Highlights

HEAT: Hierarchical Emotion Adaptation with Progressive Thresholding for EEG Emotion and Consciousness Detection

Zhepei Hong, Rongtao Chen, Liting Li, Jiajun Chen, Wei Gao, Jiahui Pan

- A hierarchical adversarial alignment adaptively balances global and local feature adaptation.
- A curriculum-based pseudo-labeling strategy effectively leverages unlabeled data.
- SOTA cross-subject emotion recognition is achieved on the SEED and SEED-IV datasets.
- Consciousness levels in DOC patients can be assessed by detecting emotion-specific neural responses.

HEAT: Hierarchical Emotion Adaptation with Progressive Thresholding for EEG Emotion and Consciousness Detection

Zhepei Hong[a,1], Rongtao Chen[a,1], Liting Li[a], Jiajun Chen[a], Wei Gao[a,b,*] and Jiahui Pan[a,b,*]

[a]School of Artificial Intelligence, South China Normal University, Guangzhou, 510631, China
[b]Research Center for Brain-Computer Interface, Pazhou Laboratory, Guangzhou, 510330, China

ARTICLE INFO

Keywords:
EEG-based emotion recognition
consciousness detection
cross-subject
hierarchical domain adaptation
progressive pseudo-labeling

Abstract

Electroencephalography (EEG)-based emotion recognition is a powerful tool for affective computing and brain-computer interfaces but is hindered by cross-subject variability. Existing domain-adaptation methods often only align marginal distributions and ignore emotion-specific discrepancies, while many semi-supervised approaches rely on rigid pseudo-labeling rules, which limit model generalization. To address these issues, we propose Hierarchical Emotion Adaptation with Progressive Thresholding (HEAT), a novel framework for enhancing cross-subject EEG analysis. HEAT introduces two key innovations: (1) Hierarchical Adversarial Alignment that dynamically balances global and local feature adaptation using adaptive weighting, ensuring domain invariance while preserving emotion-specific neural patterns; and (2) Progressive Pseudo-Labeling that employs a curriculum-based adaptive confidence threshold and probability calibration to effectively leverage unlabeled data, thereby improving model robustness and generalization. Our experiments demonstrate that HEAT achieves state-of-the-art cross-subject emotion recognition accuracy of 95.37% ± 5.17% on SEED and 82.50% ± 9.28% on SEED-IV. Furthermore, HEAT successfully distinguishes between minimally conscious state (MCS) and vegetative state (VS) using emotion-evoked EEG responses in a self-collected dataset. The resulting detection scores correlate with consciousness levels, highlighting HEAT's potential for objective consciousness assessment. Additionally, HEAT reveals anterior-temporal activation patterns consistent with emotional hemispheric specialization theories. Overall, HEAT effectively mitigates the impact of cross-subject EEG variability, advances affective computing, and provides a robust tool for consciousness evaluation.

1. Introduction

Affective brain-computer interfaces (BCIs) hold significant promise for affective computing, as well as for diagnosing and treating various neurological diseases. Among the physiological signals used for these interfaces, Electroencephalography (EEG) is particularly valuable as it directly reflects central nervous system (CNS) activity. Neural activity as captured via EEG cannot be voluntarily suppressed, thus providing a more transparent window into the neural underpinnings of both affective processes and pathological brain activity [1]. Leveraging this direct insight, EEG-based emotion recognition has become a major research focus, aiming to decode the neural signatures of discrete emotional states. Recent studies indicate a substantial performance gap between cross-subject and within-subject emotion recognition [2]. For instance, Zheng et al. achieved within-subject accuracy of 86.08% but only 70.56% in cross-subject scenarios on the SEED dataset [3]. Similarly, Mai et al. obtained 85.81% within-subject and 78.52% cross-subject accuracy [4]. This performance disparity highlights the critical need to address the cross-subject gap in EEG-based emotion recognition.

A primary challenge in EEG-based emotion recognition is the significant cross-subject variability. This variability stems from distribution discrepancies among different individuals, making the challenge fundamentally one of domain adaptation. To address this fundamental discrepancy and mitigate the impact of overall cross-subject variability, numerous studies have employed Domain-Adversarial Neural Networks (DANN), a method specifically designed to align domain distributions [5, 6, 7, 8]. Zhu et al. [9] introduced a multi-source, multi-task adaptation scheme that leverages hierarchical feature learning and homoscedastic uncertainty weighting. On the SEED and SEED-IV datasets [10], this approach achieved cross-subject classification accuracies of 91.36% and 72.25%, respectively. Wang

*Corresponding author
hongzhepei@m.scnu.edu.cn (Z. Hong); chenrongtao@m.scnu.edu.cn (R. Chen); llting@m.scnu.edu.cn (L. Li); chenjiajun@m.scnu.edu.cn (J. Chen); gaow@m.scnu.edu.cn (W. Gao); panjiahui@m.scnu.edu.cn (J. Pan)
[1]Zhepei Hong and Rongtao Chen contributed equally to this work.

et al. [11] proposed a selective graph domain-adaptation network that combines dynamic graph convolution with region-level alignment to capture domain-invariant functional connectivity. This model achieved cross-subject emotion recognition accuracies of 89.8% and 74.7% on the SEED and SEED-IV datasets, respectively. However, conventional domain adversarial alignment approaches primarily focus on aligning marginal probability distributions [7, 12], while neglecting the conditional probability distributions of specific emotion categories. This oversight can trigger negative transfer when source and target domains differ markedly, thereby limiting generalization.

Another significant challenge lies in the limitations of existing pseudo-labeling methods, which lack a dynamic strategy to update confidence thresholds in accordance with the model's evolving capabilities. As a key technique in semi-supervised learning, pseudo-labeling significantly improves data utilization efficiency and enhances cross-subject generalization by extending the training scope to unlabeled data [13, 14, 15]. For example, Du et al. [16] proposed the Residual Pseudo-Label Distance-aware Dual-Classifier model, which mitigated cross-subject mismatches and achieved 89.0% accuracy on the SEED dataset. Additionally, Zhong et al. [12] presented an unsupervised domain adaptation approach with pseudo-label propagation. After correction-and-smoothing refinement, their method achieved 89.4% accuracy on the SEED dataset. However, these methods typically fail to adaptively select suitable training samples at different stages of learning, often resulting in suboptimal utilization of unlabeled data and potential error propagation from misidentified pseudo-labels. Consequently, there is a pressing need for a more adaptive and dynamic pseudo-labeling strategy that can effectively leverage unlabeled data while accommodating the model's progression during training.

BCI paradigms for disorders of consciousness (DOC) assessment have mainly focused on detecting command-following and residual awareness through various approaches. These include motor imagery tasks where patients with DOC are asked to imagine moving their hands or feet [17], as well as paradigms based on event-related potentials, such as using an auditory oddball task to probe cognitive processing and communication abilities [18]. While these BCI methods have successfully identified signs of awareness in some patients with DOC, their approach centers on cognitive functions such as command-following rather than affective processing. Consequently, research has increasingly turned to the neurobiological link between emotion and consciousness. Because emotional processing is considered a high-level cognitive function and a fundamental aspect of conscious experience [19, 20], emotion recognition has emerged as a promising BCI paradigm for consciousness assessment. For instance, Qiu et al. applied a semi-supervised fine-tuning self-supervised graph attention network (SFT-SGAT) for consciousness detection by training the model on healthy subjects and applying it to DOC patients, reporting that the model achieved accuracies exceeding 60% for several patients [20]. In parallel, Cai et al. applied a self-supervised model trained on healthy subjects to decode music-evoked activity in patients with DOC [21]. They found classifiable responses in seven individuals, including six patients in the minimally conscious state (MCS) and one patient with unresponsive wakefulness syndrome (UWS), who subsequently showed clinical improvement and brain patterns that resembled healthy controls [21]. These studies underscore the potential of using emotion-decoding paradigms for consciousness assessment. However, while these frameworks learn globally robust representations, they lack dedicated mechanisms to explicitly align the fine-grained, class-conditional distributions that are crucial for preserving emotion-specific patterns across domains.

To address the substantial inter-subject variability in EEG-based emotion recognition and the limited utilization of unlabeled target-domain data, we propose Hierarchical Emotion Adaptation with Progressive Thresholding (HEAT). The framework is designed to preserve emotion-related neural information during cross-subject transfer while improving generalization in emotion recognition, and it further provides a promising technical route for objective assessment in patients with DOC using emotion-evoked EEG responses. The main contributions of this work are as follows:

- **Hierarchical Adversarial Alignment Strategy:** We present a hierarchical adversarial alignment mechanism that jointly models the global distribution shared across subjects and the local class-conditional distributions associated with different emotions. Unlike conventional adversarial transfer methods that emphasize only overall distribution consistency, this mechanism adaptively adjusts the alignment strength according to domain discrepancies at different levels, thereby alleviating cross-subject distribution shift while better preserving emotion-related discriminative neural representations.

- **Progressive Pseudo-Labeling Strategy:** We design a progressive pseudo-labeling strategy to improve the utilization quality of unlabeled target-domain samples. By combining curriculum-based confidence adjustment with probability calibration, the strategy allows pseudo-label selection to adapt to the model's training state,

reducing the adverse influence of noisy supervision in early stages, improving the stability of pseudo-label learning, and further enhancing cross-subject generalization.

- **State-of-the-Art Cross-Subject Emotion Recognition:** HEAT was systematically evaluated on the two public cross-subject EEG emotion recognition datasets, SEED and SEED-IV, and achieved accuracies of 95.37% ± 5.17% and 82.50% ± 9.28%, respectively, outperforming existing comparison methods overall. Ablation studies further showed that both hierarchical adversarial alignment and progressive pseudo-labeling contributed consistently to the performance gains, validating the effectiveness of the proposed framework design.

- **Objective Consciousness Assessment in DOC Patients:** We further extend the proposed framework to a cross-domain consciousness assessment task in patients with DOC, demonstrating its potential in clinically relevant settings. The results show that HEAT can capture neural representations associated with consciousness levels from emotion-evoked EEG responses, and that its recognition performance is positively correlated with Coma Recovery Scale-Revised (CRS-R) scores, suggesting that it may serve as an objective complement to conventional behavior-based assessment.

2. Related Work

2.1. Domain Adaptation Networks

A significant challenge in EEG-based emotion recognition is the inherent cross-subject variability of EEG signals, which creates domain differences between individuals. In recent years, domain adaptation techniques have emerged as effective tools to minimize distribution disparities and extract transferable features [22]. For example, Wang et al. [11] proposed a multi-source selective graph domain adaptation network that reduces cross-subject variability by aligning graph structures and feature representations across multiple source subjects. This approach captures consistent functional connectivity and regional brain states. Similarly, Imtiaz and Khan [23] introduced a gradual proximity-guided method to iteratively select reliable target domain samples for training, avoiding negative transfer from unreliable data. Zhou et al. [24] presented a prototypical-representation-based pairwise learning framework that redefines emotion recognition as a pairwise task. It learns subject-generalized emotion prototypes and employs pairwise comparisons to resist label noise and individual differences, thereby improving cross-subject generalization.

However, most adversarial adaptation studies focus only on aligning marginal probability distributions while leaving individual distribution disparities unresolved. This neglect of conditional probability distributions can lead to significant conditional shifts across domains due to substantial cross-subject variability. When an adaptation method fails to address this conditional shift, it often triggers negative transfer. Therefore, aligning both marginal and conditional distributions is crucial for reliable cross-subject knowledge transfer.

2.2. Pseudo-Labeling Technology

Pseudo-labeling is a widely used semi-supervised approach. It assigns provisional labels to large amounts of unlabeled data based on model predictions [25]. For instance, Peng et al. [26] proposed an EEG feature-transfer model for semi-supervised cross-subject emotion recognition. This model jointly optimizes a shared subspace projection matrix and pseudo target labels, iteratively refining its estimates of target emotional states and using them as soft labels to better align the conditional distributions of source and target EEG data. Similarly, Ren et al. [27] developed a semi-supervised pairwise learning method within a multi-source domain adaptation framework. Their method dynamically updates thresholds for unlabeled pairs during training, allowing the model to adaptively classify unlabeled data more reliably. Furthermore, Jimenez-Guarneros and Fuentes-Pineda [28] introduced a semi-supervised framework that uses a few labeled target samples to perform a fine-grained multi-source alignment of feature distributions at both subject and class levels.

However, conventional pseudo-labeling methods are often limited by their inability to dynamically adapt to the model's evolving capabilities. In early training stages, the model may encounter ambiguous samples beyond its capacity, leading to inefficient learning. As the model's proficiency increases, the same non-adaptive mechanism fails to provide more challenging, information-rich data, thereby underutilizing the model's capabilities. This persistent mismatch between sample difficulty and model capability results in a suboptimal learning trajectory, ultimately limiting the model's potential and final performance.

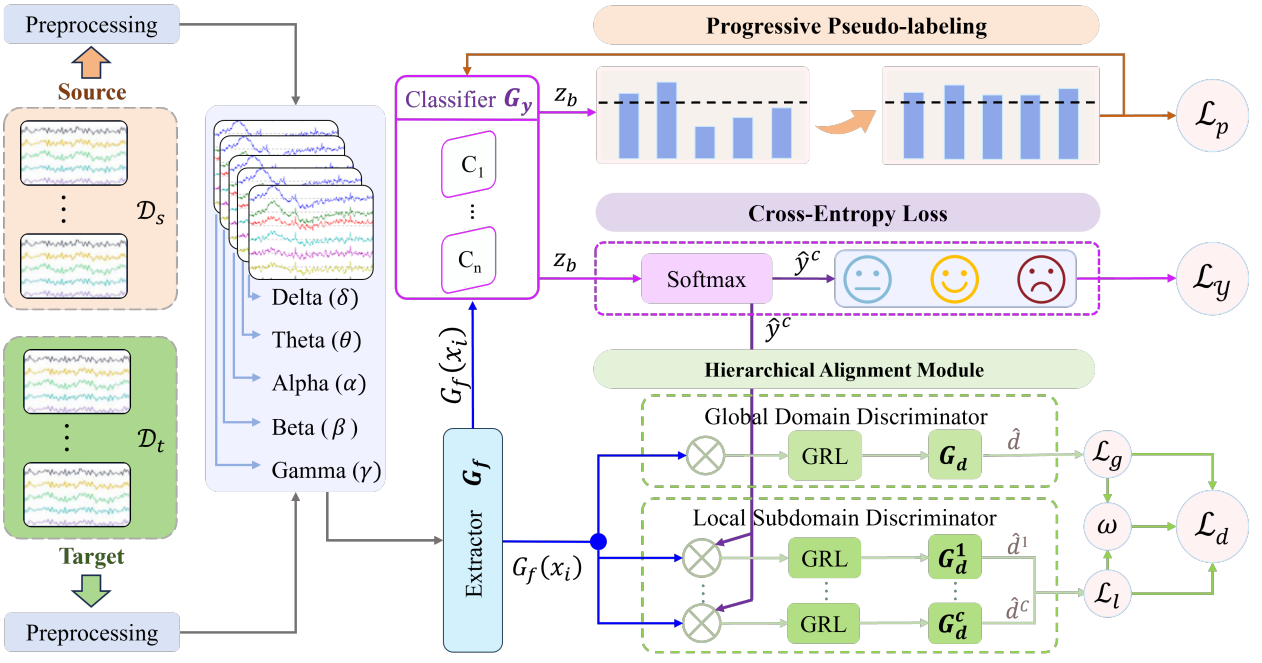

Figure 1: Overview of the proposed Hierarchical Emotion Adaptation with Progressive Thresholding (HEAT) framework: (A) An encoder (G_f) processes differential-entropy (DE) features to produce features $G_f(x_i)$, and a classifier (G_y) generates logits z_b. The logits yield class probabilities $\hat{y}^c$ used to compute the supervised classification loss ($\mathcal{L}_y$). (B) A hierarchical adversarial alignment mechanism that computes the domain adversarial loss ($\mathcal{L}_d$). It employs a global domain discriminator (G_d) and multiple emotion-specific local discriminators (G_d^c), with their respective losses dynamically balanced by an adaptive weighting factor (ω). (C) A progressive pseudo-labeling module takes the target domain logits z_t as input and generates a pseudo-label loss ($\mathcal{L}_p$) from the unlabeled target data. (D) The model is trained end-to-end by minimizing the overall training objective, $\mathcal{L}$, which is a weighted sum of $\mathcal{L}_y$, $\mathcal{L}_d$ and $\mathcal{L}_p$.

2.3. Consciousness Assessment of Patients with DOC

The clinical assessment of consciousness in patients with DOC primarily relies on standardized behavioral scales like the CRS-R, which is considered the gold standard. However, this reliance on observable motor responses is a critical limitation, often leading to misdiagnosis in patients with severe motor impairments who may retain residual awareness [29]. To overcome the limitations of behavioral observation scales like the CRS-R, researchers have increasingly turned to EEG-based BCIs as a promising objective assessment tool [30]. BCI paradigms for DOC assessment have mainly focused on detecting command-following and residual awareness through various approaches. These include motor imagery tasks where patients are asked to imagine moving their hands or feet [17], as well as paradigms based on event-related potentials, such as using an auditory oddball task to probe cognitive processing and communication abilities [18]. While these BCI methods have successfully identified signs of awareness in some patients, they are generally not designed to capture affective processing. Consequently, the systematic application of EEG-based emotion recognition as a primary BCI paradigm for consciousness assessment has received limited attention. In this context, the neurobiological link between emotion and consciousness offers a new opportunity. Consciousness is understood to consist of two core components: arousal and awareness [31, 19]. Emotional processing is not merely a passive response but a high-level cognitive function and a fundamental aspect of conscious experience [19, 20]. The ability to differentiate emotional stimuli serves as a direct indicator of residual awareness, providing a motor-independent window into a patient's inner world. Consequently, leveraging EEG-based emotion recognition to detect these emotion-specific patterns in DOC patients holds significant promise for a more accurate and nuanced assessment of consciousness.

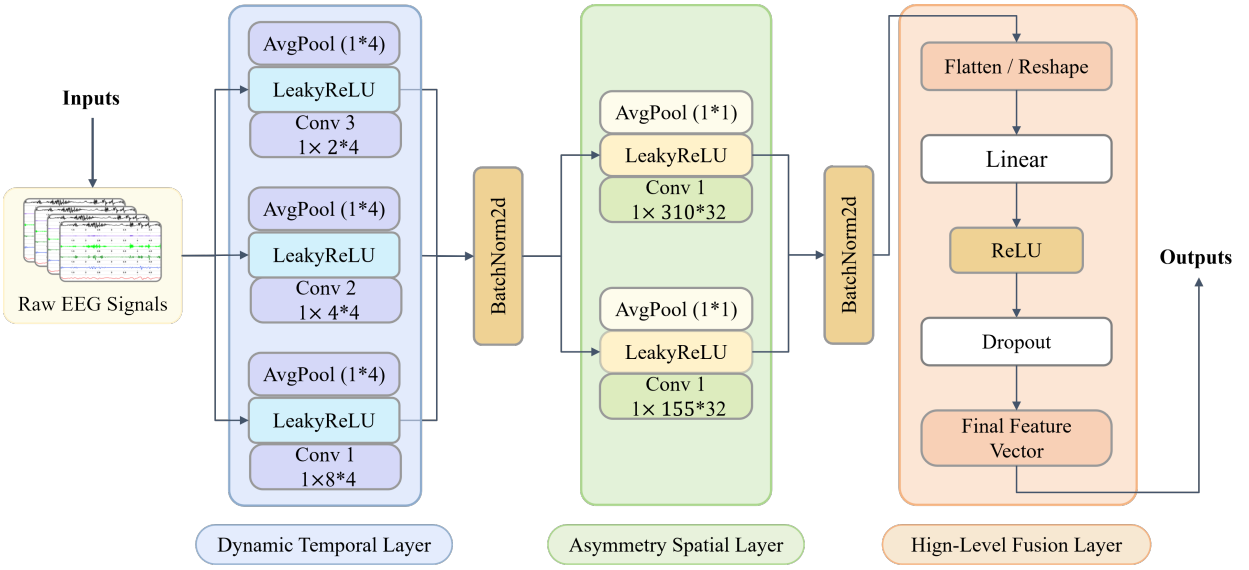

Figure 2: Architecture of the extractor, consisting of three modules: dynamic temporal, asymmetric spatial, and high-level dense.

3. Method

In this study, we propose a novel model designed to learn robust domain-invariant and emotion-specific patterns for cross-subject EEG analysis. Our model leverages this capability to address a critical challenge in the objective assessment of DOC: the robust detection of emotion-specific neural patterns in patients to reflect residual consciousness, a task complicated by significant cross-subject variability. The overall architecture of our proposed HEAT framework is illustrated in Fig. 1. For a cross-subject task, the model takes EEG trials from a source domain $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$ and a target domain $\mathcal{D}_t = \{(\mathbf{x}_i^t, y_i^t)\}_{i=1}^{N_t}$ as input. Initially, raw EEG signals undergo preprocessing to extract DE features [32]. As depicted in Fig. 1, these features are fed into a shared feature extractor (G_f) to generate spatiotemporal representations. Subsequently, these representations are processed through three parallel pathways: (1) For labeled source data, the features are passed to a classifier (G_y) to compute a supervised classification loss ($\mathcal{L}_y$). (2) Features from both domains are fed into the hierarchical alignment module, which aligns both marginal and class-conditional distributions to calculate the domain adversarial loss ($\mathcal{L}_d$). (3) For unlabeled target data, a curriculum-inspired progressive pseudo-labeling strategy is applied. This strategy adapts the confidence threshold and calibrates probabilities to generate a robust pseudo-label loss ($\mathcal{L}_p$). Finally, these loss components are integrated to form the overall training objective, enabling the framework to be trained end-to-end.

3.1. Feature Extraction

The feature extraction process is essential for capturing relevant spatiotemporal patterns in EEG signals, which are critical for effective cross-subject alignment and classification of emotional states. To effectively extract discriminative features, we employ TSCeption [33], a multi-scale convolutional neural network specifically designed to learn from EEG data by modeling its temporal dynamics and spatial asymmetry. As depicted in Fig. 2, the TSCeption architecture is composed of three sequential modules: a dynamic temporal layer, an asymmetric spatial layer, and a high-level fusion layer. The dynamic temporal layer employs three parallel branches to extract multi-scale temporal features from raw EEG signals, with each branch consisting of average pooling, LeakyReLU activation, and convolution. The outputs of these branches are concatenated along the channel dimension, followed by normalization, to obtain fused temporal features. Subsequently, the asymmetric spatial layer uses two parallel branches to capture global spatial features and hemisphere-specific spatial features, thereby modeling inter-hemispheric asymmetry that is particularly important for emotion recognition. Finally, the high-level fusion layer integrates the spatiotemporal features into a compact representation, which is subsequently fed into the semi-supervised learning and domain adaptation modules.

Consequently, the supervised classification loss is formulated as

$$\mathcal{L}_y = \frac{1}{|\mathcal{D}_s|} \sum_{(\mathbf{x}_s^i, y_s^i) \in \mathcal{D}_s} CE(G_y(G_f(\mathbf{x}_i^s)), y_i^s), \tag{1}$$

where $\mathcal{L}_y$ represents the classification loss function, CE is the cross-entropy loss, and G_f and G_y refer to the TSCeption-based feature extraction network and classification network, respectively. These representations enable the subsequent modules to handle cross-subject variability while capturing emotion-specific features.

3.2. Hierarchical Adversarial Alignment

Cross-subject EEG variability manifests not only as overall baseline physiological shifts, reflected in marginal distribution discrepancy, but also as variations in how specific emotions are neurally expressed across individuals, manifested as conditional distribution discrepancy. Conventional domain adaptation methods primarily focus on aligning marginal distributions while often neglecting emotion-specific structural differences, which can trigger negative transfer. To address this issue, we propose a Hierarchical Adversarial Alignment module that simultaneously enforces global domain invariance and local emotion-wise alignment. The complete computational flow of this module corresponds directly to Lines 8–14 in Algorithm 1.

First, to mitigate the global marginal distribution shift, we implement a global domain discriminator G_d paired with a Gradient Reversal Layer (GRL). As executed in Line 8 of Algorithm 1, this discriminator is optimized through adversarial training to make features from different subjects indistinguishable at the domain level. The global alignment loss is formulated as follows:

$$\mathcal{L}_g = \frac{1}{|\mathcal{D}_s| + |\mathcal{D}_t|} \sum_{\mathbf{x}_i \in \mathcal{D}_s \cup \mathcal{D}_t} BCE(G_d(G_f(\mathbf{x}_i)), d_i), \tag{2}$$

where $\mathcal{L}_g$ represents the global domain alignment loss, G_d is the global domain discriminator that attempts to distinguish between domains, G_f is the feature extraction network, BCE refers to the binary cross-entropy loss, and $d_i \in \{0, 1\}$ is the domain label, with 0 denoting the target domain and 1 denoting the source domain.

Parallel to the global alignment, we deploy dedicated local domain discriminators G_d^c for each emotion class $c \in \{1, \ldots, C\}$ to explicitly align the conditional probability distributions $p(\mathbf{x} \mid y = c)$ across domains. As executed in Line 9 of Algorithm 1, the objective of each local discriminator is to distinguish the domain origin of the input features specifically within the c-th class. The local discriminator loss is formulated as follows:

$$\mathcal{L}_l = \frac{1}{|\mathcal{D}_s| + |\mathcal{D}_t|} \sum_{c=1}^{C} \sum_{\mathbf{x}_i \in \mathcal{D}_s \cup \mathcal{D}_t} BCE(G_d^c(\hat{y}_i^c G_f(\mathbf{x}_i)), d_i), \tag{3}$$

where $\mathcal{L}_l$ represents the total local alignment loss across all emotion classes, $\mathcal{L}_l^c$ denotes the class-specific local discriminator loss for emotion class c, C is the total number of emotion classes, and $\hat{y}_i^c$ is the predicted probability that sample $\mathbf{x}_i$ belongs to class c.

To adaptively balance the contributions of global and local alignment, we quantify the domain divergence at both levels during training using the $\mathcal{A}$-distance metric [34]. As detailed in Lines 11–12 of Algorithm 1, the $\mathcal{A}$-distance for the global domain divergence $d_{\mathcal{A},g}$ and the emotion-specific local divergence $d_{\mathcal{A},l}$ for class c are calculated as follows:

$$d_{\mathcal{A},g}(\mathcal{D}_s, \mathcal{D}_t) = 2(1 - 2(\mathcal{L}_g)), \tag{4}$$

$$d_{\mathcal{A},l}(\mathcal{D}_s^c, \mathcal{D}_t^c) = 2(1 - 2(\mathcal{L}_l^c)), \tag{5}$$

where $\mathcal{L}_g$ represents the global domain discriminator loss, and $\mathcal{L}_l^c$ is the local discriminator loss for emotion class c.

Based on these divergence metrics, we formulate a dynamic adversarial factor ω in Line 13 of Algorithm 1 to adaptively weight the hierarchical components:

$$\omega = \frac{d_{\mathcal{A},g}(\mathcal{D}_s, \mathcal{D}_t)}{d_{\mathcal{A},g}(\mathcal{D}_s, \mathcal{D}_t) + \frac{1}{C}\sum_{c=1}^{C} d_{\mathcal{A},l}(\mathcal{D}_s^c, \mathcal{D}_t^c)}, \tag{6}$$

Algorithm 1 Training procedure of HEAT.

Require: Source data $\mathcal{D}_s$, Target data $\mathcal{D}_t$, Models $G_f, G_y, G_d, \{G_d^c\}_{c=1}^C$, Hyperparameters α, β, η, Epochs E

Initialize parameters θ for all models
for epoch = 1 to E **do**
  Sample batch $(\mathbf{x}_s, \mathbf{y}_s) \sim \mathcal{D}_s$ and batch $\mathbf{x}_t \sim \mathcal{D}_t$
  $f_s, f_t \leftarrow G_f(\mathbf{x}_s), G_f(\mathbf{x}_t)$ {*Forward Pass*}
  $\hat{\mathbf{y}}_s, \hat{\mathbf{y}}_t \leftarrow G_y(f_s), G_y(f_t)$
  $\mathcal{L}_y \leftarrow \frac{1}{|\mathbf{x}_s|} \sum_{(\mathbf{x}_s^i, \mathbf{y}_s^i)} \text{CrossEntropy}(\hat{\mathbf{y}}_s^i, \mathbf{y}_s^i)$ {*Classification Loss*}
  {*Hierarchical Alignment Module*}
  $\mathcal{L}_g \leftarrow \text{BCE}(G_d(f_s), 1) + \text{BCE}(G_d(f_t), 0)$
  $\mathcal{L}_l \leftarrow \sum_{c=1}^C \left(\text{BCE}(G_d^c(\hat{y}_s^c f_s), 1) + \text{BCE}(G_d^c(\hat{y}_t^c f_t), 0)\right)$
  {*Dynamic Weight Calculation*}
  $d_{\mathcal{A},g} \leftarrow 2(1 - 2\mathcal{L}_g)$
  $d_{\mathcal{A},l} \leftarrow \frac{1}{C} \sum_{c=1}^C 2(1 - 2\mathcal{L}_l^c)$
  $\omega \leftarrow d_{\mathcal{A},g}/(d_{\mathcal{A},g} + d_{\mathcal{A},l})$
  $\mathcal{L}_t \leftarrow (1 - \omega)\mathcal{L}_g + \omega\mathcal{L}_l$
  {*Progressive Pseudo-Labeling Module*}
  $\tau(t) \leftarrow \tau_{\text{init}} + \min(1, \frac{t}{T_{\text{total}}})(\tau_{\max} - \tau_{\text{init}})$
  $\mathbf{p}_t \leftarrow \text{Softmax}(\hat{\mathbf{y}}_t/T)$
  $M_b \leftarrow \mathbb{I}(\max(\mathbf{p}_t) \geq \tau(t))$
  $\mathcal{L}_p \leftarrow \frac{1}{|\mathbf{x}_t|} \sum_{\mathbf{x}_t^i} M_b^i \cdot \text{CrossEntropy}(\text{argmax}(\mathbf{p}_t^i), \hat{\mathbf{y}}_t^i)$
  {*Overall Objective & Parameter Update*}
  $\mathcal{L} \leftarrow \mathcal{L}_y + \alpha\mathcal{L}_t + \beta\mathcal{L}_p$
  $\theta \leftarrow \theta - \eta\nabla_\theta\mathcal{L}$
end for

Ensure: Trained models G_f, G_y

where $\omega \in (0, 1)$ quantifies the dynamic adversarial weighting factor, $d_{\mathcal{A},g}(\mathcal{D}_s, \mathcal{D}_t)$ measures the global distributional divergence between the source and target domains, and $d_{\mathcal{A},l}(\mathcal{D}_s^c, \mathcal{D}_t^c)$ measures the class-conditional divergence for emotion class c. The numerator captures the global domain divergence, whereas the denominator combines it with the average emotion-specific divergence across all classes.

Through the hierarchical alignment strategy and dynamic weighting mechanism, we integrate both global and local domain adaptation into a unified transfer loss function:

$$\mathcal{L}_t = (1 - \omega)\mathcal{L}_g + \omega\mathcal{L}_l. \tag{7}$$

where $\mathcal{L}_t$ denotes the total transfer loss; $\mathcal{L}_g$ denotes the global alignment loss; $\mathcal{L}_l$ denotes the local (emotion-specific) alignment loss. This hierarchical adversarial alignment framework effectively addresses both global domain shifts and emotion-specific variations, enabling more robust cross-subject emotion recognition by leveraging complementary information from different levels of feature representations.

3.3. Progressive Pseudo-Labeling with Confidence Curriculum Learning

In pseudo-labeling, a confidence threshold is crucial for selecting reliable predictions from unlabeled data. In cross-subject EEG scenarios, conventional pseudo-labeling can suffer from error propagation because the model may become overconfident during the early stages of training. To address this issue, we introduce a progressive pseudo-labeling strategy inspired by the "easy-to-hard" principle of curriculum learning [35]. Our pseudo-label generation follows the standard hard pseudo-labeling paradigm introduced by Lee [36], where the class with the maximum predicted probability is assigned as the provisional label for an unlabeled sample. Specifically, the strategy starts from a relatively low confidence threshold to encourage broad exploration of unlabeled target-domain samples and then progressively tightens the threshold as the model matures, thereby filtering noisy pseudo-labels more effectively. Furthermore, to prevent the model from assigning artificially high confidence to ambiguous EEG patterns, we incorporate temperature

Algorithm 2 Progressive pseudo-labeling strategy.

Require: Unlabeled target-domain logits z_t, current step t, total steps T_{total}, temperature T, thresholds τ_{init} and $\tau_{\max}$
Ensure: Progressive pseudo-label loss $\mathcal{L}_p$

1: $\tau(t) \leftarrow \tau_{\text{init}} + \min\left(1, \frac{t}{T_{\text{total}}}\right)(\tau_{\max} - \tau_{\text{init}})$ {Dynamic thresholding}
2: $\mathbf{p}(\mathbf{x}_i) \leftarrow \text{Softmax}(z_t/T)$ {Probability calibration}
3: $\hat{y}_i \leftarrow \arg\max(\mathbf{p}(\mathbf{x}_i))$ {Hard label generation}
4: $M_b \leftarrow \mathbb{I}\left(\max(\mathbf{p}(\mathbf{x}_i)) \geq \tau(t)\right)$ {Confidence masking}
5: $\mathcal{L}_p \leftarrow \frac{1}{\mu_B}\sum_{b=1}^{\mu_B} M_b \cdot H(\hat{y}_b, z_b)$ {Masked loss calculation}

scaling to calibrate the predicted probabilities. The complete computational flow of this progressive pseudo-labeling module is summarized in Algorithm 2.

Specifically, as outlined in Step 1 of Algorithm 2, the progressive confidence threshold is formulated as:

$$\tau(t) = \tau_{\text{init}} + \min\left(1, \frac{t}{T_{\text{total}}}\right)(\tau_{\max} - \tau_{\text{init}}), \tag{8}$$

where $\tau(t)$ represents the confidence threshold at training step t, τ_{init} denotes the initial threshold value, $\tau_{\max}$ is the maximum threshold value, t signifies the current training step, and T_{total} indicates the total number of training steps. This mechanism ensures that the threshold gradually increases from τ_{init} to $\tau_{\max}$ as the training progresses.

As shown in Step 2 of Algorithm 2, we use temperature scaling to suppress overconfident predictions while avoiding an excessively flattened probability distribution in cross-subject EEG scenarios. Therefore, for an input sample $\mathbf{x}_i$, we calibrate the probability distribution $\mathbf{p}(\mathbf{x}_i)$ using temperature scaling as follows:

$$\mathbf{p}(\mathbf{x}_i) = \text{Softmax}\left(\frac{G_y(G_f(\mathbf{x}_i))}{T}\right), \tag{9}$$

where $\mathbf{p}(\mathbf{x}_i)$ is the resulting probability distribution vector, $G_y(G_f(\mathbf{x}_i))$ represents the raw logits, and T is the temperature scaling coefficient.

Using the calibrated probabilities $\mathbf{p}(\mathbf{x}_i)$ and the current threshold $\tau(t)$, we select reliable unlabeled samples by computing a binary confidence mask M_b in Step 4 of Algorithm 2:

$$M_b = \mathbb{I}\left(\max(\mathbf{p}(\mathbf{x}_i)) \geq \tau(t)\right), \tag{10}$$

where $\mathbb{I}(\cdot)$ is the indicator function and $\max(\mathbf{p}(\mathbf{x}_i))$ denotes the maximum value within the predicted probability distribution for sample $\mathbf{x}_i$. This mask identifies samples whose prediction confidence meets the dynamically adjusted threshold.

Finally, the pseudo-label loss $\mathcal{L}_p$ is calculated using the confidence mask M_b and hard pseudo-labels $\hat{y}_b$ derived from the calibrated predictions in Step 5 of Algorithm 2:

$$\mathcal{L}_p = \frac{1}{\mu_B}\sum_{b=1}^{\mu_B} M_b \cdot H(\hat{y}_b, z_b) \tag{11}$$

where μ_B represents the number of unlabeled samples considered in the batch, M_b is the confidence mask for the b-th sample, $H(\cdot,\cdot)$ denotes the standard cross-entropy loss, $\hat{y}_b$ is the hard pseudo-label obtained via arg max from the calibrated probability, and z_b represents the corresponding model output logits $G_y(G_f(\mathbf{x}_i))$.

In essence, this approach implements an adaptive curriculum to leverage unlabeled EEG data in cross-subject scenarios. By selectively incorporating pseudo-labels through a calibrated and progressive mechanism, the model systematically balances initial exploration with focused exploitation of high-confidence affective signals. This ultimately improves the discriminability and robustness of the model against cross-subject EEG variability.

3.4. Overall Training Objective

The overall training objective $\mathcal{L}$ integrates the supervised classification loss, the hierarchical domain transfer loss, and the progressive pseudo-labeling loss into a unified framework. This combined objective aims to simultaneously

Table 1
Summary of the clinical status of patients.

Patients	Age	Gender	Clinical diagnosis	Etiology	Best CRS-R Score before One Week	Best CRS-R Score after One Month
P1	36	M	VS	NTBI	9 (2-1-2-2-0-2)	9 (2-1-2-2-0-2)
P2	26	F	VS	NTBI	6 (1-0-2-1-0-2)	6 (1-0-2-1-0-2)
P3	14	M	VS	NTBI	7 (1-1-2-1-0-2)	9 (1-3-2-1-0-2)
P4	55	M	MCS	TBI	11 (1-2-5-1-0-2)	16 (2-5-6-1-0-2)
P5	54	F	MCS	NTBI	16 (3-4-5-2-0-2)	17 (3-4-6-2-0-2)
P6	59	M	MCS	NTBI	17 (3-1-6-3-2-2)	17 (3-1-6-3-2-2)
P7	56	F	MCS	TBI	8 (0-3-2-1-0-2)	8 (0-3-2-1-0-2)
P8	58	M	MCS	TBI	6 (0-1-3-1-0-1)	12 (3-1-5-1-0-2)

Note: vegetative state (VS), traumatic brain injury (TBI), non-traumatic brain injury (NTBI). CRS-R scores are reported as total (subscores: auditory-visual-motor-oromotor-arousal-communication).

learn discriminative features from labeled source data, align feature distributions across domains, and leverage information from unlabeled target data in a controlled manner. The final objective function is formulated as a weighted sum:

$$\mathcal{L} = \mathcal{L}_y + \alpha\mathcal{L}_t + \beta\mathcal{L}_p, \tag{12}$$

where $\mathcal{L}$ is the final composite loss function guiding the model optimization, $\mathcal{L}_y$ represents the supervised classification loss (Eq. (1)) computed on labeled source data, $\mathcal{L}_t$ denotes the hierarchical adversarial alignment loss (Eq. (7)) for domain adaptation, $\mathcal{L}_p$ is the progressive pseudo-label loss (Eq. (11)) derived from unlabeled target data, and α, β are hyperparameters balancing the contribution of the domain transfer and pseudo-labeling objectives, respectively.

In our implementation, we set $\alpha = 15$ and $\beta = 0.8$. The selection of these values was guided by observations on a dedicated validation set. A relatively substantial weight for α was found necessary to ensure that the hierarchical adversarial alignment effectively steers the model towards domain-invariant representations. The value for β was carefully selected to harness the benefits of progressive pseudo-labeling in leveraging unlabeled target data, while also accounting for the sensitivity of this component to potential noise in the pseudo-labels.

Minimizing $\mathcal{L}$ facilitates end-to-end training, driving the HEAT model to simultaneously learn discriminative representations for accurate classification, achieve domain invariance crucial for cross-subject identification of consistent EEG emotion patterns, and effectively utilize unlabeled data to refine the understanding of subtle, emotion-specific neural features.

4. Experiments

We designed our experimental framework to validate the efficacy of the proposed method across three distinct experimental settings. The first experiment aims to demonstrate cross-subject emotion recognition performance using the public SEED [32] and SEED-IV [10] datasets. The second experiment focuses on assessing the method's effectiveness in detecting residual consciousness in patients with DOC, utilizing a self-recorded dataset comprising data from healthy controls and DOC patients. The third experiment investigates the impact of each model component through ablation studies conducted on the SEED and SEED-IV datasets.

4.1. Datasets

1) SEED: The SEED dataset is a comprehensive emotional database comprising EEG data from 15 participants (7 males and 8 females). It is organized around three emotional valences: positive, negative, and neutral. These emotional states are elicited using 15 video clips, each approximately 4 minutes in duration. Immediately following each clip, participants confirmed the elicited emotion by completing a self-assessment questionnaire. For each participant, EEG signals were recorded across three separate sessions conducted on different days, with each session consisting of 15 trials corresponding to each video stimulus.

Table 2
Model performance on SEED using LOSO cross-validation.

Methods	P_{acc} (%)	F1 score	Methods	P_{acc} (%)	F1 score
DANN [34]	81.65 ± 9.92	81.52 ± 10.01	JTSR [26]	84.69 ± 6.65	84.81 ± 6.53
PLMSDANet [27]	90.09 ± 9.58	89.87 ± 9.71	MFA-LR [25]	89.11 ± 7.72	89.26 ± 7.60
ST-SCGNN [37]	85.90 ± 4.90	85.74 ± 4.81	SFT-SGAT [20]	92.04 ± 3.28	92.15 ± 3.39
SSCDG [21]	87.90 ± 4.30	87.75 ± 4.18	MSGDA [11]	89.80 ± 3.40	89.92 ± 3.51
EEGMatch [38]	91.35 ± 7.03	91.14 ± 7.18	PR-PL [24]	93.06 ± 5.12	93.21 ± 5.04
MAS-DGAT-Net [39]	80.02 ± 5.79	79.86 ± 5.92	LGDAAN-Nets [40]	89.09 ± 7.76	84.85 ± 5.41
Ours	$\mathbf{95.37 \pm 5.17}$	$\mathbf{95.68 \pm 5.01}$			

2) SEED-IV: As an advancement on the SEED, SEED-IV provides EEG data related to four distinct emotional categories: happiness, sadness, fear, and neutrality. Emotion elicitation was performed using 24 culturally appropriate film segments, with six unique stimuli assigned to each of the four target emotions. The dataset retains the participant structure (15 subjects) and session framework (three sessions per subject) established in the original SEED. Furthermore, the neurophysiological recording procedures and overall experimental approach are comparable to those used in SEED.

3) Self-Recorded Dataset: EEG data were acquired using a 32-channel Neuroscan system (Computer Medicine Inc., Australia) with a sampling rate of 1000 Hz, a right mastoid reference, and a forehead ground, ensuring electrode impedance remained below 5 kΩ. Participants included ten healthy controls (HC) and eight patients with DOC recruited from Zhujiang Hospital. All patients underwent CRS-R assessments within one week prior to the study and one month afterward; the results are shown in Table 1. Inclusion criteria for patients with DOC were: (1) age over 16 years; (2) a diagnosis of chronic DOC; (3) clinical stability without neuromuscular blockers in the preceding 24 hours; and (4) signed informed consent from legal guardians. Exclusion criteria were: (1) severe or unstable medical conditions; (2) pre-injury CNS disabilities such as epilepsy; (3) a history of substance abuse; (4) severe facial injuries impacting expression; (5) pre-existing psychiatric conditions; or (6) use of sedatives, analgesics, barbiturates, antipsychotics, or antidepressants within 24 hours prior to the study. Healthy controls completed 18 trials, while patients with DOC completed 12 trials over two sessions. Each trial consisted of a 5-second auditory cue, followed by a 30-second emotional video stimulus, and a minimum 15-second rest interval. The research protocol adhered to the Declaration of Helsinki and received approval from the Zhujiang Hospital Ethics Committee (No: 2023-KY-174-01), with written informed consent obtained from the patients' legal guardians.

4.2. Data Preprocessing

The preprocessing pipeline began by applying an anti-aliasing low-pass filter to the original 1000 Hz data from all three datasets (SEED, SEED-IV, and self-recorded) to prevent signal distortion. The filtered signals were then downsampled to 200 Hz to allow for faster calculations. After downsampling, signals were band-pass filtered into five ranges: delta (14 Hz), theta (48 Hz), alpha (814 Hz), beta (1430 Hz), and gamma (3050 Hz). A linear dynamic system method was employed thereafter to effectively eliminate noise and artifacts, ensuring cleaner signals. The pipeline concluded with the application of the proposed channel normalization technique and the subsequent extraction of DE features, computed over consecutive, non-overlapping 1-second intervals.

4.3. Experiment I: Cross-Subject Emotion Recognition

To evaluate the cross-subject emotion recognition performance of our model, we utilized the SEED and SEED-IV datasets. Our approach involved employing leave-one-subject-out (LOSO) cross-validation. In each iteration of LOSO, data from one participant was designated as the test set, while the remaining 14 subjects were used for training. The results obtained from this evaluation are presented in Tables 2 and 3, which provide a comprehensive comparison with state-of-the-art (SOTA) methods. Our model demonstrated superior performance on both datasets, achieving mean accuracies of 95.37% ± 5.17% on SEED, and 82.50% ± 9.28% on SEED-IV. These results significantly outperformed recently proposed deep learning methods. Notably, our model exhibited robust generalization across most participants, achieving accuracies of 100% for subjects 2, 4, 7, and 9 on SEED, and accuracies above 90% for subjects 1, 2, and

Table 3
Model performance on SEED-IV using LOSO cross-validation.

Methods	P_{acc} (%)	F1 score	Methods	P_{acc} (%)	F1 score
DANN [34]	54.63 ± 8.03	54.48 ± 8.15	JTSR [26]	78.85 ± 3.47	78.71 ± 3.62
PLMSDANet [27]	73.08 ± 10.33	72.95 ± 10.18	MFA-LR [25]	74.99 ± 12.10	74.83 ± 11.95
ST-SCGNN [37]	76.37 ± 5.77	76.51 ± 5.65	MSGDA [11]	74.70 ± 7.20	74.58 ± 7.33
EEGMatch [38]	65.53 ± 8.31	65.68 ± 8.12	PR-PL [24]	81.32 ± 8.53	81.45 ± 8.42
MAS-DGAT-Net [39]	70.22 ± 9.12	70.36 ± 9.25	LGDAAN-Nets [40]	74.12 ± 08.45	75.12 ± 7.21
Ours	$\mathbf{82.50 \pm 9.28}$	$\mathbf{82.76 \pm 9.04}$			

Table 4
Performance and statistical comparison of different EEG-specific feature extractors on SEED.

Feature Extractor	P_{acc}	F1 score	P value
EEGNet	92.53±6.24	92.87±6.62	0.012
BiLSTM	93.12±5.91	93.32±5.87	0.043
DGCNN	93.45±5.35	93.15±5.19	0.047
TSCeption (Ours)	95.37±5.17	95.68±5.01	-

Note: P_{acc}: Overall correctness; F1: F1 Score. P values are computed using a two-tailed paired t-test between the proposed method and each baseline based on classification accuracy.

13 on SEED-IV. To provide a more intuitive demonstration of our model's generalization performance, we evaluated multiple metrics across several datasets, including the HBUED dataset [41] and the DEAP dataset [52] under cross-subject binary classification settings, as shown in Table 9. These findings underscore the model's strong generalization capability across individual subjects.

4.4. Experiment II: Cross-Subject Consciousness Detection

This experiment involved two distinct evaluations. First, to establish a baseline performance on healthy individuals, the model's emotion recognition accuracy was assessed using a LOSO cross-validation within the healthy control group. Second, for the consciousness detection task, a cross-domain setting was formulated. This setting utilized the 10 healthy individuals as the source domain and each of the 8 patients with DOC as an individual target domain. The model was trained on the data from all healthy subjects and subsequently tested on each patient to determine if emotion-specific neural patterns, indicative of residual consciousness, could be detected. Table 6 shows accuracies for healthy subjects ranging from 75% to 91.67%, with a mean of 82.64%. These results demonstrate that our HEAT model exhibits excellent performance and robust generalization in emotion recognition tasks among healthy subjects. For the cross-domain consciousness detection task, the results for the eight DOC patients are presented in Table 7. All patients performed above the chance level (50%), and four patients (P3, P4, P5, and P8) achieved a more stringent level of statistical significance ($p < 0.01$). These findings demonstrate the model's potential to identify emotion-specific neural patterns in a subset of DOC patients, offering objective evidence to supplement clinical consciousness assessments.

Building on this potential for consciousness assessment, we further investigated the relationship between emotion recognition accuracy and the best CRS-R score measured one month after the experiment in patients with DOC. As shown in Fig. 3, a Spearman rank correlation analysis revealed a significant positive association between recognition accuracy and CRS-R score ($\rho = 0.774$, $p = 0.024$, $n = 8$). This result suggests that patients with higher levels of consciousness tended to exhibit clearer emotion-discriminative neural responses, further indicating that HEAT may reflect graded differences in conscious processing across the DOC spectrum.

4.5. Experiment III: Ablation Study

To examine the contribution of each component, we performed ablation studies on SEED, removing Hierarchical Adversarial Alignment, fixed pseudo-labeling, and the progressive pseudo-labeling strategy in turn. The ablation results are presented in Table 8.

Table 5
Experimental hyperparameter settings.

Hyperparameters	Assigned Value
Segmentation Window Length	1 s
Batch Size	128
Total Epochs	500
Optimizer	RMSprop
Learning Rate	2.5e-4
Momentum	0.975
Weight Decay	3e-5
LR Scheduler Gamma	0.0015
LR Scheduler Decay Rate	0.6
Transfer Loss Weight (α)	0.5
Pseudo-label Loss Weight (β)	0.8
Temperature Coefficient (T)	1.8
Initial Threshold (τ_{init})	0.5
Maximum Threshold (τ_{max})	0.95
Adversarial Max Iters	10000
Temporal Filters (num_T)	4
Spatial Filters (num_S)	32
Hidden Dimension	64
Dropout Rate	0.6

Table 6
LOSO emotion recognition accuracy for healthy subjects.

Subjects	P_{acc} (%)
H1	91.67
H2	83.36
H3	75.00
H4	79.17
H5	78.49
H6	79.17
H7	87.50
H8	83.33
H9	85.42
H10	83.33
Mean	**82.64±4.87**

We tested four configurations to examine the contribution of each component. These configurations, denoted C1 through C4, are defined as follows: C1 is a baseline DANN model; C2 consists of our hierarchical alignment structure alone; C3 combines the hierarchical alignment with a fixed-threshold pseudo-labeling strategy; and C4 represents the full HEAT model, which incorporates Hierarchical Adversarial Alignment with the progressive pseudo-labeling strategy. As depicted in Table 8, the results of the ablation study validate the contribution of each component through progressive performance improvements. The transition from the C1 baseline (89.23% accuracy) to C2 (92.03% accuracy) demonstrates the significant benefit of our Hierarchical Adversarial Alignment architecture for improving cross-subject generalization. Building upon this, the inclusion of a fixed pseudo-labeling strategy in C3 further elevates the performance to 92.65%. Finally, the full HEAT model (C4), which employs the progressive, curriculum-guided

Table 7
Cross-domain emotion recognition accuracy for patients with DOC.

Patients	P_{acc} (%)	95% CI (%)	P value
P1	58.34	[48.68, 68.00]	0.047
P2	58.33	[48.67, 67.99]	0.048
P3	66.67	[57.43, 75.91]	<0.01
P4	75.00	[66.51, 83.49]	<0.01
P5	75.00	[66.51, 83.49]	<0.01
P6	59.36	[49.73, 68.99]	0.030
P7	58.33	[48.67, 67.99]	0.048
P8	70.84	[61.94, 79.74]	<0.01

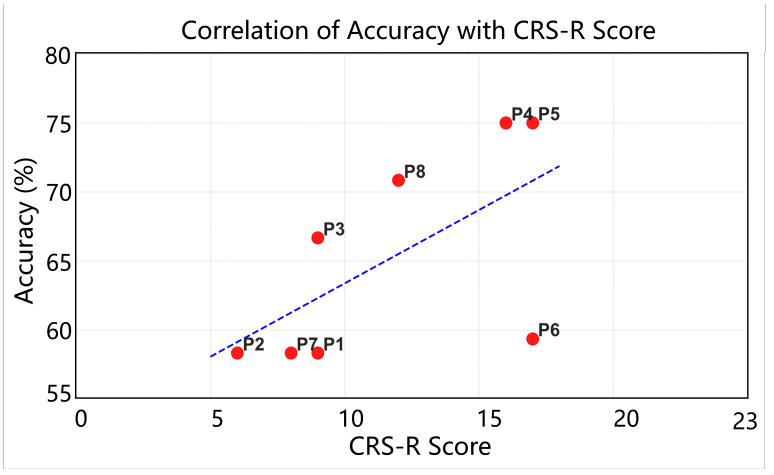

Figure 3: Correlation between the best CRS-R score measured one month after the experiment and emotion recognition accuracy in patients with DOC (Spearman rank correlation, $\rho = 0.774$, $p = 0.024$, $n = 8$).

Table 8
Ablation study results for different model configurations (C1-C4).

Datasets	Methods	P_{acc} (%)	AUC	P value
SEED	C1	89.23±6.42	0.9379	0.012
	C2	92.03±5.45	0.9636	0.037
	C3	92.65±7.13	0.9684	0.041
	C4 (Ours)	95.37±5.17	0.9704	-

pseudo-labeling strategy, achieves the highest accuracy of 95.37%. Moreover, the improvement from C3 to C4 is statistically significant ($p = 0.041$), indicating that the progressive strategy provides a practical advantage over fixed-threshold pseudo-labeling. This step-by-step enhancement confirms that each proposed component provides a marked and collective improvement to the model's overall emotion-recognition capabilities.

5. Discussion

In this work, we introduced HEAT, a novel framework designed to overcome cross-subject EEG variability by coupling hierarchical adversarial alignment with a curriculum-based progressive pseudo-labeling strategy. This approach preserves essential emotion-specific neural signatures, leading to state-of-the-art emotion recognition. In a clinical application, HEAT effectively identified residual consciousness in patients with DOC, with detection accuracy

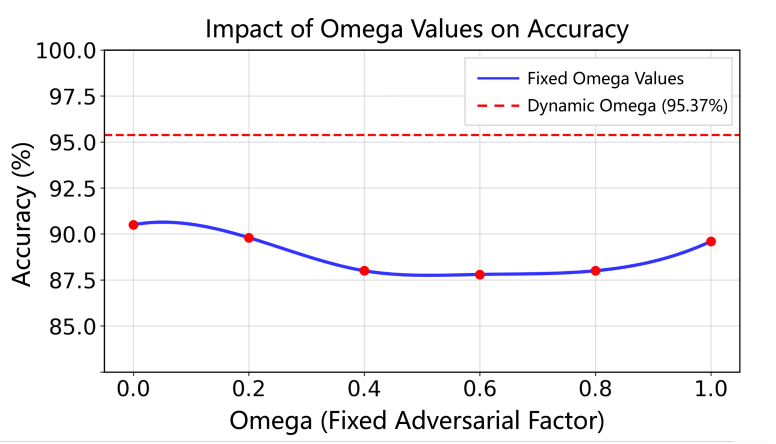

Figure 4: Effect of fixed vs. dynamic ω settings on model accuracy.

correlating with CRS-R scores. These findings, supported by neurophysiologically plausible brain activation patterns, establish HEAT as a robust tool for advancing label-efficient affective BCIs and translational neurodiagnostics.

5.1. Efficacy of Hierarchical Emotional Alignment

The effectiveness of the hierarchical alignment architecture in mitigating cross-subject variability is directly evidenced by our ablation studies, with results presented in Table 8. A clear performance improvement over a conventional DANN approach is demonstrated when comparing the results of configurations C1 and C2. The baseline DANN model (C1) achieved an accuracy of 89.23% on the SEED dataset. By incorporating our full hierarchical alignment structure (C2), the accuracy notably increased to 92.03%. This highlights our model's enhanced capability to learn domain-invariant features through hierarchical alignment, compared to methods that rely solely on global marginal distribution alignment.

The advantage of hierarchical alignment is particularly important in cross-subject EEG emotion recognition because the domain shift is not uniform across emotion categories. In this context, emotion-specific discrepancies refer to class-conditional differences in how neural responses associated with different emotions are organized across subjects. Beyond the shared inter-subject variation caused by individual physiology and recording conditions, different emotions can exhibit partially distinct spatial patterns and discriminative structures. Therefore, aligning only the marginal distribution may reduce the overall source-target gap while still mixing samples with different emotional meanings in the latent space. Such insufficient alignment can blur class boundaries, weaken emotion discriminability, and ultimately lead to negative transfer. By jointly constraining global domain invariance and local emotion-wise alignment, HEAT better preserves class-discriminative structure during transfer, which is consistent with the performance gain from C1 to C2 in Table 8 and the clearer cluster separation in Fig. 7.

Another important design choice is that the global discriminator and the local class-specific discriminators share the same basic architecture, although they operate on different feature scopes and optimization targets. We intentionally kept the discriminator backbone identical because both modules solve the same binary domain-classification problem, and using different architectures would introduce an additional model-capacity confound when comparing the effects of global and local alignment. Their functional difference therefore lies not in network complexity, but in the granularity of the alignment constraint: the global discriminator acts on the overall feature distribution, whereas each local discriminator receives emotion-weighted features and enforces class-conditional alignment. This controlled design allows the performance gain of HEAT to be more directly attributed to hierarchical alignment itself rather than to unequal discriminator capacity.

The advantage of our dynamic adversarial factor, ω, is clearly demonstrated by the comparative experiments illustrated in Fig. 4. As the results show, the dynamic weighting strategy achieved a peak accuracy of 95.37%, markedly outperforming all tested configurations that utilized fixed ω values, which yielded accuracies between approximately 87.8% and 90.6%. This performance gap highlights a key challenge in adversarial domain adaptation: the difficulty of pre-selecting a single, universally optimal fixed weight to balance global and local feature alignment. Manually tuning

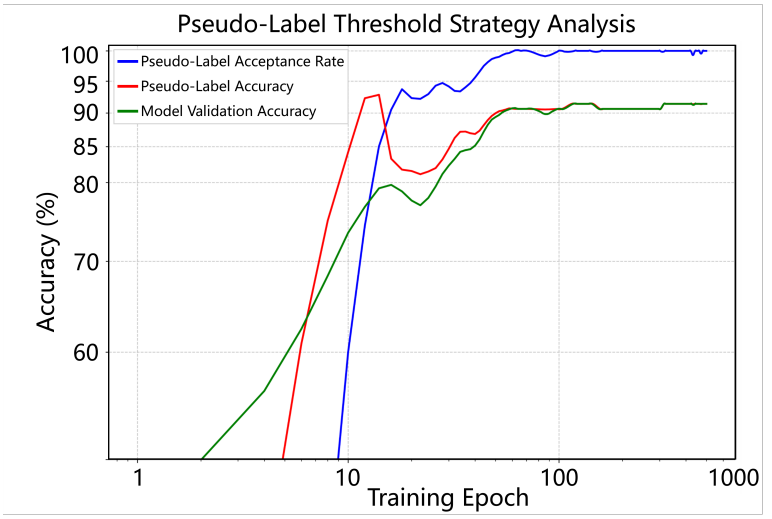

Figure 5: Evolution of pseudo-label metrics and model performance over training epochs.

this balance is often suboptimal, as the relative importance of different alignment objectives can change throughout the training process. Our adaptive mechanism, based on real-time $\mathcal{A}$-distance estimations, addresses this issue by automatically adjusting the focus between achieving robust global domain invariance and preserving vital emotion-specific features. This principle of dynamically weighting losses has proven effective in other complex learning scenarios[42, 34]. This adaptive mechanism, in turn, optimizes the alignment trade-off, leading to superior cross-subject generalization.

To verify the stability of the dynamic weighting mechanism, we visualized the trajectory of the adversarial factor ω during training on the cross-subject SEED dataset, as shown in Fig. 6. The factor first exhibits adaptive fluctuations while the model balances global and local alignment objectives, and then gradually converges to a stable region near 0.5. This behavior indicates that the data-driven weighting scheme does not introduce instability, but instead provides a controlled adjustment process that supports the final performance gain.

5.2. Progressive Pseudo-Labeling

To empirically validate this dynamic, curriculum-driven learning mechanism, we illustrate its operational training dynamics in Fig. 5. This visualization tracks three essential indicators: the pseudo-label acceptance rate, defined as the proportion of unlabeled data with predictions confident enough to serve as training targets; the pseudo-label accuracy, indicating the ratio of correctly assigned labels among the accepted ones; and the model validation accuracy evaluated on a separate test set.

Rather than using a fixed or abruptly changing confidence threshold, we adopt a linearly increasing schedule because the reliability of pseudo-labels in cross-subject EEG transfer improves progressively rather than discontinuously during training. This schedule provides a smooth transition from broader sample exploration in the early stage to stricter quality control in the later stage: it retains enough target samples to support adaptation before the decision boundary becomes stable, and then gradually suppresses noisy supervision as the model matures. As shown in Fig. 5, the pseudo-label acceptance rate continues to rise under this increasingly strict criterion, while pseudo-label accuracy remains high and the validation accuracy improves steadily. This coordinated trend suggests that the linear schedule is better matched to the model's gradual maturation than a fixed threshold, an overly aggressive schedule that would prematurely reject informative samples, or an overly slow schedule that would retain noisy pseudo-labels for too long.

These results further suggest that the performance gains of HEAT arise from the coordinated integration of hierarchical adversarial alignment and progressive pseudo-label refinement. Unlike prototypical representation based pairwise learning (PR-PL) [24], which reformulates emotion recognition through prototype-guided pairwise relations, HEAT retains a direct classification objective and employs hierarchical adversarial alignment to jointly preserve global domain invariance and local emotion-specific structure, with the adaptive weight ω dynamically balancing these two objectives. Likewise, unlike SFT-SGAT [20], which relies on a two-stage semi-supervised fine-tuning scheme with a

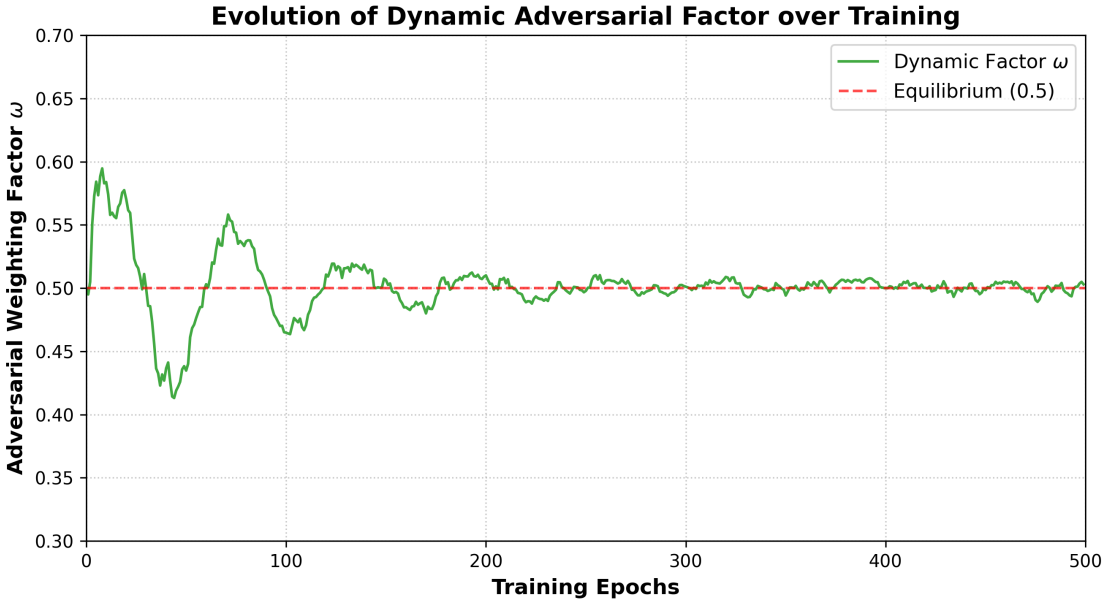

Figure 6: Evolutionary trajectory of the dynamic adversarial weighting factor ω during training in the cross-subject SEED experiment.

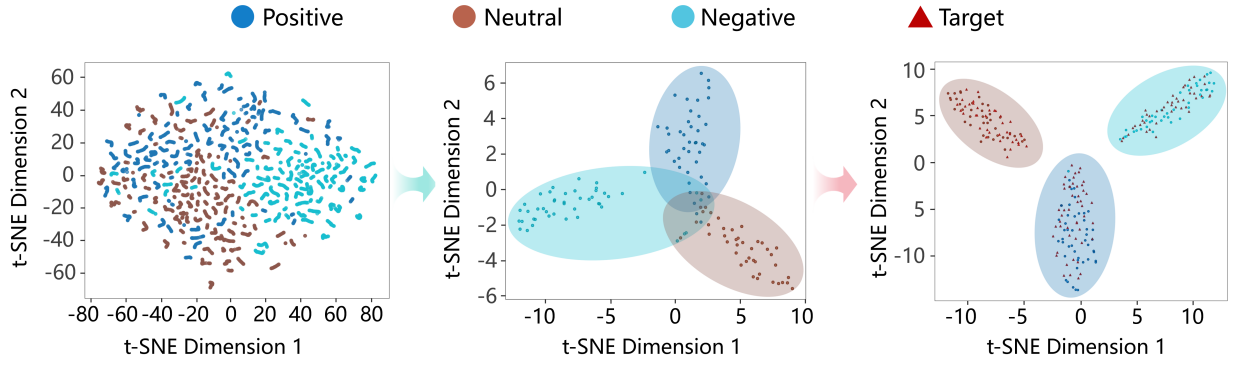

Figure 7: t-distributed stochastic neighbor embedding (t-SNE) visualization of the output logits from the classification layer at different training stages: (a) epoch 0, (b) epoch 50, and (c) epoch 500.

fixed confidence threshold, HEAT integrates hierarchical alignment and curriculum-based progressive pseudo-labeling into a unified end-to-end optimization process, allowing the pseudo-label selection criterion to become increasingly strict during training. More broadly, recent EEG-based emotion recognition studies have highlighted the value of richer spatio-temporal dependency modeling and more precise cross-domain adaptation. For example, STRFLNet emphasizes spatio-temporal representation fusion through dynamic-static graph modeling and hierarchical transformer-based feature integration [43], whereas FMLAN improves cross-subject and cross-session recognition through multi-source domain adaptation, mutual learning, and fine-grained alignment [44]. Taken together, these comparisons suggest that robust EEG emotion recognition benefits from both expressive representation learning and principled domain adaptation, while the effectiveness of HEAT mainly arises from the coordinated optimization between class-aware hierarchical alignment and progressive pseudo-label refinement.

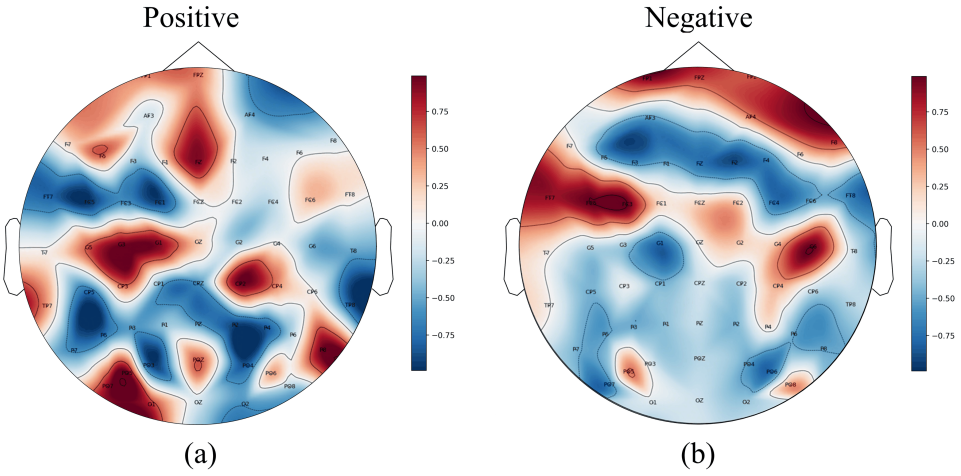

Figure 8: Brain activation topography of HEAT for (a) a positive emotion sample, and (b) a negative emotion sample, on the SEED dataset.

5.3. HEAT's Feature Extraction Capability

Fig. 7 presents t-SNE visualizations of the output logits from the classification layer at different training epochs. At epoch 0 (Fig. 7(a)), the representations exhibit substantial overlap among emotional categories, indicating that the model has not yet formed clear discriminative decision boundaries. By epoch 50 (Fig. 7(b)), an initial class-wise organization begins to emerge. After full training at epoch 500 (Fig. 7(c)), the representations become markedly better separated, showing that HEAT progressively organizes the learned decision space into more discriminative emotion-related clusters. This trend is consistent with the model's improved cross-subject alignment and class discrimination.

5.4. Interpreting Emotion-Specific Neural Patterns

To neurophysiologically investigate the features learned by HEAT, particularly its capability to capture emotion-specific information, we analyzed the importance of different EEG channels in the emotion recognition task. Channel importance was assessed using a Monte Carlo random subset masking method. The core of this approach involves evaluating the change in model performance when signals from subsets of randomly selected channels are masked, thereby inferring the contribution of each channel to the accurate prediction of emotional states.

The brain topography maps presented in Fig. 8 visually depict the importance scores of different channels identified by the HEAT model when processing positive and negative emotions. In these maps, warmer colors (e.g., red) indicate channels with higher contributions to the classification of the respective emotion, while cooler colors (e.g., blue) represent channels with lower contributions.

The positive emotion map, shown in Fig. 8(a), reveals a notable left-hemispheric pattern, with strong activations encompassing left frontal (e.g., Fp1, AF3, F3 and F1), medial frontal (e.g., FZ), and left fronto-central channels (e.g., FC1, FC3 and C3). Some involvement of left temporal (e.g., T7 and FT7) and left posterior parietal/occipital regions (e.g., P3, P7, PO3 and PO7) is also apparent. This pronounced left-lateralized frontal and central activity aligns well with theories of emotional hemispheric specialization, which posit greater involvement of the left prefrontal cortex in the processing of positive affect and approach motivation [45, 46]. The engagement of temporal regions may also reflect the semantic understanding and contextual associations of positive emotional stimuli.

In contrast, the brain topography map for negative emotions, shown in Fig. 8(b), exhibits distinct activation characteristics. Channels in the right hemisphere show higher importance, particularly in the right frontal (e.g., Fp2, AF4, F4, F2 and F6), right fronto-temporal (e.g., F8, FC6 and FT8), and right central/parietal regions (e.g., FC2, FC4, C4, CP2, CP4, P4, P6 and P8). This right-lateralized pattern is consistent with perspectives on brain asymmetries related to emotion, suggesting greater right prefrontal cortex involvement in processing negative emotions and withdrawal motivation [45, 46]. Notably, under negative emotional states, the temporal cortices also exhibit differential importance; the higher importance observed in the right temporal regions (e.g., T8 and FT8) could be associated with processing auditory or visual emotional stimuli with negative valence, or with the retrieval of unpleasant experiential

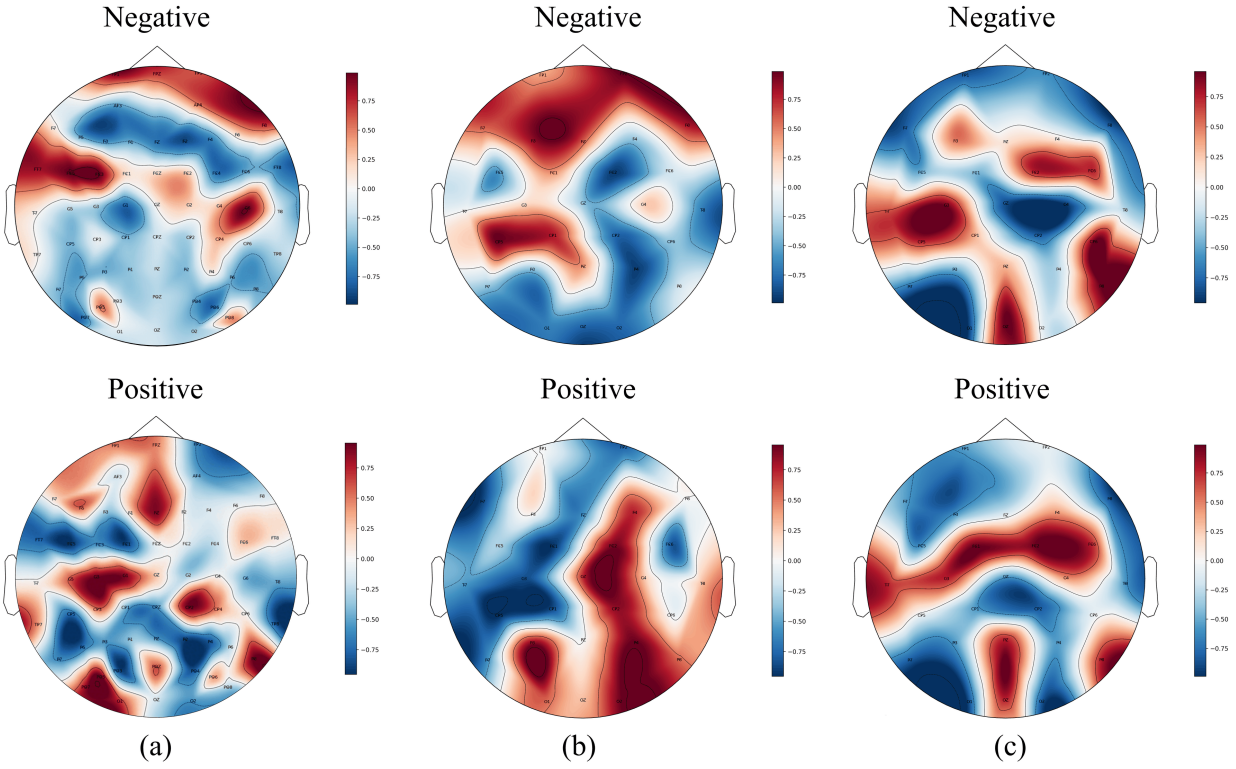

Figure 9: HEAT-derived brain activation topographies showing a graded attenuation of emotion-specific neural patterns for positive (bottom row) and negative (top row) states across (a) HC, (b) MCS, and (c) VS.

memories, while the left temporal lobe appears relatively less active. This offers a potential perspective for exploring the unique contributions of specific brain regions during negative emotions.

Crucially, positive and negative emotions exhibit distinct importance patterns in the brain topography maps generated from the HEAT model. These differentiated topographies indicate that HEAT not only learns general neural activity features related to emotional arousal but also successfully extracts and utilizes neural representations that are emotion-specific. This capability is closely linked to its model design, which emphasizes the preservation and alignment of local, class-specific features (e.g., through hierarchical adversarial alignment). These neurophysiological observations provide valuable evidence for the efficacy and interpretability of HEAT's learned features, and suggest underlying brain functional network patterns it relies upon to distinguish between different emotional states. These differentiated topographies provide a neurophysiological interpretation of the class-conditional discrepancies discussed above: cross-subject adaptation in EEG emotion recognition should not only remove subject-related nuisance variation, but also preserve the emotion-dependent spatial organization of neural responses.

5.5. Objective Consciousness Assessment in Patients with DOC

HEAT's core contribution to DOC assessment lies in its capacity for objective evaluation. By analyzing a patient's emotion-specific neural representations, the framework moves beyond behavioral observation to directly probe underlying cognitive function. Our results substantiate the principle that the clarity of these neural features correlates with the severity of consciousness impairment, offering a supplementary diagnostic tool that can reveal evidence of consciousness in patients whose awareness may not be captured by the CRS-R.

The objectivity and feasibility of our model for consciousness assessment are grounded in the intimate neurobiological relationship between emotion and consciousness [47, 48, 49], as the capacity to process emotional stimuli is considered a key facet of conscious awareness [50, 51]. Therefore, the detection of reliable emotional activity in patients can provide powerful, objective evidence of residual consciousness [20, 19]. This stands in contrast to standard clinical

Table 9
Performance of the proposed HEAT framework on SEED, SEED-IV, DEAP, and HBUED datasets

Datasets	P_{acc} (%)	F1 score	Recall	Precision	AUC
SEED	95.37±5.17	95.68±5.01	95.43±5.07	95.86±5.21	96.35±5.71
SEED-IV	82.50±9.28	82.76±9.04	82.46±9.42	82.79±9.78	83.34±9.36
DEAP (Valence)	72.63±9.32	72.29±9.77	72.21±9.24	72.58±9.25	72.19±9.34
DEAP (Arousal)	72.35±9.34	72.24±9.12	72.88±9.03	72.34±9.21	73.13±9.68
HBUED	84.21±05.25	84.34±05.12	84.43±05.21	84.32±05.23	84.56±05.32

Table 10
Parameters and inference times of the proposed and baseline models under the cross-subject SEED and SEED-IV settings.

Methods	Parameters (M)	SEED	SEED-IV
		Inference time (ms)	Inference time (ms)
DANN [34]	0.15	3.245	3.120
JTSR [26]	0.25	3.856	3.642
ST-SCGNN [37]	0.35	8.241	8.015
MFA-LR [25]	0.38	4.125	3.958
PLMSDANet [27]	0.42	4.512	4.320
MSGDA [11]	0.46	5.104	4.895
EEGMatch [38]	0.48	5.324	5.102
PR-PL [24]	0.51	6.616	5.598
SSCDG [21]	0.55	6.153	–
SFT-SGAT [20]	0.58	7.882	–
MAS-DGAT-Net [39]	0.62	7.452	7.214
LGDAAN-Nets [40]	0.65	8.153	7.846
Ours	**0.45**	**4.932**	**4.768**

evaluations, such as the CRS-R, which assess consciousness by interpreting a patient's observable behavioral responses. Such assessments are inherently limited and prone to misdiagnosis, as patients with severe motor impairments may retain awareness but be unable to demonstrate it. By analyzing involuntary neural responses, our EEG-based paradigm bypasses the need for motor output, offering a more direct and objective window into a patient's preserved cognitive functions.

As presented in Section 4.4, the model demonstrated its efficacy in differentiating among HC, MCS, and VS patients. This efficacy was evidenced by significant emotion recognition in patients with DOC, and the accuracy scores showed a clear correlation with clinical CRS-R scores. The EEG topographic maps in Fig. 9 provide further neurophysiological substantiation for this framework. They reveal a discernible, graded attenuation of emotion-specific neural patterns progressing from HC to MCS, and further to VS patients. In essence, HEAT provides an objective and neurophysiologically-grounded tool that can identify evidence of consciousness in patients whose awareness is not captured by the CRS-R, thereby serving as a crucial supplementary tool for clinicians.

5.6. Model Complexity and Inference Efficiency

For a more practical perspective, we further compare the parameter counts and inference times of HEAT and representative baseline methods under the cross-subject SEED and SEED-IV settings, as summarized in Table 10. Although HEAT is not the smallest model, it maintains a compact parameter scale and competitive inference speed while achieving the best recognition performance, indicating a favorable trade-off between effectiveness and computational cost.

5.7. Parameter Sensitivity Analysis

To investigate the impact of core hyperparameters on model performance and to balance the optimization objectives of each module, we conducted a comprehensive sensitivity analysis on the SEED dataset. As shown in Fig. 10, we tested different values for the transfer loss weight α, pseudo-label loss weight β, initial threshold τ_{init}, maximum threshold

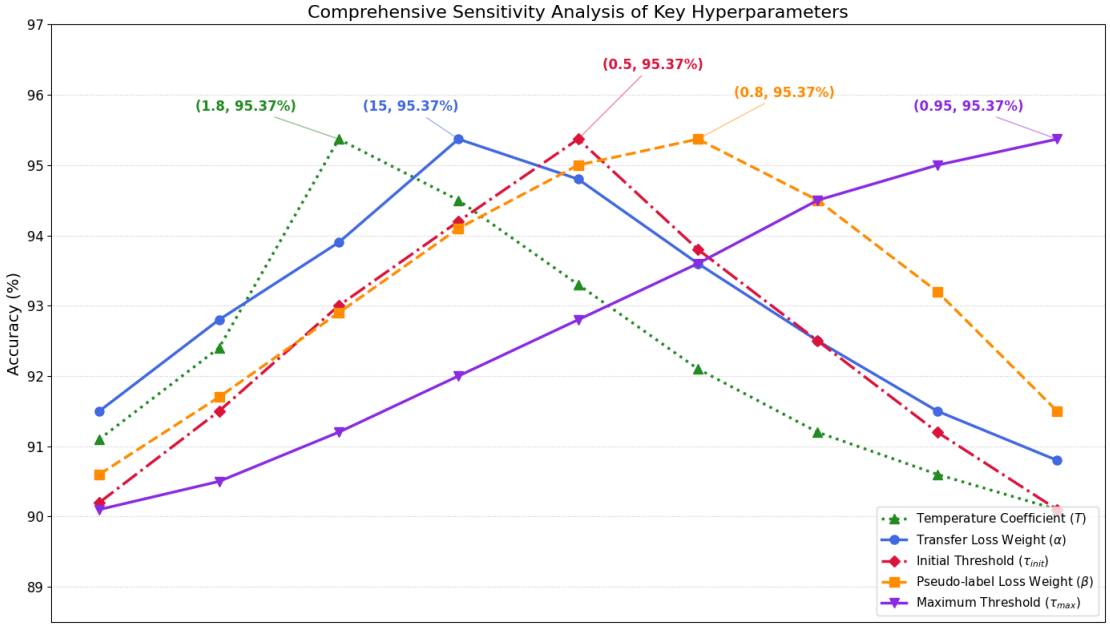

Figure 10: Comprehensive sensitivity analysis of key hyperparameters on the SEED dataset. Note that the horizontal axis represents the increasing search sequence for each parameter rather than an absolute numerical scale.

τ_{max}, and temperature coefficient T, respectively. As illustrated, with the adjustment of each parameter value, the model performance exhibits a trend of initially increasing and then decreasing, ultimately peaking at an accuracy of 95.37% in all cases.

We selected the optimal parameters based on the validation set, specifically setting $\alpha = 0.5$ and $\beta = 0.8$. In this section, we conduct a parameter sensitivity analysis. This analysis further verifies that our selected parameters achieve an effective performance balance. Specifically, it confirms that a relatively large value for α ensures that the hierarchical adversarial alignment mechanism effectively steers the model toward learning domain-invariant representations. Meanwhile, it demonstrates that our chosen setting for β not only fully harnesses the benefits of the progressive pseudo-labeling strategy in leveraging unlabeled target data but also effectively controls this component's sensitivity to potential noise in the pseudo-labels. Regarding the pseudo-label confidence threshold, τ_{init} achieves the best result at 0.5, a value that strikes a reasonable compromise between generating a sufficient volume of pseudo-labels and maintaining their reliability during the model's early exploratory stage. Simultaneously, τ_{max} is set to 0.95 to impose rigorous quality control, ensuring that only highly confident predictions guide the fine-tuning process in the advanced stages of training. Finally, the temperature coefficient T reaches its optimum at 1.8. This calibration mechanism, by smoothing the probability distribution, provides a reliable confidence basis for the progressive thresholding strategy, thereby ensuring the quality and stability of pseudo-label selection. When these hyperparameters are too large or too small, model performance degrades due to imbalanced feature learning or error accumulation. Therefore, we adopt this set of parameters in all experiments to achieve the optimal performance trade-off.

5.8. Limitations and Future Work

Despite its promising performance, HEAT has certain limitations. First, the progressive pseudo-labeling module still depends on sufficiently reliable predictions at the early stage of training. When the source-target discrepancy is substantial or the target data are particularly noisy, pseudo-label errors may still accumulate and reduce the benefit of curriculum refinement. Second, the present validation is mainly based on the SEED and SEED-IV datasets and a relatively small DOC cohort, so the generalizability of the framework to broader clinical populations and more heterogeneous recording conditions remains to be established. In addition, relying solely on EEG signals may limit the comprehensive characterization of emotional and conscious states. Future work will therefore focus on improving robustness under harder transfer settings, validating the framework in broader cohorts, and extending it with multimodal information.

6. Conclusion

In this paper, we propose HEAT, a novel framework designed to enhance cross-subject EEG emotion recognition and consciousness detection. HEAT establishes a cohesive workflow where a hierarchical adversarial alignment module first learns domain-invariant yet class-discriminative representations by dynamically balancing global and local feature adaptation. Simultaneously, a progressive pseudo-labeling strategy effectively leverages unlabeled data through a curriculum-based adaptive confidence threshold and probability calibration. These components are optimized jointly through a composite loss function to perform robust cross-domain classification, ensuring that emotion-specific neural patterns are preserved throughout the transfer process. Our experiments demonstrate that HEAT achieves state-of-the-art performance in cross-subject emotion recognition on the SEED and SEED-IV datasets. Furthermore, HEAT shows great potential for objective consciousness assessment in patients with DOC by analyzing emotion-specific neural representations.

Data availability

The source code for HEAT is publicly available at `https://github.com/Peregrine123/HEAT`. The SEED and SEED-IV datasets are available from their original providers. The self-collected DOC dataset is available from the corresponding author on reasonable request, subject to ethical and privacy restrictions.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Zhepei Hong: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing original draft, Writing review & editing, Visualization; **Rongtao Chen:** Software, Validation, Investigation, Writing review & editing; **Liting Li:** Investigation, Writing review & editing, Visualization; **Jiajun Chen:** Investigation, Writing review & editing; **Wei Gao:** Writing review & editing, Supervision, Investigation; **Jiahui Pan:** Conceptualization, Investigation, Resources, Data curation, Writing review & editing, Supervision, Project administration, Funding acquisition.

Acknowledgments

This work was supported in part by the Major Projects of Colleges and Universities in Guangdong Province (Grant No. 2023ZDZX2021) and the Guangdong Basic, Applied Basic Research Foundation (Grant No. 2024A1515010524), and the National Natural Science Foundation of China (Grant No. 62576142).

References

[1] Z. He, Y. Zhong, J. Pan, An adversarial discriminative temporal convolutional network for EEG-based cross-domain emotion recognition, Comput. Biol. Med. 141 (2022) 105048.

[2] D. Nath, Anubhav, M. Singh, et al., A comparative study of subject-dependent and subject-independent strategies for EEG-based emotion recognition using LSTM network, in: Proc. 4th Int. Conf. Compute Data Anal., 2020, pp. 142–147.

[3] W.-L. Zheng, B.-L. Lu, Investigating critical frequency bands and channels for EEG-based emotion recognition with deep neural networks, IEEE Trans. Auton. Mental Dev. 7 (2015) 162–175.

[4] N.-D. Mai, B.-G. Lee, W.-Y. Chung, Affective computing on machine learning-based emotion recognition using a self-made EEG device, Sensors 21 (2021) 5135.

[5] Y. Wang, Q. Li, J. Jia, et al., A novel transfer learning model for cross-subject emotion recognition using EEGs, in: Proc. 6th Int. Conf. Comput. Sci. Artif. Intell., 2022, pp. 217–223.

[6] L.-M. Zhao, X. Yan, B.-L. Lu, Plug-and-play domain adaptation for cross-subject EEG-based emotion recognition, in: Proc. AAAI Conf. Artif. Intell., volume 35, 2021, pp. 863–870.

[7] Y. Ganin, E. Ustinova, H. Ajakan, et al., Domain-adversarial training of neural networks, J. Mach. Learn. Res. 17 (2016) 1–35.

[8] J. Li, S. Qiu, C. Du, et al., Domain adaptation for EEG emotion recognition based on latent representation similarity, IEEE Trans. Cogn. Devel. Syst. 12 (2019) 344–353.

[9] Y. Zhu, Y. Ma, M. Meng, et al., Application of a multi-source multi-task weight adaptation framework for cross-domain EEG emotion recognition (MS-MWA), Signal Image Video Process. 18 (2024) 9105–9118.
[10] W.-L. Zheng, et al., EmotionMeter: A multimodal framework for recognizing human emotions, IEEE Trans. Cybern. 49 (2018) 1110–1122.
[11] J. Wang, X. Ning, W. Xu, et al., Multi-source selective graph domain adaptation network for cross-subject EEG emotion recognition, Neural Netw. 180 (2024) 106742.
[12] X.-C. Zhong, Q. Wang, R. Li, et al., Unsupervised domain adaptation with pseudo-label propagation for cross-domain EEG emotion recognition, IEEE Trans. Instrum. Meas. (2025) 1–11.
[13] W. Li, L. Fan, S. Shao, et al., Generalized contrastive partial label learning for cross-subject EEG-based emotion recognition, IEEE Trans. Instrum. Meas. (2024) 1–11.
[14] G. Zhang, V. Davoodnia, A. Etemad, PARSE: Pairwise alignment of representations in semi-supervised EEG learning for emotion recognition, IEEE Trans. Affect. Comput. 13 (2022) 2185–2200.
[15] K. Sohn, D. Berthelot, N. Carlini, et al., Fixmatch: Simplifying semi-supervised learning with consistency and confidence, Adv. Neural Inf. Process. Syst. 33 (2020) 596–608.
[16] J. Du, L. Yang, H. He, et al., Cross-subject emotion classification with residual pseudo-label distance-aware dual-classifier, in: Proc. IEEE Int. Conf. Bioinf. Biomed. (BIBM), IEEE, 2024, pp. 5683–5690.
[17] D. Cruse, S. Chennu, C. Chatelle, T. A. Bekinschtein, D. Fernández-Espejo, J. D. Pickard, S. Laureys, A. M. Owen, Bedside detection of awareness in the vegetative state: a cohort study, Lancet 378 (2011) 2088–2094.
[18] D. Lulé, Q. Noirhomme, S. C. Kleih, C. Chatelle, S. Halder, A. Demertzi, M.-A. Bruno, O. Gosseries, A. Vanhaudenhuyse, C. Schnakers, et al., Probing command following in patients with disorders of consciousness using a brain–computer interface, Clin. Neurophysiol. 124 (2013) 101–106.
[19] H. Huang, Q. Xie, J. Pan, Y. He, Z. Wen, R. Yu, Y. Li, An EEG-based brain computer interface for emotion recognition and its application in patients with disorder of consciousness, IEEE Trans. Affect. Comput. 12 (2019) 832–842.
[20] L. Qiu, L. Zhong, J. Li, et al., SFT-SGAT: A semi-supervised fine-tuning self-supervised graph attention network for emotion recognition and consciousness detection, Neural Netw. 180 (2024) 106643.
[21] H. Cai, J. Pan, Q. Xiao, et al., Decoding musical neural activity in patients with disorders of consciousness through self-supervised contrastive domain generalization, IEEE Trans. Affective Comput. (2024) 726–743.
[22] H. Li, Y.-M. Jin, W.-L. Zheng, et al., Cross-subject emotion recognition using deep adaptation networks, in: Proc. 25th Int. Conf. Neural Information Processing (ICONIP), Springer, 2018, pp. 403–413.
[23] M. N. Imtiaz, N. Khan, Enhanced cross-dataset electroencephalogram-based emotion recognition using unsupervised domain adaptation, Comput. Biol. Med. 184 (2025) 109394.
[24] R. Zhou, Z. Zhang, H. Fu, et al., PR-PL: A novel prototypical representation based pairwise learning framework for emotion recognition using EEG signals, IEEE Trans. Affective Comput. 15 (2023) 657–670.
[25] M. Jiménez-Guarneros, G. Fuentes-Pineda, Learning a robust unified domain adaptation framework for cross-subject EEG-based emotion recognition, Biomed. Signal Process. Control 86 (2023) 105138.
[26] Y. Peng, H. Liu, W. Kong, et al., Joint EEG feature transfer and semisupervised cross-subject emotion recognition, IEEE Trans. Ind. Inform. 19 (2022) 8104–8115.
[27] C. Ren, J. Chen, R. Li, et al., Semi-supervised pairwise transfer learning based on multi-source domain adaptation: A case study on EEG-based emotion recognition, Knowl.-Based Syst. 305 (2024) 112669.
[28] M. Jiménez-Guarneros, G. Fuentes-Pineda, Cross-subject EEG-based emotion recognition via semisupervised multisource joint distribution adaptation, IEEE Trans. Instrum. Meas. 72 (2023) 1–11.
[29] C. Schnakers, A. Vanhaudenhuyse, J. Giacino, M. Ventura, M. Boly, S. Majerus, G. Moonen, S. Laureys, Diagnostic accuracy of the vegetative and minimally conscious state: clinical consensus versus standardized neurobehavioral assessment, BMC Neurol. 9 (2009) 35.
[30] Q. Noirhomme, D. Lesenfants, R. Lehembre, Z. Lugo, C. Chatelle, A. Vanhaudenhuyse, S. Laureys, Detecting consciousness with a brain-computer interface, in: Converging Clinical and Engineering Research on Neurorehabilitation, Springer, 2013, pp. 1261–1264.
[31] M. Lee, L. R. Sanz, A. Barra, A. Wolff, J. O. Nieminen, M. Boly, M. Rosanova, S. Casarotto, O. Bodart, J. Annen, et al., Quantifying arousal and awareness in altered states of consciousness using interpretable deep learning, Nat. Commun. 13 (2022) 1064.
[32] R.-N. Duan, J.-Y. Zhu, B.-L. Lu, Differential entropy feature for EEG-based emotion classification, in: Proc. 6th Int. IEEE/EMBS Conf. Neural Eng. (NER), IEEE, 2013, pp. 81–84.
[33] Y. Ding, N. Robinson, S. Zhang, et al., TSception: Capturing temporal dynamics and spatial asymmetry from EEG for emotion recognition, IEEE Trans. Affect. Comput. 14 (2022) 2238–2250.
[34] C. Yu, J. Wang, Y. Chen, et al., Transfer learning with dynamic adversarial adaptation network, in: Proc. IEEE Int. Conf. Data Mining (ICDM), IEEE, 2019, pp. 778–786.
[35] Y. Bengio, J. Louradour, R. Collobert, et al., Curriculum learning, in: Proc. 26th Annu. Int. Conf. Mach. Learn., 2009, pp. 41–48.
[36] D.-H. Lee, et al., Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks, in: Workshop on challenges in representation learning, ICML, volume 3, Atlanta, 2013, p. 896.
[37] J. Pan, R. Liang, Z. He, et al., ST-SCGNN: a spatio-temporal self-constructing graph neural network for cross-subject EEG-based emotion recognition and consciousness detection, IEEE J. Biomed. Health Inform. 28 (2023) 777–788.
[38] R. Zhou, W. Ye, Z. Zhang, et al., EEGMatch: Learning with incomplete labels for semisupervised EEG-based cross-subject emotion recognition, IEEE Trans. Neural Netw. Learn. Syst. (2024) 12991–13005.
[39] S. Liu, X. Wang, M. Jiang, Y. An, Z. Gu, B. Li, Y. Zhang, MAS-DGAT-Net: A dynamic graph attention network with multibranch feature extraction and staged fusion for EEG emotion recognition, Knowledge-Based Systems 305 (2024) 112599.
[40] Y. An, S. Hu, S. Liu, X. Wang, Z. Gu, Y. Zhang, LGDAAN-Nets: A local and global domain adversarial attention neural networks for EEG emotion recognition, Knowledge-Based Systems 318 (2025) 113613.

[41] S. Liu, X. Wang, Y. An, Z. Wang, Z. Gu, Y. Zhang, S. Zhao, HBUED: An EEG dataset for emotion recognition, J. Affect. Disord. 385 (2025) 119397.
[42] Y. Wang, H. Chen, Q. Heng, et al., Freematch: Self-adaptive thresholding for semi-supervised learning, arXiv preprint arXiv:2205.07246 (2022).
[43] F. Hu, K. He, C. Wang, Q. Zheng, B. Zhou, G. Li, Y. Sun, STRFLNet: Spatio-temporal representation fusion learning network for EEG-based emotion recognition, IEEE Trans. Affect. Comput. 17 (2026) 204–218.
[44] P. Yu, X. He, H. Li, H. Dou, Y. Tan, H. Wu, B. Chen, FMLAN: A novel framework for cross-subject and cross-session EEG emotion recognition, Biomed. Signal Process. Control 100 (2025) 106912.
[45] R. J. Davidson, 11 affect, cognition, and hemispheric specialization, in: Emotions, cognition, and behavior, Cambridge Univ. Press, 1984, pp. 320–365.
[46] R. J. Davidson, A. J. Shackman, J. S. Maxwell, Asymmetries in face and brain related to emotion, Trends Cogn. Sci. 8 (2004) 389–391.
[47] H. Brace, The feeling of what happens: Body and emotions in the making of consciousness, Psychiatr. Serv. 51 (2000) 1579–1579.
[48] A. Damasio, Feelings of emotion and the self, Ann. N. Y. Acad. Sci. 1001 (2003) 253–261.
[49] N. Tsuchiya, R. Adolphs, Emotion and consciousness, Trends Cogn. Sci. 11 (2007) 158–167.
[50] M. Balconi, C. Lucchiari, Consciousness and emotional facial expression recognition: Subliminal/supraliminal stimulation effect on n200 and p300 erps, J. Psychophysiol. 21 (2007) 100–108.
[51] J. Megill, Emotion, cognition and artificial intelligence, Minds Mach. 24 (2014) 189–199.
[52] S. Koelstra, C. Mühl, M. Soleymani, J.-S. Lee, A. Yazdani, T. Ebrahimi, T. Pun, A. Nijholt, I. Patras, DEAP: A database for emotion analysis using physiological signals, IEEE Trans. Affect. Comput. 3 (2012) 18–31.