1137–1148.

[22] Xu L, Xu M, Ke Y, An X, Liu S, Ming D. Cross-dataset variability problem in EEG decoding with deep learning. *Frontiers in Human Neuroscience*. 2020;14:103.

[23] Ramoser H, Muller-Gerking J, Pfurtscheller G. Optimal spatial filtering of single trial EEG during imagined hand movement. *IEEE Transactions on Rehabilitation Engineering*. 2000;8(4):441–446.

[24] Ang KK, Chin ZY, Wang C, Guan C, Zhang H. Filter bank common spatial pattern algorithm on BCI competition IV datasets 2a and 2b. *Frontiers in Neuroscience*. 2012;6:39.

[25] Barachant A, Bonnet S, Congedo M, Jutten C. Multiclass brain-computer interface classification by Riemannian geometry. *IEEE Transactions on Biomedical Engineering*. 2012;59(4):920–928.

[26] Congedo M, Barachant A, Bhatia R. Riemannian geometry for EEG-based brain-computer interfaces; a primer and a review. *Brain-Computer Interfaces*. 2017;4(3):155–174.

[27] Lotte F, Bougrain L, Cichocki A, Clerc M, Congedo M, Rakotomamonjy A, Yger F. A review of classification algorithms for EEG-based brain-computer interfaces: a 10 year update. *Journal of Neural Engineering*. 2018;15(3):031005.

[28] Lawhern VJ, Solon AJ, Waytowich NR, Gordon SM, Hung CP, Lance BJ. EEGNet: a compact convolutional neural network for

EEG-based brain-computer interfaces. *Journal of Neural Engineering*. 2018;15(5):056013.

[29] Schirrmeister RT, Springenberg JT, Fiederer LDJ, Glasstetter M, Eggensperger K, Tangermann M, Hutter F, Burgard W, Ball T. Deep learning with convolutional neural networks for EEG decoding and visualization. *Human Brain Mapping*. 2017;38(11):5391–5420.

[30] Sakhavi S, Guan C, Yan S. Learning temporal information for brain-computer interface using convolutional neural networks. *IEEE Transactions on Neural Networks and Learning Systems*. 2018;29(11):5619–5629.

[31] Fahimi F, Zhang Z, Goh WB, Lee TS, Ang KK, Guan C. Inter-subject transfer learning with an end-to-end deep convolutional neural network for EEG-based BCI. *Journal of Neural Engineering*. 2019;16(2):026007.

[32] Song Y, Zheng Q, Liu B, Gao X. EEG conformer: convolutional transformer for EEG decoding and visualization. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2023;31:710–719.

[33] Miao Z, Zhao M, Zhang X, Ming D. LMDA-Net: a lightweight multi-dimensional attention network for general EEG-based brain-computer interface paradigms and interpretability. *NeuroImage*. 2023;276:120209.

[34] Zhao X, Zhang H, Zhu G, You F, Kuang S, Sun L. A multi-branch 3D convolutional neural network for EEG-based motor imagery classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2023;31:1608–1617.

[35] Altaheri H, Muhammad G, Alsulaiman M. Physics-informed attention temporal convolutional network for EEG-based motor imagery classification. *IEEE Transactions on Industrial Informatics*. 2023;19(2):2249–2258.

[36] Kostas D, Aroca-Ouellette S, Rudzicz F. BENDR: using transformers and a contrastive self-supervised learning task to learn from massive amounts of EEG data. *Frontiers in Human Neuroscience*. 2021;15:653659.

[37] Banville H, Chehab O, Hyvärinen A, Engemann DA, Gramfort A. Uncovering the structure of clinical EEG signals with self-supervised learning. *Journal of Neural Engineering*. 2021;18(4):046020.

[38] Jayaram V, Barachant A. MOABB: trustworthy algorithm benchmarking for BCIs. *Journal of Neural Engineering*. 2018;15(6):066011.

[39] Kwon OY, Lee MH, Guan C, Lee SW. Subject-independent brain-computer interfaces based on deep convolutional neural networks. *IEEE Transactions on Neural Networks and Learning Systems*. 2020;31(10):3839–3852.

[40] Dose H, Møller JS, Iversen HK, Puthusserypady S. An end-to-end deep learning approach to MI-EEG signal classification for BCIs. *Expert Systems with Applications*. 2018;114:532–542.

[41] Yger F, Berar M, Lotte F. Riemannian approaches in brain-computer interfaces: a review. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2017;25(10):1753–1762.

[42] Zanini P, Congedo M, Jutten C, Said S, Berthoumieu Y. Transfer learning: a Riemannian geometry framework with applications to brain-computer interfaces. *IEEE Transactions on Biomedical Engineering*. 2018;65(5):1107–1116.

[43] Rodrigues PLC, Jutten C, Congedo M. Riemannian procrustes analysis: transfer learning for brain-computer interfaces. *IEEE Transactions on Biomedical Engineering*. 2019;66(8):2390–2401.

[44] He H, Wu D. Transfer learning for brain-computer interfaces: a Euclidean space data alignment approach. *IEEE Transactions on Biomedical Engineering*. 2020;67(2):399–410.

[45] He Z, Zhuang N, Bao G, Zeng Y, Yan B. Cross-day EEG-based emotion recognition using transfer component analysis. *Electronics*. 2022;11(4):651.

[46] Chevallier S, Kalunga EK, Barthélemy Q, Monacelli E. Review of Riemannian distances and divergences, applied to SSVEP-based BCI. *Neuroinformatics*. 2021;19(1):93–106.

[47] Xiao X, Fang Y, Liu Q, Zhou J. Source-free domain adaptation for cross-subject EEG motor imagery classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2023;31:3034–3045.

[48] Gaur P, McCreadie KA, Pachori RB, Wang H, Prasad G. An automatic subject specific channel selection method for enhancing motor imagery classification in EEG-BCI using correlation. *Biomedical Signal Processing and Control*. 2021;68:102574.

[49] Lotte F, Guan C. Regularizing common spatial patterns to improve BCI designs: unified theory and new algorithms. *IEEE Transactions on Biomedical Engineering*. 2011;58(2):355–362.

[50] Ganin Y, Ustinova E, Ajakan H, Germain P, Larochelle H, Laviolette F, Marchand M, Lempitsky V. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*. 2016;17(1):2096–2030.

[51] Long M, Cao Z, Wang J, Jordan MI. Conditional adversarial domain adaptation. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2018;31:1640–1650.

[52] Fu B, Li F, Ji Y, Li Y, Xie X, Shi G. SCDAN: learning common feature representation of brain activation for intersubject motor imagery EEG decoding. *IEEE Transactions on Instrumentation and Measurement*. 2023;72:2518315.

[53] Li H, Zhang D, Xie J. MI-DABAN: a dual-attention-based adversarial network for motor imagery classification. *Computers in Biology and Medicine*. 2023;152:106420.

[54] Chen Y, Yang R, Huang M, Wang Z, Liu X. Single-source to single-target cross-subject motor imagery classification based on multisubdomain adaptation network. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2022;30:1992–2002.

[55] Jeon E, Ko W, Suk HI. Domain adaptation with source selection for motor-imagery based BCI. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2022;30:1602–1612.

[56] Roy Y, Banville H, Albuquerque I, Gramfort A, Falk TH, Faubert J. Deep learning-based electroencephalography analysis: a systematic review. *Journal of Neural Engineering*. 2019;16(5):051001.

[57] Dolzhikova IA, Sameni R, Zollanvari A. Subject-independent classification of motor imagery tasks in EEG using multisubject ensemble CNN. *IEEE Access*. 2022;10:81355–81363.

[58] Hang W, Feng W, Liang S, Yu Q, Wang Q, Liu K, Chen L, Li Y. Cross-subject EEG signal recognition using deep domain adaptation network. *IEEE Access*. 2019;7:128273–128282.

[59] Gretton A, Borgwardt KM, Rasch MJ, Schölkopf B, Smola A. A kernel two-sample test. *Journal of Machine Learning Research*. 2012;13:723–773.

[60] Arjovsky M, Chintala S, Bottou L. Wasserstein generative adversarial networks. In: *International Conference on Machine Learning (ICML)*. 2017:214–223.

[61] Sun B, Saenko K. Deep CORAL: correlation alignment for deep domain adaptation. In: *European Conference on Computer Vision (ECCV) Workshops*. 2016:443–450.

[62] Zhu Y, Zhuang F, Wang J, Ke G, Chen J, Bian J, Xiong H, He Q. Deep subdomain adaptation network for image classification. *IEEE Transactions on Neural Networks and Learning Systems*. 2021;32(4):1713–1722.

[63] Kang G, Jiang L, Yang Y, Hauptmann AG. Contrastive adaptation network for unsupervised domain adaptation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019:4893–4902.

[64] Li S, Xie M, Gong K, Liu CH, Wang Y, Li W. Transferable semantic augmentation for domain adaptation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021:11516–11525.

[65] Saito K, Kim D, Sclaroff S, Darrell T, Saenko K. Semi-supervised domain adaptation via minimax entropy. In: *IEEE International Conference on Computer Vision (ICCV)*. 2019:8050–8058.

[66] Brunner C, Leeb R, Müller-Putz GR, Schlögl A, Pfurtscheller G. BCI Competition 2008–Graz data set A. *Institute for Knowledge Discovery (Laboratory of Brain-Computer Interfaces), Graz University of Technology*. 2008;16:1–6.

[67] Ingolfsson TM, Hersche M, Wang X, Kobayashi N, Cavigelli L, Benini L. EEG-TCNet: an accurate temporal convolutional network for embedded motor-imagery brain-machine interfaces. In: *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 2020:2958–2965.

[68] Li Y, Guo L, Liu Y, Liu J, Meng F. A temporal-spectral-based squeeze-and-excitation feature fusion network for motor imagery EEG decoding. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2021;29:1534–1545.

[69] Zhao W, Zhang B, Zhou H, Wei D, Huang C, Lan Q. Multi-scale convolutional transformer network for motor imagery brain-computer interface. *Scientific Reports*. 2025;15:1234.

[70] van der Maaten L, Hinton G. Visualizing data using t-SNE. *Journal of Machine Learning Research*. 2008;9:2579–2605.

[71] Yi W, Qiu S, Qi H, Zhang L, Wan B, Ming D. EEG feature comparison and classification of simple and compound limb motor imagery. *PLoS One*. 2014;9(12):e114853.

[72] Chevallier S, Carrara I, Aristimunha B, Guetschel P, Sedlar S, Lopes B, Velut S, Khazem S, Moreau T. The largest EEG-based BCI reproducibility study for open science: the MOABB benchmark. *arXiv preprint arXiv:2404.15319*. 2024.

[73] Lévy J, Zhang M, Pinet S, Rapin J, Banville H, d'Ascoli S, King JR. Brain-to-text decoding: a non-invasive approach via typing. *arXiv preprint arXiv:2502.17480*. 2025.