EEG-PIEDA: REINFORCEMENT LEARNING FOR EEG-DRIVEN PERSONALIZED VIDEO INTERVENTION WITH EMOTIONAL DYNAMICS AWARENESS

Anonymous authors
Paper under double-blind review

ABSTRACT

Achieving personalized and adaptive mental stress regulation in dynamic environments remains challenging due to substantial variations in individual emotional responses and the temporal dependency of affective states. Although video-based interventions provide an accessible approach for stress alleviation, existing methods mainly rely on behavioral feedback and cannot characterize individual neural responses or the stability of emotional regulation processes. To address these challenges, we propose EEG-PIEDA, an EEG-guided personalized video regulation framework that integrates neural response pattern modeling, continuous emotion perception, and reinforcement learning. First, we analyze dynamic EEG differential entropy and functional connectivity variations to discover individual emotional response patterns during video stimulation. Then, we introduce a pattern-aware dual-branch affective perception network that adaptively fuses instantaneous EEG representations with individual response characteristics for continuous emotional state estimation. Furthermore, we formulate emotional regulation as a reinforcement learning problem and model emotional state transitions under different perturbation intensities to quantify individual emotional anti-interference capability. Based on the learned policy, EEG-PIEDA enables personalized video selection according to users' neural response dynamics and regulatory characteristics. Experiments on real-world EEG datasets demonstrate that our framework accurately identifies heterogeneous emotional response patterns, effectively captures dynamic emotional evolution, and achieves superior personalized stress regulation performance compared with existing approaches.

1 INTRODUCTION

In high-intensity study and work environments, stress has become an increasingly important factor affecting cognitive performance and psychological well-being (Kuppens et al., 2010). Prolonged stress not only impairs attentional control and executive functions (Kozakevich Arbel et al., 2024; Girotti et al., 2024), but also increases emotional instability and reduces daily life quality. Recent psychological studies have demonstrated that empathy-based interventions can effectively alleviate stress by enhancing emotional resonance, social connection, and cognitive reappraisal (Lattie et al., 2022). Accordingly, empathic videos have emerged as an accessible and non-invasive approach for stress regulation (Stieger et al., 2023). However, the effectiveness of empathic video intervention varies substantially across individuals (Kannadasan et al., 2026). The same video stimulus may induce significant emotional recovery in some users, while producing limited or unstable responses in others. This phenomenon suggests that the key challenge of video-based stress regulation is not merely selecting effective content, but understanding how individuals dynamically respond to the same stimulus.

Current video recommendation systems mainly optimize content selection based on behavioral feedback, including click-through rate, dwell time, and completion rate (Lubos et al., 2023). However, these signals mainly reflect user preferences rather than actual stress recovery states, making it difficult to characterize individual regulatory responses during intervention. Meanwhile, emotional states exhibit temporal dependency, where current affect is influenced by preceding experiences, a phenomenon known as emotional inertia (Kuppens et al., 2010). This property suggests that emo-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

tional regulation is not an instantaneous transition induced by external stimuli, but a dynamic process constrained by previous emotional trajectories. During intervention, individuals with different emotional dynamics may show distinct abilities to maintain the desired regulatory direction when their states are affected by internal fluctuations or external perturbations. Therefore, modeling the stability of emotional evolution is essential for understanding individual regulatory capability and achieving personalized intervention.

Electroencephalography (EEG), with high temporal resolution and sensitivity to neural fluctuations, provides an effective approach for characterizing individual-specific neural response dynamics during emotional regulation (Wu et al., 2025). EEG enables continuous observation of how users' neural states evolve under video stimulation, thereby revealing differences in emotional response trajectories, regulatory stability, and susceptibility to state fluctuations among individuals. Existing EEG-based affective computing studies mainly focus on static tasks such as emotion classification and stress recognition, while limited attention has been devoted to modeling the temporal evolution of neural responses during video-based regulation and investigating how individual neural response characteristics influence regulatory outcomes. Therefore, identifying personalized neural response patterns from EEG signals and incorporating them into adaptive stress regulation remains a challenge.

To address this challenge, we propose an EEG-guided personalized video regulation framework that integrates neural response pattern modeling and reinforcement learning. First, we analyze dynamic EEG Differential Entropy (DE) and Functional Connectivity (FC) variations during empathic video viewing to characterize individual response patterns and categorize users into three neural response types. Second, we introduce an EEG emotion perception model that combines instantaneous neural states with individual response characteristics for reliable emotional state estimation. Finally, reinforcement learning is employed to model emotional state transitions under different video interventions and quantify users' emotional anti-interference capability, enabling adaptive regulation strategies tailored to individual characteristics. The main contributions of this paper are summarized as follows:

- **Individual emotional response pattern discovery:** A neural response pattern discovery framework is developed to characterize individual differences under identical video stimulation. Temporal DE and FC features are modeled to identify distinct EEG response patterns through unsupervised analysis.
- **Personalized EEG affective state perception:** A pattern-aware EEG affective perception framework is proposed to fuse EEG features and individual response features for continuous emotional state estimation.
- **EEG-driven personalized video recommendation:** A neural state-aware recommendation strategy is established by modeling emotional resistance under dynamic perturbations, enabling video suitability evaluation and adaptive selection based on emotional dynamics.

2 Related Work

EEG-based Emotion and Stress Recognition. Machine learning algorithms have provided important approaches for analyzing human emotional and stress-related states. Tan et al. (2025) and Zhao et al. (2026) extracted EEG signal features, including DE and FC, respectively, as important representations for emotion and stress recognition. With the development of deep learning techniques, Wang et al. (2026) and Pi et al. (2026) further employed neural networks to automatically learn spatiotemporal representations from multi-channel EEG signals, thereby improving recognition accuracy. However, these methods mainly focus on the mapping relationship between EEG and categories, while paying limited attention to the dynamic evolution of emotional responses during continuous stimulation.

Video-based Emotion Regulation and Recommendation. Conventional recommendation methods mainly construct user preference models based on users' historical interactions, contextual information and video content, and leverage deep learning or reinforcement learning to optimize the matching relationship between users and videos (Ren et al., 2025). Nevertheless, such approaches are optimized predominantly according to user behavioral feedback and content relevance. Behavioral feedback mostly reflects users' selection preferences, yet cannot directly measure the actual

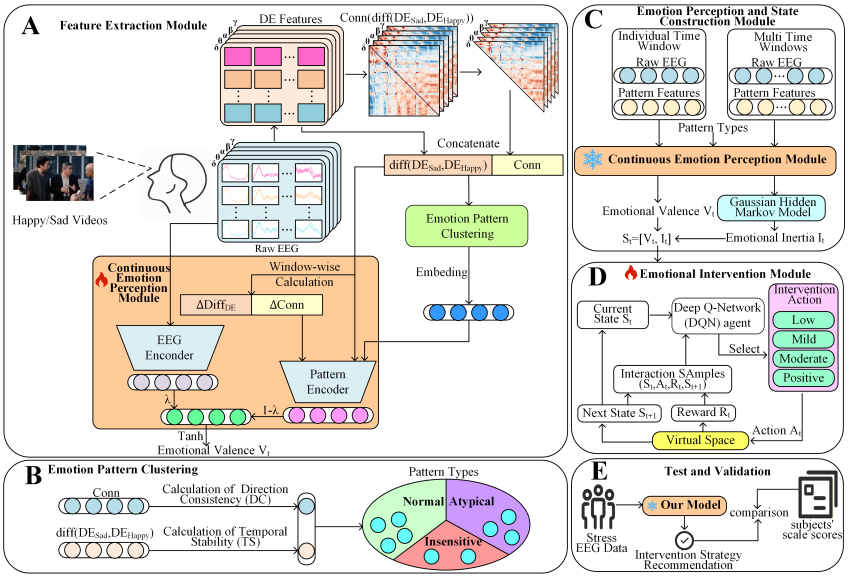

Figure 1: Overall architecture of the proposed EEG-PIEDA framework for personalized emotion regulation.

regulatory effect exerted by videos on individuals' emotional states. To address the feedback bias issue in the recommendation process, Wang et al. (2025) proposed the Wasserstein Balancing Metric. This method reduces observational bias via a causal balancing strategy and enhances the reliability of recommendation models. Even so, existing methods pay insufficient attention to dynamic emotional fluctuations during user regulation and fail to uncover the correlation between video content and the individual emotional regulation outcomes.

Reinforcement Learning for Adaptive Recommendation. Reinforcement learning has been widely adopted for dynamic policy learning in recommendation systems by virtue of its strengths in sequential decision-making. Zhang et al. (2025) proposed a large language model-driven user simulator to boost the training efficiency of recommendation systems. Li et al. (2025) put forward GraphDRL, which effectively models complex user-item interactions and alleviates the data sparsity issue. Nevertheless, these approaches generally take external behavioral feedback from users as their primary optimization objective and fail to characterize changes in users' internal states.

3 METHODS

3.1 EEG-BASED EMOTION REPRESENTATION LEARNING

The overall architecture of the proposed framework is illustrated in Figure 1. The EEG signal recorded during video viewing is denoted as $\boldsymbol{X} \in \mathbb{R}^{C\times L}$, where C and L represent the number of channels and sampling points, respectively. To capture dynamic emotional changes, the continuous EEG signal is segmented into T non-overlapping windows $\{\boldsymbol{X}_1, \ldots, \boldsymbol{X}_T\}$, where each window represents a brain state. To consider the frequency features of EEG oscillations, each window is further decomposed into five frequency bands $\mathcal{B} = \{\delta, \theta, \alpha, \beta, \gamma\}$, and for each time window the EEG signal is denoted as $\boldsymbol{X}_t \in \mathbb{R}^{C\times W}$, where W represents the number of sampling points. Then, we extract DE and FC features from EEG signals to capture temporal-spectral variations and inter-regional neural interaction patterns, respectively. The DE feature of the t-th time window is formulated as $\boldsymbol{F}_t^{\mathrm{DE}} \in \mathbb{R}^{C\times K}$, where $K = 5$ is the total number of frequency bands. The FC feature is represented as $\boldsymbol{F}_t^{\mathrm{FC}} \in \mathbb{R}^{K\times C\times C}$, where the connectivity strength between channel pairs is estimated within each frequency band and time window. The detailed definitions of DE and FC, together with the specific division of the frequency bands, are provided in Appendix A.

3.2 EMOTION RESPONSE PATTERN DISCOVERY

To characterize individual variations in emotional responses, the DE differences between the Happy and Sad conditions are computed as $\Delta\mathrm{DE} = \mathrm{DE}_{\mathrm{sad}} - \mathrm{DE}_{\mathrm{happy}}$. To further quantify emotion-related alterations in brain functional organization, the FC discrepancy between the two emotional conditions is calculated as $\Delta\mathrm{FC} = \mathrm{FC}_{\mathrm{sad}} - \mathrm{FC}_{\mathrm{happy}}$. Since individuals may exhibit diverse neural responses to identical video stimuli, average emotional trends cannot fully reflect individual differences; therefore, user-level emotional response features are extracted from two perspectives.

1. Direction Consistency (DC). DC is used to measure the temporal consistency of users' emotional changes. For the t-th time window, the direction of emotional feature variation is defined as $d_t = \mathrm{sign}(x_t - x_{t-1})$, where d_t denotes the feature variation direction between adjacent windows and x_t denotes the averaged emotional difference feature at the t-th window. DC is calculated in Equation equation 1:

$$\mathrm{DC} = \frac{1}{T-1} \sum_{t=2}^{T} \mathbb{I}\left(d_t = d_{t-1}\right), \tag{1}$$

where $\mathbb{I}(\cdot)$ denotes the indicator function. A higher DC indicates more consistent temporal response trends during emotional stimulation.

2. Temporal Stability (TS). TS is introduced to measure the stationarity of emotional response amplitude, with its calculation given in Equation equation 2:

$$\mathrm{TS} = 1 - \frac{\sigma(\boldsymbol{x})}{\mu(|\boldsymbol{x}|) + \epsilon}, \tag{2}$$

where $\sigma(\boldsymbol{x})$ denotes the standard deviation of the EEG feature sequence $\boldsymbol{x}$ across time windows, $\mu(|\boldsymbol{x}|)$ represents the mean absolute value of the sequence, and ϵ is a small constant. A higher TS indicates lower fluctuations of EEG responses across windows and more stable emotional responses. Finally, the user-level features $P_i = [\Delta\mathrm{DE}, \Delta\mathrm{FC}, \mathrm{DC}, \mathrm{TS}]$ are constructed for each user i. Based on these features, a two-stage unsupervised clustering method is employed to divide users into distinct categories. Details are provided in Appendix B.

3.3 PATTERN-AWARE CONTINUOUS EMOTION PERCEPTION

To integrate these EEG features, we propose a dual-branch affective perception network (PADAFN). Given a time window $\boldsymbol{X}_t$, the EEG Encoder is first employed to extract temporal features: $\boldsymbol{f}_{\mathrm{eeg}}^t = f_{\mathrm{eeg}}(\boldsymbol{X}_t)$. Meanwhile, the subject-specific DE and FC features are processed by the Pattern Encoder to obtain response features: $\boldsymbol{f}_{\mathrm{P}}^i = f_{\mathrm{P}}([\mathrm{DE}^i; \mathrm{FC}^i])$. To dynamically balance the contributions of EEG features and response features, an adaptive fusion module is introduced, which generates a fusion weight $\lambda_t = \sigma\left(\boldsymbol{W}\left[\boldsymbol{f}_{\mathrm{eeg}}^t; \boldsymbol{f}_{\mathrm{P}}^i\right] + \boldsymbol{b}\right)$ and produces the fused feature $\boldsymbol{f}_t$, as formulated in Equation equation 3:

$$\boldsymbol{f}_t = \sigma\left(\boldsymbol{W}\left[\boldsymbol{f}_{\mathrm{eeg}}^t; \boldsymbol{f}_{\mathrm{P}}^i\right] + \boldsymbol{b}\right) \boldsymbol{f}_{\mathrm{eeg}}^t + \left(1 - \sigma\left(\boldsymbol{W}\left[\boldsymbol{f}_{\mathrm{eeg}}^t; \boldsymbol{f}_{\mathrm{P}}^i\right] + \boldsymbol{b}\right)\right) \boldsymbol{f}_{\mathrm{P}}^i, \tag{3}$$

where $[\cdot\,;\,\cdot]$ denotes feature concatenation and $\sigma(\cdot)$ is the sigmoid function. The adaptive weight λ_t dynamically balances the EEG features and individual response features in the fused representation. The continuous valence score v_t is obtained from the fused features and normalized to $[-1, 1]$ using the tanh function, as expressed in Equation equation 4:

$$v_t = \tanh\left(\boldsymbol{W}_v \boldsymbol{f}_t + \boldsymbol{b}_v\right). \tag{4}$$

3.4 REINFORCEMENT LEARNING-BASED EMOTION RESISTANCE MODELING

During each video stimulation session, the PADAFN module outputs a continuous emotional response sequence $\{v_1, \ldots, v_T\}$, where v_t denotes the emotional state at the t-th time window. To capture the temporal dynamics of EEG, emotional inertia is introduced as a state feature. Specifically, a first-order autoregressive model is adopted to model the emotional trajectory, as formulated in Equation equation 5:

$$i_t = \phi_t\, v_{t-1} + c + \epsilon_t, \tag{5}$$

where ϕ_t denotes the inertia coefficient and ϵ_t is the stochastic fluctuation term.

In this paper, emotional anti-interference capacity estimation is formulated as a continuous-action reinforcement learning task. The agent interacts with a virtual perturbation environment, where external perturbations with different intensities are simulated to learn emotional state transitions. At time step t, the agent selects a perturbation action a_t according to the current state $\boldsymbol{s}_t$, where $a_t \in [0, 1]$ denotes the perturbation intensity. Smaller and larger values represent weaker and stronger perturbations, respectively. Considering the continuous evolution of emotional states, PPO is adopted for optimization. The policy network outputs the parameters of the action distribution $(\alpha_t, \beta_t) = \pi_\theta(\boldsymbol{s}_t)$, and the perturbation action is sampled from the Beta distribution: $a_t \sim \mathrm{Beta}(\alpha_t, \beta_t)$. Within the reinforcement learning environment, the emotional response trajectory generated by PADAFN is used, and the emotional variation between adjacent time windows is calculated as $\Delta v_t = v_{t+1} - v_t$. After the agent executes the perturbation action a_t, the environment updates the emotional state as formulated in Equation equation 6:

$$\Delta v'_t = (1 - a_t)\,(v_{t+1} - v_t) + \epsilon, \tag{6}$$

where ϵ is a small random perturbation term. The environment models the relationship between perturbation intensity and emotional state transition. Stable emotional states require stronger perturbations to induce transitions, whereas unstable states are more sensitive to weaker perturbations. Through iterative interactions, the agent learns to estimate the perturbation threshold associated with emotional state transition. To guide the agent toward effective perturbation policies, the reward function is designed as formulated in Equation equation 7:

$$R_t = |\Delta v_t| - |\Delta v'_t| - \lambda\, a_t, \tag{7}$$

where $|\Delta v_t| - |\Delta v'_t|$ measures the suppression effect of perturbations on emotional variation, and $\lambda\, a_t$ penalizes excessive perturbation intensity. The reward function encourages the agent to achieve effective emotional state transitions with lower perturbation levels. Consequently, the learned policy obtains the minimal perturbation intensity a_t^*.

Based on the optimal perturbation policy obtained by reinforcement learning, the Emotion Resistance Threshold (ERT) is defined as $\mathrm{ERT}(t) = a_t^*$. Here, $\mathrm{ERT}(t)$ denotes the minimum perturbation intensity required to induce emotional state transition at timestamp t. A higher ERT indicates stronger emotional stability under video stimulation, while a lower ERT suggests greater susceptibility to disturbances. The video-level emotional anti-interference capacity is further calculated according to Equation equation 8:

$$\mathrm{ERT}_{\mathrm{mean}} = \frac{1}{T} \sum_{t=1}^{T} \mathrm{ERT}(t). \tag{8}$$

Here, $\mathrm{ERT}_{\mathrm{mean}}$ represents the average anti-interference capacity throughout the video stimulation process. Meanwhile, the maximum anti-interference capacity is defined as $\mathrm{ERT}_{\mathrm{peak}} = \max_t \mathrm{ERT}(t)$, with the corresponding timestamp $t^* = \arg\max_t \mathrm{ERT}(t)$. This timestamp indicates the moment of strongest emotional stability during video stimulation.

3.5 Personalized Video Suitability Evaluation

To evaluate the regulatory effect of each video, we extract video-level features from the emotional trajectory. For each subject, the regulation features of a video are formulated as $\boldsymbol{f}^{\mathrm{v}} = [F^{\mathrm{v}}_{\mathrm{val}}, F^{\mathrm{v}}_{\mathrm{nat}}, F^{\mathrm{v}}_{\mathrm{eff}}, F^{\mathrm{v}}_{\mathrm{ERT}}]$, where $F^{\mathrm{v}}_{\mathrm{val}}$, $F^{\mathrm{v}}_{\mathrm{nat}}$, $F^{\mathrm{v}}_{\mathrm{eff}}$, and $F^{\mathrm{v}}_{\mathrm{ERT}}$ represent valence dynamics, natural emotional evolution, intervention effectiveness, and reinforcement-learning-based emotional stability features, respectively. Since different videos may induce different regulation effects for the same individual, the relative regulation features between candidate videos are further calculated: $\Delta \boldsymbol{f} = \boldsymbol{f}^{\mathrm{happy}} - \boldsymbol{f}^{\mathrm{sad}}$. The obtained features are then fed into a logistic regression model with L1 regularization to estimate the personalized suitability probability $P(y \mid \boldsymbol{f}) = \sigma\left(\boldsymbol{w}^\top \boldsymbol{f} + \boldsymbol{b}\right)$, based on which the optimal video is selected, as formulated in Equation equation 9:

$$v^* = \arg\max_{v \in \mathcal{V}} \; P(y \mid \boldsymbol{f}), \tag{9}$$

where $\mathcal{V}$ denotes the candidate video set. During inference, the video with the highest predicted suitability is recommended to the subject.

The overall procedure of the proposed EEG-PIEDA framework is summarized in Algorithm 1 in Appendix C.

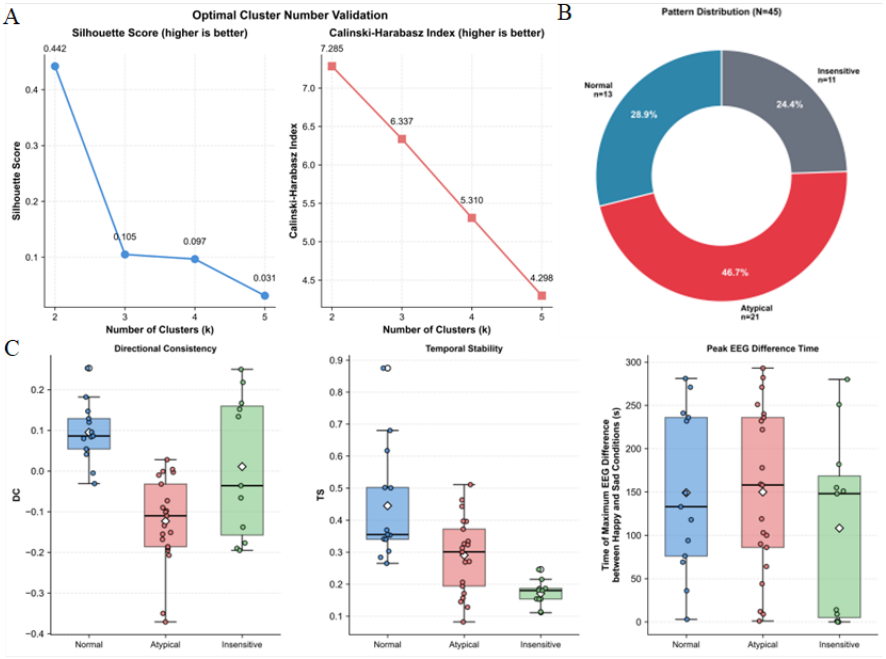

Figure 2: Emotional response clustering and EEG neural feature quantification. (A) Verification of the optimal cluster number in the second stage based on the Silhouette Score and CH Index. (B) Sample distribution of three emotional response patterns obtained via two-stage unsupervised clustering. (C) Quantitative comparison of DC, TS, and peak time of happy-sad EEG differences across different response patterns.

4 Experiments

Detailed descriptions of the three datasets are provided in Appendix D.

4.1 Analysis of Individual Emotional Response Patterns

With the constructed user-level feature representation P_i, subjects are partitioned into different categories through a two-stage unsupervised clustering procedure, as detailed in Appendix B. As shown in Figure 2(A), the silhouette score and CH Index achieved their highest values at $k = 2$, reaching 0.442 and 7.285, respectively. As k increased, both metrics gradually decreased, indicating reduced intra-cluster compactness and inter-cluster separability. Therefore, the remaining subjects were divided into two categories, Normal and Atypical. Combined with the Insensitive group identified in the first stage, all 45 subjects were finally partitioned into three groups, with the distribution shown in Figure 2(B).

To investigate the neural response differences among groups, we further analyzed the distributions of response indicators and differential FC. As shown in Figure 2(C), the three groups exhibited significant differences in DC, TS, and the peak EEG difference time between Happy and Sad conditions. Regarding DC, the three patterns exhibit different temporal consistency of happy-sad EEG differences. The Normal pattern shows the highest DC (median ≈ 0.09), with a positive-biased distribution, indicating consistent variation directions across time windows. The Atypical pattern has a lower DC (median ≈ -0.11), suggesting variable temporal trajectories despite distinguishable emotional responses. The Insensitive pattern shows DC values around zero with larger dispersion, reflecting weak directional consistency and limited emotional differentiation. For TS, the Normal pattern achieves the highest value (median ≈ 0.36), indicating more stable EEG difference responses over time. The Atypical pattern shows moderate temporal stability (median ≈ 0.29), while the Insensitive pattern presents the lowest TS (median ≈ 0.17), reflecting unstable and weakly maintained emotional differences. Together, DC and TS indicate that Normal exhibits more coherent temporal dynamics, whereas Atypical and Insensitive patterns show reduced temporal consistency.

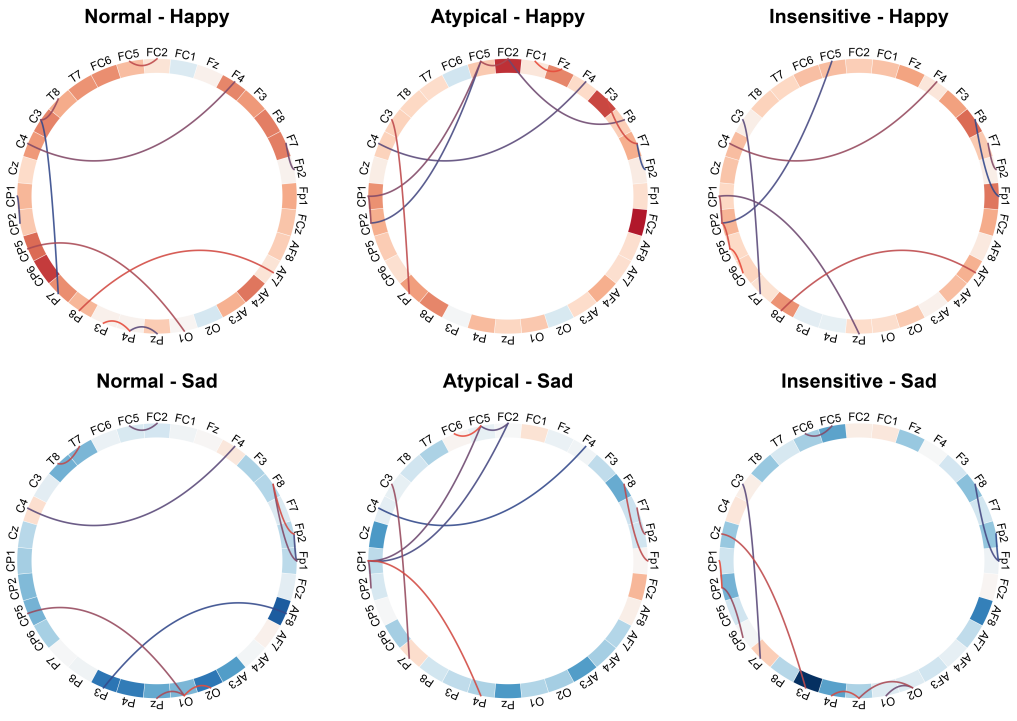

Figure 3: Group-level EEG channel connectivity across three emotional response patterns under happy and sad stimuli.

Table 1: Performance comparison of different EEG feature combinations and fusion strategies.

Selected Features	EEG	DE	FC	DE+FC	EEG+DE	EEG+FC	Full
Fixed Fusion	0.933 ± 0.004	0.608 ± 0.004	0.913 ± 0.003	0.915 ± 0.001	0.740 ± 0.008	0.911 ± 0.003	0.919 ± 0.003
Adaptive Fusion	–	–	–	–	0.728 ± 0.008	0.910 ± 0.003	$\mathbf{0.957 \pm 0.039}$

Furthermore, we analyzed the peak time at which the EEG difference between happy and sad conditions reached its maximum for each subject, with the detailed calculation procedure provided in Appendix E. As illustrated in Figure 2(C), the three categories (Normal, Atypical, Insensitive) yield distinct peak time points of EEG differential signals, with median peak latencies of roughly 150 s, 160 s and 140 s correspondingly. Specifically, the Atypical group has the latest peak response, the Insensitive group peaks ahead of the others, and the Normal group falls in between. Such temporal discrepancies demonstrate inter-subject temporal heterogeneity of emotion-evoked EEG activity.

Figure 3 visualizes interregional brain FC derived from the MIST dataset. The Normal pattern mainly activates the frontal lobe, parietal lobe, and frontotemporal cortices, brain areas tightly coupled with attentional cognition (Liu et al., 2016). Its dominant functional links shift with emotion: key connections are F6–Oz, FC4–P7, T8–P4 and CP5–Pz under happiness, whereas sad emotion features Fz–Pz, FC4–CP4 and F8–T8. The Atypical pattern recruits broader brain areas covering frontal, central, parietal, occipital and temporal lobes, core substrates for emotional processing (Gao et al., 2021). Happy cues strengthen FC5–C3, FC1–Cz, CP5–P3 and P4–O2, alongside robust frontal-central and central-posterior connectivity. Sadness induces prominent couplings including Fz–CPz, FC6–P8 and T8–Pz. The Insensitive pattern only engages the frontal, central and parietal lobes. Happy emotion relies on FC3–C3, Cz–CPz and CP4–P4, dominated by local cortico-cortical interactions within these three regions; sad emotion yields primary connections FC2–C4, C3–CP3 and CP2–Pz concentrated in central and parietal areas. Unlike the other two patterns, happy and sad states produce nearly identical connectivity layouts for the Insensitive pattern, with no emotion-dependent connectivity differences.

Table 2: Classification accuracy of PADAFN across three emotional response groups.

	Accuracy	Precision	Recall	F1	AUC	Kappa
Normal	$\mathbf{0.973 \pm 0.021}$	$\mathbf{0.978 \pm 0.026}$	$\mathbf{0.972 \pm 0.022}$	$\mathbf{0.972 \pm 0.022}$	$\mathbf{0.994 \pm 0.012}$	$\mathbf{0.946 \pm 0.042}$
Atypical	0.956 ± 0.044	0.952 ± 0.053	0.959 ± 0.038	0.955 ± 0.041	0.987 ± 0.027	0.911 ± 0.087
Insensitive	0.939 ± 0.036	0.942 ± 0.046	0.931 ± 0.044	0.936 ± 0.039	0.985 ± 0.018	0.877 ± 0.073
Full Data	0.957 ± 0.039	0.957 ± 0.047	0.955 ± 0.042	0.955 ± 0.038	0.988 ± 0.022	0.913 ± 0.077

Table 3: Performance comparison with state-of-the-art EEG representation methods on the MIST, DEAP, and SEED-IV datasets. * indicates marginal statistical significance ($p < 0.05$) compared with PADAFN on accuracy.

Method	MIST		DEAP		SEED-IV	
	Accuracy	F1	Accuracy	F1	Accuracy	F1
NeuroBolt*	0.636 ± 0.188	0.606 ± 0.211	0.521 ± 0.163	0.468 ± 0.171	0.328 ± 0.110	0.318 ± 0.102
LaBraM*	0.634 ± 0.109	0.632 ± 0.110	0.544 ± 0.060	0.540 ± 0.057	0.312 ± 0.033	0.306 ± 0.033
Our Method	$\mathbf{0.957 \pm 0.039}$	$\mathbf{0.955 \pm 0.038}$	$\mathbf{0.722 \pm 0.060}$	$\mathbf{0.674 \pm 0.010}$	$\mathbf{0.621 \pm 0.097}$	$\mathbf{0.620 \pm 0.098}$

4.2 Evaluation of Pattern-aware Continuous Emotion Perception

To evaluate the effectiveness of PADAFN, we compared different EEG feature combinations and fusion strategies. The results are shown in Table 1. The EEG-only model achieves an accuracy of 0.933 ± 0.004, demonstrating that the EEG encoder can effectively extract discriminative features from temporal EEG signals. The DE-only model obtains lower performance (0.608 ± 0.004), while the FC-only model achieves higher accuracy (0.913 ± 0.003). When combined with EEG features, EEG+DE and EEG+FC achieve 0.740 ± 0.008 and 0.911 ± 0.003, respectively, indicating that FC features provide more effective information for emotion discrimination. The full feature combination with fixed fusion achieves 0.919 ± 0.003, while PADAFN with the adaptive fusion module further improves accuracy to 0.957 ± 0.039. This improvement indicates that individual emotional response features provide additional subject-related information, and adaptive fusion can better integrate different feature contributions.

We evaluate PADAFN across different emotional response patterns. As shown in Table 2, PADAFN achieves accuracies of 0.973 ± 0.021, 0.956 ± 0.044, and 0.939 ± 0.036 on the Normal, Atypical, and Insensitive groups, respectively. The performance differences are consistent with the strength of emotion-related EEG responses. These results demonstrate the robustness of PADAFN in modeling heterogeneous emotional responses.

We further compare our method with NeuroBolt and LaBraM on three EEG datasets. As shown in Table 3, our method consistently achieves superior performance across all datasets. Specifically, it obtains an accuracy of 0.957 ± 0.039 and an F1-score of 0.955 ± 0.038 on MIST, as well as accuracies of 0.722 ± 0.060 and 0.621 ± 0.097 on DEAP and SEED-IV, respectively. The improvements indicate that incorporating individual response patterns provides more discriminative representations than conventional single-branch EEG encoders, allowing the model to better capture task-related neural variations. NeuroBolt and LaBraM adopt single-branch EEG encoders to map heterogeneous EEG signals into a unified feature space. This structure cannot distinguish task-relevant features from task-irrelevant features, restricting precise extraction of task-specific EEG features.

4.3 Evaluation of Reinforcement Learning-based Adaptive Regulation

To further evaluate whether the reinforcement learning policy can effectively utilize personalized emotional states for video intervention, we analyze the learned policy from three perspectives: individual emotional response variability, validity of state representation, and final recommendation performance. As shown in Figure 4(A), subjects exhibit distinct ERT values under happy and sad video interventions. Most subjects have ERT values around 0.048, while several subjects show higher resistance levels. For example, Subject 7 reaches 0.080 under the sad video condition, and Subject 35 reaches 0.074. In addition, within-subject differences are observed across video conditions; for instance, Subject 6 achieves ERT values of 0.100 and 0.063 under happy and sad videos, respectively. These results indicate significant individual variability in responses to video interventions. As shown in Figure 4(B), we further investigate the relationship between ERT differences and scores. The results show a consistent trend between ERT variations and scores, where subjects with higher scores are mainly distributed in regions with larger ERT differences, whereas subjects with lower scores exhibit relatively smaller ERT variations. This result suggests that ERT differences

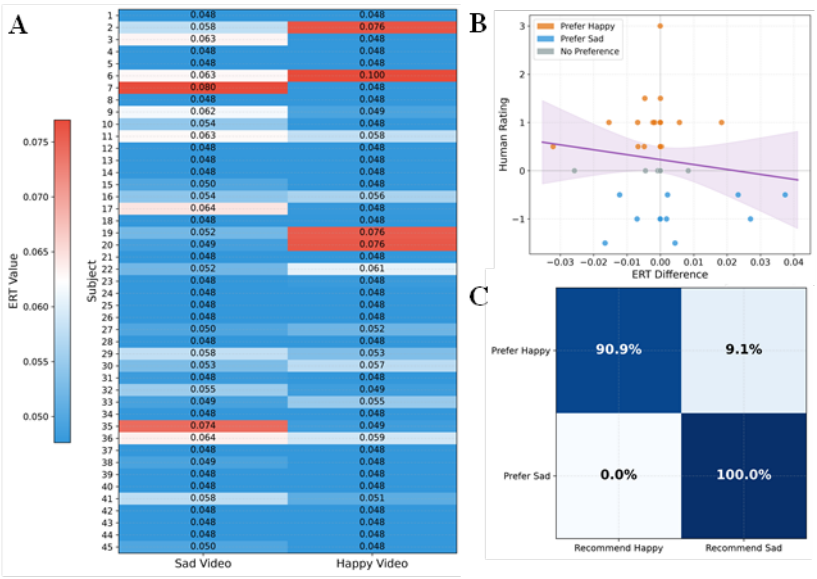

Figure 4: Validation of personalized state-driven reinforcement learning for adaptive video intervention. (A) Inter-subject variation of emotion resistance threshold. (B) Consistency between ERT-based state estimation and subject scores. (C) Confusion matrix of reinforcement learning-based video recommendation.

can effectively characterize emotional response variations under different video interventions and provide informative state feedback for reinforcement learning. As shown in Figure 4(C), the reinforcement learning policy achieves high consistency between recommended videos and individual emotional preferences, reaching 90.9% accuracy for users preferring happy interventions and 100% for users preferring sad interventions. These results demonstrate that the personalized state-feedback reinforcement learning strategy can effectively optimize video intervention decisions and improve the alignment between recommendations and individual emotional responses.

5 Discussion and Conclusion

This study investigates personalized emotional regulation by modeling individual differences in neural responses during empathic video stimulation. The results demonstrate that individuals exhibit distinct EEG feature patterns under the same emotional stimuli, revealing that emotional regulation is influenced not only by instantaneous neural states but also by individual-specific neural characteristics. By incorporating these features into PADAFN, EEG-PIEDA achieves more reliable continuous emotional state estimation across heterogeneous users. Furthermore, the reinforcement learning-based intervention module quantifies individual emotional resistance under different perturbation intensities and utilizes this information to support personalized video selection according to users' regulatory characteristics.

Technical Impact. These findings highlight the importance of modeling individual differences for personalized emotional regulation. EEG-PIEDA incorporates individual-specific neural features to characterize differences in emotional responses and regulatory capabilities. This method provides a more adaptive approach for understanding and supporting personalized emotion regulation.

Limitation and Future Direction. First, although EEG effectively captures neural dynamics during emotional regulation, incorporating additional physiological signals may provide complementary information for more comprehensive personalized modeling. Second, the current framework is evaluated using video stimuli only, and its effectiveness under diverse intervention modalities requires further investigation. Future work will focus on multimodal physiological integration, extending personalized regulation to diverse stimulation modalities, and developing more adaptive intervention strategies for practical applications.

AI USE STATEMENT

In this work, generative AI tools were used only for language assistance, namely translation and polishing of the manuscript text. We have not used generative AI tools for research ideation, experimental design, data collection, data analysis, or the generation of scientific results. All AI-assisted text has been reviewed and verified by the authors. We take responsibility for the final content of this work, including all text, claims, and artifacts.

ETHICS STATEMENT

This study involves EEG data collected from human subjects. All participants provided written informed consent before the experiment. The EEG data were anonymized prior to analysis.

REPRODUCIBILITY STATEMENT

For model development, the data are split into training and test sets with a ratio of 8:2. Model selection is performed via 5-fold cross-validation (CV) on the training set, and the model achieving the best validation performance is adopted. The selected model is then retrained on the entire training set and finally evaluated on the held-out test set.

REFERENCES

Yunyuan Gao, Zhen Cao, Jia Liu, and Jianhai Zhang. A novel dynamic brain network in arousal for brain states and emotion analysis. *Mathematical Biosciences and Engineering*, 18(6):7440–7463, 2021. doi: 10.3934/mbe.2021368.

Milena Girotti, Sarah E. Bulin, and Flavia R. Carreno. Effects of chronic stress on cognitive function – From neurobiology to intervention. *Neurobiology of Stress*, 33:100670, 2024. doi: 10.1016/j.ynstr.2024.100670.

K. Kannadasan, B. Shameedha Begum, and N. Ramasubramanian. Neurological insights into personality and emotional intelligence: Analysing EEG spectral bands through machine learning models. *Biomedical Signal Processing and Control*, 113:109190, 2026. doi: 10.1016/j.bspc.2025.109190.

Elena Kozakevich Arbel, Simone G. Shamay-Tsoory, and Uri Hertz. Adaptive empathic response selection is sensitive to multiple dimensions of social interaction. *Communications Psychology*, 2(1):112, 2024. doi: 10.1038/s44271-024-00164-8.

Peter Kuppens, Nicholas B. Allen, and Lisa B. Sheeber. Emotional inertia and psychological maladjustment. *Psychological Science*, 21(7):984–991, 2010. doi: 10.1177/0956797610372634.

Emily G. Lattie, Colleen Stiles-Shields, and Andrea K. Graham. An overview of and recommendations for more accessible digital mental health services. *Nature Reviews Psychology*, 1(2):87–100, 2022. doi: 10.1038/s44159-021-00003-1.

Wenxin Li, Xiao Song, and Yuchun Tu. GraphDRL: GNN-based deep reinforcement learning for interactive recommendation with sparse data. *Expert Systems with Applications*, 273:126832, 2025. doi: 10.1016/j.eswa.2025.126832.

Yuelu Liu, Jesse Bengson, Haiqing Huang, George R. Mangun, and Mingzhou Ding. Top-down modulation of neural activity in anticipatory visual attention: Control mechanisms revealed by simultaneous EEG-fMRI. *Cerebral Cortex*, 26(2):517–529, 2016. doi: 10.1093/cercor/bhu204.

Sebastian Lubos, Alexander Felfernig, and Markus Tautschnig. An overview of video recommender systems: state-of-the-art and research issues. *Frontiers in Big Data*, 6:1281614, 2023. doi: 10.3389/fdata.2023.1281614.

Yao Pi, Chen Yang, Xianbin Zhang, Richard Millham, Guibin Bian, Shenglin Wen, and Wanqing Wu. Knowledge-guided attention fusion of multi-paradigm EEG for depressive episode detection. *Neurocomputing*, 700:134344, 2026. doi: 10.1016/j.neucom.2026.134344.

Xubin Ren, Wei Wei, Lianghao Xia, and Chao Huang. A comprehensive survey on self-supervised learning for recommendation. *ACM Computing Surveys*, 58(1):22, 2025. doi: 10.1145/3746280.

Stefan Stieger, Irina Schmid, Philip Altenburger, and David Lewetz. Short-, medium-, and long-term impact of watching humorous video clips on stress and well-being: An experience sampling method-based field experiment. *Computers in Human Behavior Reports*, 10:100270, 2023. doi: 10.1016/j.chbr.2023.100270.

Weilong Tan, Hongyi Zhang, Yingbei Wang, Weimin Wen, Liang Chen, Han Li, Xingen Gao, and Nianyin Zeng. SEDA-EEG: A semi-supervised emotion recognition network with domain adaptation for cross-subject EEG analysis. *Neurocomputing*, 622:129315, 2025. doi: 10.1016/j.neucom.2024.129315.

Hao Wang, Zhichao Chen, Honglei Zhang, Zhengnan Li, Licheng Pan, Haoxuan Li, and Mingming Gong. Debiased recommendation via Wasserstein causal balancing. *ACM Transactions on Information Systems*, 43(6):145, 2025. doi: 10.1145/3725731.

Yiming Wang, Bin Zhang, and Lamei Di. Learning from semantic ambiguity: A dual-noise robust framework for EEG partial label emotion recognition. *Neurocomputing*, 701:134587, 2026. doi: 10.1016/j.neucom.2026.134587.

Mengqi Wu, C. L. Philip Chen, Bianna Chen, and Tong Zhang. GROP: Graph orthogonal purification network for EEG emotion recognition. *IEEE Transactions on Affective Computing*, 16(1): 319–332, 2025. doi: 10.1109/TAFFC.2024.3433613.

Zijian Zhang, Shuchang Liu, Ziru Liu, Rui Zhong, Qingpeng Cai, Xiangyu Zhao, Chunxu Zhang, Qidong Liu, and Peng Jiang. LLM-powered user simulator for recommender system. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 13339–13347, Palo Alto, CA, 2025. AAAI Press.

Sha Zhao, Yitian Liu, Shijian Li, and Gang Pan. Cross-subject EEG-based emotion recognition leveraging multi-source domain adaptation with curriculum leaning strategy. *Neurocomputing*, 697:134158, 2026. doi: 10.1016/j.neucom.2026.134158.