

# Personalized Dynamic Music Emotion Recognition with Dual-Scale Attention-Based Meta-Learning

Dengming Zhang<sup>1</sup>, Weitao You<sup>2\*</sup>, Ziheng Liu<sup>2</sup>, Lingyun Sun<sup>2</sup>, Pei Chen<sup>2</sup>

<sup>1</sup>School of Software Technology, Zhejiang University, China

<sup>2</sup>College of Computer Science and Technology, Zhejiang University, China  
{dmz, weitao\_you, northstar\_lzh, sunly, chenpei}@zju.edu.cn

## Abstract

Dynamic Music Emotion Recognition (DMER) aims to predict the emotion of different moments in music, playing a crucial role in music information retrieval. The existing DMER methods struggle to capture long-term dependencies when dealing with sequence data, which limits their performance. Furthermore, these methods often overlook the influence of individual differences on emotion perception, even though everyone has their own personalized emotional perception in the real world. Motivated by these issues, we explore more effective sequence processing methods and introduce the Personalized DMER (PDMER) problem, which requires models to predict emotions that align with personalized perception. Specifically, we propose a Dual-Scale Attention-Based Meta-Learning (DSAML) method. This method fuses features from a dual-scale feature extractor and captures both short and long-term dependencies using a dual-scale attention transformer, improving the performance in traditional DMER. To achieve PDMER, we design a novel task construction strategy that divides tasks by annotators. Samples in a task are annotated by the same annotator, ensuring consistent perception. Leveraging this strategy alongside meta-learning, DSAML can predict personalized perception of emotions with just one personalized annotation sample. Our objective and subjective experiments demonstrate that our method can achieve state-of-the-art performance in both traditional DMER and PDMER.

**Code & Case** — <https://littleor.github.io/PDMER>

## Introduction

Music Emotion Recognition (MER) technology focuses on identifying emotions conveyed by music, applying to music therapy (Dingle et al. 2015), music recommendation (Liu et al. 2023; Tran et al. 2023), and music generation (Huang and Huang 2020; Ji and Yang 2024). To describe the music emotion, Russell’s two-dimensional valence-arousal (V-A) emotional model (Russell 1980) is widely used in MER, where valence describes the extent to which an emotion is positive or negative, and arousal refers to its intensity. Existing MER tasks are generally divided into static MER (SMER) and dynamic MER (DMER) (Han et al. 2022).

\*Corresponding author




Figure 1: The differences between traditional DMER and PDMER. All charts represent the emotion valence/arousal (V/A) curve of music, where the x-axis represents time and the y-axis represents V/A values.

SMER inputs music and outputs only one V-A label to describe the emotion, which fails to describe the variations in emotion within the music. For example, Beethoven’s Symphony No. 5 can’t be described as simply positive and intense, as it also contains moments of sadness and tranquility. In contrast, DMER predicts the V-A label sequence, using a sequence label to describe the emotional changes in the music, which can more accurately express the emotions.

Existing DMER work focuses on utilizing sequential information to predict, as the emotion of each moment in music is related to the emotions before and after. Specifically, long short-term memory (LSTM) (Hochreiter and Schmidhuber 1997) attracted the attention of researchers due to its superiority in sequence data processing. These researchers use LSTM to extract sequence features and have shown some effectiveness in DMER (He and Ferguson 2020; Zhang et al. 2023). However, previous studies (Khandelwal et al. 2018; Li et al. 2019; Grigsby et al. 2021) have found that LSTM struggles to capture long-term dependencies, which limits the capture of the global emotion of music, resulting in poor performance in DMER.

More importantly, existing work often assumes that all

individuals perceive music emotions in the same way, neglecting the significant impact of individual differences on DMER. For example, MER1101 dataset (Zhang et al. 2023) involves multiple annotators for labeling and uses their average as labels to attempt to eliminate individual differences in the data as shown in Figure 1(a). However, emotions are personalized, and different individuals have different perceptions of emotions towards the same song (Kang and Herremans 2024). For example, a song that makes one person feel happy may make another person feel sad. These existing works eliminate individual differences in data while eliminating bias, using group emotional perceptions instead of individual emotional perceptions to avoid personalization issues. To make matters worse, real-world applications often face challenges with diverse emotional perceptions from each individual, and removing personalization differences in datasets cannot address this issue. Therefore, we point out the personalized DMER (PDMER) problem, which requires models to predict emotions that align with individual personalized perceptions rather than group perceptions.

To address the above issues, we propose a Dual-Scale Attention-Based Meta-Learning (DSAML) method to handle PDMER tasks. Specifically, to achieve DMER, DSAML uses a pre-trained Imagebind model (Girdhar et al. 2023) to extract global audio features and introduces an adapter to extract trainable local audio features. Fusion of the two features is fed into the dual-scale attention transformer, which focus on both local and global features to capture more comprehensive emotional features. Finally, a sequence of V-A values is predicted by the sequence prediction module, completing the DMER. To achieve PDMER, the Model-Agnostic Meta-Learning (MAML) (Finn, Abbeel, and Levine 2017) is used to personalize the model. Furthermore, we propose a personalized meta-learning task construction strategy for MAML, which divided tasks by annotators instead of music samples. Samples in a task are annotated by the same annotator, ensuring consistent perception. With this strategy, DSAML preserves personality differences, effectively enhancing PDMER performance.

In summary, this paper has the following contributions:

- We propose the DSAML including a dual-scale feature extractor and a dual-scale attention transformer to capture both local and global features, improving the performance of traditional DMER.
- We recognize the importance of personalized emotional perception for DMER, introduces the PDMER task, and design a personalized prediction method based on meta-learning with a novel task construction strategy.
- Objective experiments demonstrate that our method achieves the best performance in both traditional DEMR and PDMER. Subjective experiments also show that our method better conforms to individual personalized emotional perception in the real world.

## Related Work

### Dynamic Music Emotion Recognition

DMER aims to predict the emotions of music at different moments, as the emotions at any moment are often related to

those before and after, so it is necessary to consider the temporal dependencies in the music emotions. As the recurrent neural network (RNN) is suitable for sequence data processing, researchers first used RNN to extract sequence features for DMER (Malik et al. 2017). However, due to the issue of gradient vanishing in RNNs when dealing with long sequences, researchers in the DMER field have subsequently adopted LSTM to extract sequential features, as LSTM introduces gating mechanisms for long sequence processing. Specifically, Zhang et al. (Zhang et al. 2023) integrated spatial and channel dimension features and used Bi-directional LSTM (BiLSTM) for sequence learning to predict the V-A sequence of music. He et al. (He and Ferguson 2020) used multi-view CNN as feature extractors and then used BiLSTM to capture temporal context for predicting the V-A sequence of music. These methods have shown some effectiveness in DMER, but they still struggle to capture long-term dependencies as previously studies demonstrated (Khandelwal et al. 2018; Li et al. 2019; Grigsby et al. 2021). Therefore, we propose a dual-scale feature extractor and a dual-scale attention transformer to capture both short- and long-term dependencies, thereby improving the performance of long-sequence data processing.

### Personalized Music Emotion Recognition

In the SMER field, researchers have conducted some PMER studies, Yang et al. (Yang et al. 2007) first quantitatively evaluated the impact of personality on MER and found that personalization significantly influences MER. Based on this quantitative experiment, numerous personalized MER studies have appeared in the field of SMER. To achieve PMER, researchers mainly train personalized models using samples annotated by specific users. Personalized training depends on a large amount of user-annotated data, so researchers focus on reducing the number of specific user-annotated samples. Su and Fung (Su and Fung 2012) proposed an active learning method to reduce the number of specific user-annotated samples by selecting the most informative ones for manual annotation. Wang et al. (Wang et al. 2012) and Chen et al. (Chen et al. 2014) used two-stage training, first training the background model and then adapting it using fewer specific user-annotated samples.

However, in the field of DMER, there has been no specialized research on PMER, researchers often use the mean of multiple data annotators as the emotion label which ignores the impact of individual personalization on MER (Orjesek et al. 2019; Zhang et al. 2023). In other words, many pieces of DMER research focus on group emotions rather than individual emotions, which limits the practical application performance of existing DMER. Although the SMER methods have achieved good results in PMER, they still require a large amount of user-annotated data (at least 20 samples) (Chen et al. 2014). In DMER, users need to annotate more labels (e.g., 1200 labels for 20 30-second music, assuming labeling every 0.5 seconds), making these personalized SMER methods difficult to apply in DMER. Therefore, we design a PDMER method based on meta-learning with a new meta-learning task construction strategy, which only requires one specific user-annotated sample.

## Meta-learning

Meta-learning is widely used in few-shot learning problems because it can learn to solve new tasks with a small number of samples (Thrun and Pratt 1998; Wang et al. 2020). Meta-learning methods can be generally divided into three types: metric-based, model-based, and optimization-based meta-learning methods. Among them, metric-based meta-learning methods focus on solving classification problems in the feature space (Snell, Swersky, and Zemel 2017; Sung et al. 2018) and are difficult to apply to regression tasks like DMER; model-based meta-learning methods focus on designing specific model structures to achieve fast learning goals (Wang et al. 2020), which limits the model structure; while optimization-based meta-learning methods adjust existing optimization algorithms to converge with a few samples for new tasks (Finn, Abbeel, and Levine 2017; Nichol, Achiam, and Schulman 2018). For example, MAML has no constraints on model structure and can also be applied to regression tasks, making it most suitable for application in DMER. However, directly applying MAML in PDMER still ignores individual personalized effects during the training process, limiting the performance of MAML. Therefore, we propose a new method for constructing meta-learning tasks: incorporating the annotated data of each annotator and constructing tasks according to annotators, thereby effectively improving the performance of the model in the PDMER task.

## Methods

### Problem Formulation

Given a training dataset  $\mathbf{D} = \{(x_1, y_1^1), \dots, (x_i, y_i^j), \dots\}$ , and personalized annotated data (annotator not present in the training set)  $\mathbf{S}_p = \{(x_{p1}, y_{p1}^p), \dots, (x_{pn}, y_{pn}^p)\}$ , where  $x_i$  is the  $i$ -th music sample in  $\mathbf{D}$ ,  $y_i^j$  is the label from the  $j$ -th annotator for  $x_i$ ,  $p$  represents any user with personalized emotional perception,  $x_{pn}$  is the  $n$ -th music sample in  $\mathbf{S}_p$ , and  $y_{pn}^p$  is the label for  $x_{pn}$  from the user  $p$ . In traditional DMER, the goal is to directly predict non-personalized label  $y_q$  for a new query sample  $q$  by training on  $\mathbf{D}$ . However, in PDMER, our goal is to predict the personalized label  $y_q^p$  for the specific user  $p$  by training on  $\mathbf{D}$  and adapting it with  $\mathbf{S}_p$ .

### Model Architecture

The DSAML model consists of four parts: Input Preprocessor, Dual-Scale Feature Extractor, Dual-Scale Attention Transformer, and Sequence Predictor, as shown in Figure 2.

**Input Preprocessor** To achieve DMER, the model needs to extract sequence features from the audio. Therefore, in the input preprocessor, we slice the original audio input. A music segment of length  $l$  seconds is sliced into  $k$  segments, corresponding to  $k$  sequence prediction values and a resolution of  $k/l$  Hz. The sliced music segments are then processed to calculate the log Mel-spectrogram.

**Dual-Scale Feature Extractor** The pre-trained Imagebind model has shown good performance in audio feature extraction (Zou et al. 2023; CHAKHTOUNA, SEKKATE,

and Abdellah 2024). Consequently, DSAML utilizes Imagebind to extract the global feature  $\mathbf{z}_g$  of the audio  $\mathbf{x}$ . However, the large number of parameters in Imagebind would significantly increase the personalization adaptation time in PDMER. Therefore, this paper freezes the parameters of Imagebind and introduces the Imagebind Adapter module with significantly fewer parameters. Furthermore, since the global feature extracted by Imagebind cannot represent the finer emotional changes in music, the Imagebind Adapter is designed to extract the local feature  $\mathbf{z}_l$  from the spectrogram sequence  $\mathbf{x}'$  of the short music segments. The local and global features are then fused to obtain the audio feature  $\mathbf{z} = \sigma(\mathbf{z}_l + \mathbf{z}_g)$ , where  $\sigma$  refers to the Sigmoid function. Specifically, as shown in Figure 3, the Imagebind Adapter module extracts features using two convolutional layers, reduces the number of channels to 1 through a  $1 \times 1$  convolutional layer, and finally maps the feature dimension to  $D$  through a fully connected layer.

**Dual-Scale Attention Transformer** Considering that the emotional state of music at a particular moment is often related to the music in the preceding moments, and that the overall emotion also influences the emotional state at that moment, this paper proposes a Dual Attention Transformer to extract context-aware emotional features  $\mathbf{z}'$  from audio features  $\mathbf{z}$ . The Dual Attention Transformer achieves different scales of attention through a local mask  $\mathbf{M}_l$  and a global mask  $\mathbf{M}_g$ , where the context length  $n_l$  of  $\mathbf{M}_l$  is much smaller than the context length  $n_g$  of  $\mathbf{M}_g$ , thereby enabling the Transformer’s dual attention to focus on both local and global emotions. For a mask  $\mathbf{M}_n$  with context length  $n$ , it can be expressed as:

$$\mathbf{M}_n[i, j] = \begin{cases} 1 & \text{if } |i - j| \leq n \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where  $i$  and  $j$  represent different time steps in the sequence, and  $|i - j| \leq n$  indicates that at any time step  $i$ , the information at time step  $j$  can be seen only if their distance does not exceed  $n$ . Using  $\mathbf{M}_l$  and  $\mathbf{M}_g$  as masks, local features  $\mathbf{z}'_l$  and global emotional features  $\mathbf{z}'_g$  are extracted from the audio features  $\mathbf{z}$ :

$$\mathbf{z}'_l = f_\theta(\mathbf{z}, \mathbf{M}_l); \mathbf{z}'_g = f_\theta(\mathbf{z}, \mathbf{M}_g) \quad (2)$$

where  $f$  represents the transformer module, and  $\theta$  represents the parameters of the transformer module. It is noteworthy that  $\mathbf{z}'_l$  and  $\mathbf{z}'_g$  are extracted using a transformer module with shared parameters, only utilizing different masks to extract features at different scales. Finally, the local and global emotional features are fused  $\mathbf{z}' = \sigma(\mathbf{z}'_l + \mathbf{z}'_g)$  to obtain the context-aware emotional features  $\mathbf{z}'$ .

However, our analysis of the local attention maps  $\mathbf{A}_l$  and global attention maps  $\mathbf{A}_g$  under different masks reveals that merely constraining the attention through masks does not ensure that  $\mathbf{A}_l$  focuses more on the local context and  $\mathbf{A}_g$  on the global context. As shown in Figure 4, it is possible that  $\mathbf{A}_l$  might focus on distant moments and  $\mathbf{A}_g$  on nearby moments, which results in extracted and fused features that are no longer comprehensive. To address this issue, we propose a diagonal attention map loss, which constrains the diagonal attention of  $\mathbf{A}_l$  to be higher and that of  $\mathbf{A}_g$  to be lower,

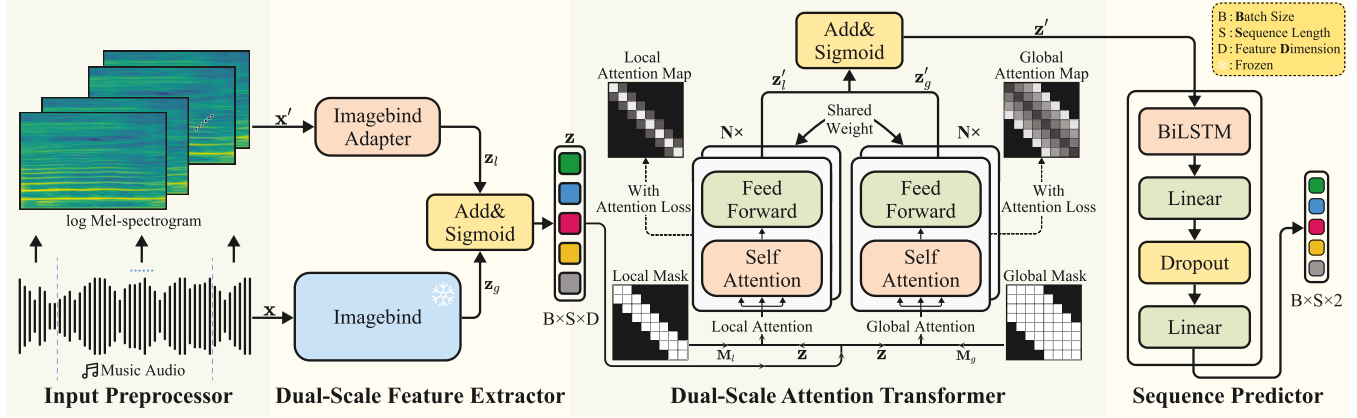


Figure 2: The architecture of the DSAML model.




Figure 3: The architecture of the Imagebind Adapter

ensuring that the two types of attention focus on different scales. It can be expressed as:

$$\mathcal{L}_{\text{attention}} = \frac{1}{n} \sum_{i=1}^n [(\text{diag}(\mathbf{A}_l)_i - \alpha)^2 + (\text{diag}(\mathbf{A}_g)_i - \beta)^2] \quad (3)$$

where  $\text{diag}(\mathbf{A})$  denotes the diagonal values of the attention map  $\mathbf{A}$ , and  $\alpha$  and  $\beta$  are hyperparameters that control the diagonal attention of  $\mathbf{A}_l$  and  $\mathbf{A}_g$ , respectively. Specifically,  $\alpha$  is set to a higher value to ensure that  $\mathbf{A}_l$  focuses more on the local context, while  $\beta$  is set to a lower value to ensure that  $\mathbf{A}_g$  focuses more on the global context.




Figure 4: Attention Map

**Sequence Predictor** In the sequence predictor, DSAML further processes the sequence features using BiLSTM and reduces the feature dimension to 2 through fully connected layers to achieve regression of the V-A values, thus completing the DMER of the music.

### Personalized Strategy

DSAML employs MAML for personalized learning and proposes a novel meta-learning task construction strategy to enhance the model’s personalized prediction performance. The objective of MAML is to find a model parameter  $\theta$  for the task distribution  $p(\mathcal{T})$ , such that the loss  $\mathcal{L}_{\tau_i}$  is minimized after  $k$  steps of learning on a randomly sampled task  $\tau_i \sim p(\mathcal{T})$ , which can be expressed as:

$$\min_{\theta} \sum_{\tau_i \sim p(\mathcal{T})} \mathcal{L}_{\mathbf{Q}_i} [U_{\mathbf{S}_i}^k(\theta)] \quad (4)$$

where  $\mathbf{S}_i$  and  $\mathbf{Q}_i$  are the support set and query set randomly sampled from  $\tau_i$  ( $\mathbf{S}_i \cap \mathbf{Q}_i = \emptyset$ ), and  $U_{\mathbf{S}_i}^k(\theta)$  is the operator that updates  $\theta$   $k$  times using the  $\mathbf{S}_i$ .

Defined in the problem formulation, we have a training dataset  $\mathbf{D} = \{(x_1, y_1^1), \dots, (x_i, y_i^j), \dots\}$  to find the optimal  $\theta$ , where each music has labels from multiple annotators. As shown in Figure 1(a), the existing DMER approaches often use the mean of all annotator labels for each music as the label, and train on the processed dataset  $\mathbf{D}' = \{(x_1, y_1'), \dots, (x_n, y_n')\}$ , where  $y_n' = \frac{1}{N} \sum_{i=1}^N y_n^i$ . The traditional task construction strategy for training with MAML using  $\mathbf{D}'$  can be represented as:

$$\mathbf{S}_i = \text{RandomSample}(\mathbf{D}'); \mathbf{Q}_i = \text{RandomSample}(\mathbf{D}' - \mathbf{S}_i) \quad (5)$$

Although this method performs well on traditional DMER and many datasets often directly provide  $\mathbf{D}'$  (Aljanaki, Yang, and Soleymani 2017; Zhang et al. 2018), this type of dataset and task construction approach loses personalized emotional preference information and essentially represents group emotional preferences. Therefore, we propose a novel meta-learning task construction strategy, which builds tasks directly from the dataset  $\mathbf{D}$  based on the annotators:

$$\mathbf{S}_i = \text{RandomSample}(\mathbf{D}_i); \mathbf{Q}_i = \text{RandomSample}(\mathbf{D}_i - \mathbf{S}_i) \quad (6)$$

where  $\mathbf{D}_i$  refers to all the music annotated by the  $i$ -th user, i.e.,  $\mathbf{D}_i = \{(x_1^i, y_1^i), (x_2^i, y_2^i), \dots, (x_m^i, y_m^i)\}$ . This method treats each user’s personality as different tasks, enabling the model to find a  $\theta$  that can adapt quickly and perform optimally across all personalities.

## Training & Inference Process

During the training process, we construct task distribution  $p(\mathcal{T})$  from the training dataset  $\mathbf{D}$  using the above personalized task construction strategy. A batch of tasks is randomly sampled from  $p(\mathcal{T})$ , and from each task  $\tau_i$ , we will randomly sample support set  $\mathbf{S}_i$  and query set  $\mathbf{Q}_i$ . Specifically,  $\mathbf{S}_i$  is used to optimize  $\theta$  and obtain  $\theta'_i = U_{\mathbf{S}_i}^k(\theta)$ , then the training loss will be calculated on  $\mathbf{Q}_i$ . After accumulating the losses of all tasks in this batch, the model parameters will be updated as  $\theta \leftarrow \theta - \eta \nabla_{\theta} \sum_{\tau_i \sim p(\mathcal{T})} \mathcal{L}_{\mathbf{Q}_i}(\theta'_i)$ , where  $\eta$  is the learning rate. This process repeats until the model converges, and the fitted parameters  $\hat{\theta}$  are used for inference.

In the inference process, we use all samples annotated by a new user  $p$  (anyone absent from  $\mathbf{D}$ ) in the personalized data  $\mathbf{S}_p$ . These samples serve as a support set to fast adapt the model  $\hat{\theta}$  to the user’s personalized emotional perception and obtain the personalized model  $\hat{\theta}'_p$  for the new user  $p$ . Finally, the model  $\hat{\theta}'_p$  can predict any music’s emotion that aligns with the personalized perception of the new user  $p$ .

## Implementation Details

The resolution of DSAML is 2Hz, indicating there is one label every 0.5 seconds. The model architecture consists of 3 layers of Transformer, with a mask context length of  $n_l = 5$  and  $n_g = 30$ . In attention loss,  $\alpha = 0.5$  and  $\beta = 0.05$ . During training, only one sample is used for fast adaptation (i.e., both  $\mathbf{S}_i$  and  $\mathbf{S}_p$  only contain 1 sample), and 15 samples are used for evaluation (i.e.,  $\mathbf{Q}_i$  contains 15 samples). The Adam optimizer (Kingma and Ba 2014) is employed with a learning rate of 0.00005. We train the model for 2000 episodes on a single NVIDIA GeForce RTX 4090 GPU.

## Experiments

### Experiment Settings

**Dataset** The performance of DSAML is evaluated using two publicly available DMER datasets, both of which provide V-A value annotations every 0.5 seconds, with all unstable annotations from the beginning to 15 seconds removed.

The first dataset is the DEAM dataset (Aljanaki, Yang, and Soleymani 2017), which includes 1744 45-second clips and 58 full-length songs with an average length of 4 minutes, all containing dynamic annotations from each annotator as illustrated in Figure 5. In our experiments, we use the 58 full-length songs as the test set, with the remaining 1744 songs as the training set. Notably, 744 of the 45-second clips do not have annotator IDs, meaning we cannot determine the annotators for these songs. Therefore, in our proposed personalized task construction strategy, we only use 1000 songs as the training set.

The second dataset is the PMemo dataset (Zhang et al. 2018), which includes 794 songs of varying lengths and the mean and standard deviation of the dynamic annotations, without individual annotator data. Thus, our personalized task construction strategy cannot be applied to this dataset, and this dataset is only used for evaluating traditional DMER tasks. We discard 122 samples with song lengths less than 25




Figure 5: Example of personalized annotations for the same song with three different annotators in the DEAM dataset.

seconds and use 40 songs longer than 65 seconds as the test set, with the remaining 632 songs as the training set.

**Compared Models** Owing to the lack of research specifically targeting PDMER, this paper utilizes the DMER methods as a baseline, which includes: (1) CRNN (Convolutional and Recurrent Neural Network) (Malik et al. 2017): This method stacks convolutional and recurrent neural networks to predict emotions using a compact architecture with fewer parameters. (2) DNN (Deep Neural Network) (Orjesek et al. 2019): This approach leverages convolutional and recurrent neural networks for feature extraction directly from raw audio without any preprocessing. (3) MCRNN (Multi-view Convolutional Recurrent Neural Network) (He and Ferguson 2020): This technique employs multi-view CNNs and BiLSTM to automatically learn feature representations from raw audio, incorporating data augmentation methods. (4) DAMFF (Dual Attention-based Multi-scale Feature Fusion) (Zhang et al. 2023): As the SOTA DMER method, it introduces a dual attention mechanism for temporal-frequency multi-scale feature fusion from spectrograms, employing BiLSTM for sequence learning.

**Objective Metrics.** The performance of DSAML is evaluated using three objective metrics: Root Mean Square Error (RMSE), Pearson Correlation Coefficient (PCC), and Concordance Correlation Coefficient (CCC). Among these, the RMSE metric is used to measure the deviation between the predicted values and the actual values. A smaller RMSE value indicates higher prediction accuracy and lower error of the model. The PCC is used to assess the linear correlation between the predicted values and the actual values, with larger values indicating a stronger positive correlation. The CCC combines both precision and consistency aspects, serving as an improvement over the PCC. It not only considers the linear correlation but also takes into account the agreement between the means and variances of the predicted and actual values. A higher CCC value indicates better predictive performance of the model.

### Objective Experiment

This paper evaluates the performance of DSAML on PDMER and traditional DMER separately. For PDMER, the test set uses personalized annotations from all annotators as labels, referred to as the personalized task. In contrast, for traditional DMER, we follow conventional validation meth-

| Model | DEAM           |                |                   |                |                |                   | PMemo          |                |                   |                |                |                   |
|-------|----------------|----------------|-------------------|----------------|----------------|-------------------|----------------|----------------|-------------------|----------------|----------------|-------------------|
|       | Arousal        |                |                   | Valence        |                |                   | Arousal        |                |                   | Valence        |                |                   |
|       | CCC $\uparrow$ | PCC $\uparrow$ | RMSE $\downarrow$ | CCC $\uparrow$ | PCC $\uparrow$ | RMSE $\downarrow$ | CCC $\uparrow$ | PCC $\uparrow$ | RMSE $\downarrow$ | CCC $\uparrow$ | PCC $\uparrow$ | RMSE $\downarrow$ |
| Ours  | <b>0.39</b>    | <b>0.675</b>   | <b>0.202</b>      | <b>0.104</b>   | <b>0.221</b>   | <b>0.26</b>       | <b>0.186</b>   | <b>0.386</b>   | <b>0.112</b>      | <b>0.108</b>   | 0.26           | <b>0.122</b>      |
| DAMFF | 0.354          | 0.621          | 0.217             | 0.055          | 0.043          | 0.299             | 0.175          | 0.378          | 0.118             | 0.091          | <b>0.335</b>   | 0.134             |
| MCRNN | 0.227          | 0.374          | 0.265             | 0.029          | -0.005         | 0.297             | 0.104          | 0.266          | 0.168             | 0.038          | 0.154          | 0.169             |
| DNN   | 0.188          | 0.362          | 0.261             | 0.033          | 0.003          | 0.281             | 0.087          | 0.217          | 0.151             | 0.03           | 0.134          | 0.159             |
| CRNN  | 0.244          | 0.517          | 0.218             | 0.017          | -0.039         | 0.27              | 0.076          | 0.231          | 0.183             | 0.034          | 0.224          | 0.169             |

Table 1: Performance of different models on two datasets of the traditional DMER task.

ods where the test set uses the mean of annotations from all annotators as labels, referred to as the traditional DMER task.

**Traditional DMER Task** In traditional DMER tasks, this paper follows the approach of other DMER studies by using the mean of all annotators as the label for validation. Table 1 presents the performance of our method compared to other methods on the DEAM and PMemo datasets. It is observed that our method outperforms other methods across all metrics on the DEAM dataset, and also performs well on the PMemo dataset. This is mainly because the emotional state at any moment in DMER may be influenced by distant moments, and our model architecture design with a dual-scale attention transformer can better capture long-term dependencies, thereby improving prediction accuracy. This is also validated in the ablation study. Moreover, the PCC value in the valence dimension of PMemo dataset is lower than the DAMFF method, which may be due to the shorter music segments in this dataset, with fewer long-term dependencies. The trend of the valence dimension can be more easily predicted by DAMFF based on short-term dependencies, whereas our model is better at capturing long-term dependencies. However, our method still outperforms other methods on other metrics of PMemo dataset, especially the more comprehensive CCC metric, demonstrating that our method is more effective in traditional DMER tasks.

| Model              | Arousal        |                |                   | Valence        |                |                   |
|--------------------|----------------|----------------|-------------------|----------------|----------------|-------------------|
|                    | CCC $\uparrow$ | PCC $\uparrow$ | RMSE $\downarrow$ | CCC $\uparrow$ | PCC $\uparrow$ | RMSE $\downarrow$ |
| Ours               | <b>0.377</b>   | <b>0.541</b>   | <b>0.236</b>      | <b>0.092</b>   | <b>0.154</b>   | <b>0.245</b>      |
| DAMFF <sup>†</sup> | 0.22           | 0.414          | 0.319             | 0.035          | 0.069          | 0.344             |
| DAMFF              | 0.232          | 0.478          | 0.342             | 0.035          | 0.008          | 0.39              |
| MCRNN <sup>†</sup> | 0.173          | 0.312          | 0.332             | 0.022          | 0.021          | 0.348             |
| MCRNN              | 0.176          | 0.301          | 0.364             | 0.015          | -0.002         | 0.391             |
| DNN <sup>†</sup>   | 0.138          | 0.273          | 0.35              | 0.01           | -0.009         | 0.367             |
| DNN                | 0.153          | 0.286          | 0.361             | 0.016          | -0.004         | 0.385             |
| CRNN <sup>†</sup>  | 0.113          | 0.29           | 0.31              | 0.008          | -0.009         | 0.337             |
| CRNN               | 0.131          | 0.365          | 0.35              | -0.001         | -0.021         | 0.366             |

<sup>†</sup> The models trained with MAML.

Table 2: Performance of different models in the PDMER task.

**PDMER Task** In personalized tasks, the test set requires annotation data from each annotator. Since PMemo dataset

does not include these data, so personalized tasks are validated only on the DEAM dataset. Table 2 presents the performance of all methods in personalized tasks. Compared to other baseline methods, our method demonstrates superior performance across all metrics. This can be attributed to the effectiveness of our personalized strategy. Our personalized strategy retains personalized information in the dataset, and our task construction strategy enables the model to differentiate between different personalized perceptions. When facing new personalized tasks, our method can effectively utilize previous knowledge, thereby improving model performance in the PDMER task. Notably, incorporating MAML reduces the performance of other methods, as they learn the support set’s personality during training but predict the query set with different personalized perceptions during the training phase, which makes it difficult for the models to fit.

### Subjective Experiment

To further demonstrate the effectiveness of DSAML for the PDMER task in real-world scenarios, we conducted a subjective user experiment to evaluate the accuracy of personalized emotion prediction. The experiment involved 22 participants, including 11 females and 11 males. Participants were first required to listen to a 45-second song and adjust the initial V-A curve generated by the DSAML to best match their perception. Based on this adjustment, participants then listened to 10 songs sequentially, all of which come from the test set of the DEAM dataset. After listening to each song, they were asked to rank four V-A curves by perceptual match. The four V-A curves were generated by the following methods: (1) Ground truth, (2) DSAML, (3) DSAML without personalized strategy, and (4) DAMFF. The ranking was based on the consistency between the V-A curve and the participant’s perception, with 1 indicating the most consistency and 4 indicating the least consistency.

As shown in Figure 6, we analyzed the consistency ranking of different methods in both arousal and valence values, and conducted a significance analysis using the paired t-test. In the arousal dimension, the average ranking of DSAML is at the forefront, and it has a high level of significance compared to other methods, indicating that our model can predict the most consistent personalized emotion with the participants’ perceptions. More notably, our model even outperforms the ground truth of the dataset, which to some extent indicates the importance of personalization in the real world for the DMER task. In the valence dimension, the average

| Model                      | DEAM         |              |              |              |              |              | PMemo        |              |              |              |             |              |
|----------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|-------------|--------------|
|                            | Arousal      |              |              | Valence      |              |              | Arousal      |              |              | Valence      |             |              |
|                            | CCC ↑        | PCC ↑        | RMSE ↓       | CCC ↑        | PCC ↑        | RMSE ↓       | CCC ↑        | PCC ↑        | RMSE ↓       | CCC ↑        | PCC ↑       | RMSE ↓       |
| Ours                       | <b>0.402</b> | <b>0.633</b> | <b>0.196</b> | <b>0.117</b> | 0.182        | <b>0.267</b> | <b>0.186</b> | <b>0.386</b> | 0.112        | <b>0.108</b> | <b>0.26</b> | <b>0.122</b> |
| Ours w.o. Local-Attention  | 0.385        | 0.624        | 0.215        | 0.077        | 0.134        | 0.268        | 0.167        | 0.341        | <b>0.112</b> | 0.087        | 0.213       | 0.123        |
| Ours w.o. Global-Attention | 0.37         | 0.594        | 0.207        | 0.084        | 0.148        | 0.283        | 0.157        | 0.352        | 0.123        | 0.075        | 0.214       | 0.132        |
| Ours w.o. Attention-Loss   | 0.338        | 0.564        | 0.221        | 0.103        | <b>0.185</b> | 0.279        | 0.148        | 0.353        | 0.135        | 0.098        | 0.254       | 0.139        |
| Ours w.o. Adapter          | 0.329        | 0.529        | 0.221        | 0.083        | 0.12         | 0.272        | 0.138        | 0.318        | 0.121        | 0.077        | 0.247       | 0.125        |

Table 3: Performance of ours on two datasets of the traditional DMER task.




Figure 6: Average ranking of users using different methods in subjective experiment (\*:  $p < .05$ , \*\*:  $p < .01$ , \*\*\*:  $p < .001$ ).

ranking of our method is also significantly higher than that of the other methods except the ground truth. This may be due to valence’s greater complexity, making it harder to predict than arousal. As Chua et al. pointed out, perceptions of arousal are primarily influenced by auditory information, while perceptions of valence can be influenced by both visual and auditory information (Chua et al. 2022). This indicates that the valence dimension is more challenging to predict using only audio information, which is consistent with our results in both objective and subjective experiments.

## Ablation Study

In the ablation study, we evaluated the effectiveness of each component of our model. We conducted ablation studies on both traditional DMER and PDMER tasks.

**Traditional DMER Task** Table 3 presents the performance of our model with different components removed on the DEAM and PMemo datasets. It can be observed that the overall performance of our model drops when the local attention, global attention, attention loss and Imagebind adapter are removed, indicating that these components are essential for the model’s performance. Moreover, the attention loss and Imagebind components have a more significant impact on the model’s performance, which is consistent with our design concept. The attention loss component ensures that the local and global attention focuses on different scales, while the Imagebind adapter enables the model can extract local features which ignoring by Imagebind.

**PDMER Task** Table 4 presents the performance of our model with different components removed in the PDMER

task. We can observe that the overall performance drops when the MAML or personalized task construction strategy is removed, indicating that these components are essential for the model’s performance. However, it is worth noting that the performance of the model without MAML but with the personalized task construction strategy is worse than the model without both components. This is because the personalized task construction strategy alone leads to multiple different labels for the same sample in the training set, making it difficult for the model to fit during training. This also explains why other traditional DMER methods often train the model using mean labels, as using the original labels directly will lead to a decrease in model performance.

| Model   | Arousal      |              |              | Valence      |              |              |
|---------|--------------|--------------|--------------|--------------|--------------|--------------|
|         | CCC ↑        | PCC ↑        | RMSE ↓       | CCC ↑        | PCC ↑        | RMSE ↓       |
| Ours    | <b>0.377</b> | 0.541        | <b>0.236</b> | <b>0.092</b> | <b>0.154</b> | <b>0.245</b> |
| Ours †  | 0.305        | 0.518        | 0.332        | 0.062        | 0.095        | 0.383        |
| Ours ‡  | 0.357        | 0.533        | 0.236        | 0.08         | 0.129        | 0.256        |
| Ours †‡ | 0.332        | <b>0.566</b> | 0.309        | 0.061        | 0.132        | 0.359        |

† The models trained w.o. MAML.

‡ The models trained w.o. personalized task construction strategy.

Table 4: Performance of ours in the PDMER task.

## Conclusion

This paper proposes the DSAML method for personalized dynamic music emotion recognition (PDMER). DSAML fuses features from a dual-scale feature extractor and captures both short and long-term dependencies using a dual-scale attention transformer, improving the performance in traditional DMER. Moreover, a personalized strategy is proposed, which apply a novel task construction strategy into the MAML training process. The proposed task construction strategy divides tasks by annotators, ensuring consistent perception. Leveraging this strategy alongside meta-learning, DSAML can predict personalized perception of emotions with just one personalized annotation sample. Objective experimental results demonstrate that DSAML outperforms previous music emotion recognition methods in both traditional DMER and PDMER tasks. Furthermore, subjective experiments validate the effectiveness of DSAML in real-world scenarios.

## Acknowledgments

This work is supported by the National Key Research and Development Program of China (2023YFF0904900), and the National Natural Science Foundation of China (No. 62272409).

## References

- Aljanaki, A.; Yang, Y.-H.; and Soleymani, M. 2017. Developing a benchmark for emotional analysis of music. *PLoS one*, 12(3): e0173392.
- CHAKHTOUNA, A.; SEKKATE, S.; and Abdellah, A. 2024. Modeling Speech Emotion Recognition via Image-Bind representations. *Procedia Computer Science*, 236: 428–435.
- Chen, Y.-A.; Wang, J.-C.; Yang, Y.-H.; and Chen, H. 2014. Linear regression-based adaptation of music emotion recognition models for personalization. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2149–2153. IEEE.
- Chua, P.; Makris, D.; Herremans, D.; Roig, G.; and Agres, K. 2022. Predicting emotion from music videos: exploring the relative contribution of visual and auditory information to affective responses. *arXiv preprint arXiv:2202.10453*.
- Dingle, G. A.; Kelly, P. J.; Flynn, L. M.; and Baker, F. A. 2015. The influence of music on emotions and cravings in clients in addiction treatment: a study of two clinical samples. *The Arts in Psychotherapy*, 45: 18–25.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Girdhar, R.; El-Nouby, A.; Liu, Z.; Singh, M.; Alwala, K. V.; Joulin, A.; and Misra, I. 2023. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15180–15190.
- Grigsby, J.; Wang, Z.; Nguyen, N.; and Qi, Y. 2021. Long-range transformers for dynamic spatiotemporal forecasting. *arXiv preprint arXiv:2109.12218*.
- Han, D.; Kong, Y.; Han, J.; and Wang, G. 2022. A survey of music emotion recognition. *Frontiers of Computer Science*, 16(6): 166335.
- He, N.; and Ferguson, S. 2020. Multi-view neural networks for raw audio-based music emotion recognition. In *2020 IEEE International Symposium on Multimedia (ISM)*, 168–172. IEEE.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Huang, C.-F.; and Huang, C.-Y. 2020. Emotion-based AI music generation system with CVAE-GAN. In *2020 IEEE Eurasia Conference on IOT, Communication and Engineering (ECICE)*, 220–222. IEEE.
- Ji, S.; and Yang, X. 2024. MusER: Musical Element-Based Regularization for Generating Symbolic Music with Emotion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12821–12829.
- Kang, J.; and Herremans, D. 2024. Are we there yet? A brief survey of Music Emotion Prediction Datasets, Models and Outstanding Challenges. *arXiv preprint arXiv:2406.08809*.
- Khandelwal, U.; He, H.; Qi, P.; and Jurafsky, D. 2018. Sharp nearby, fuzzy far away: How neural language models use context. *arXiv preprint arXiv:1805.04623*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Li, S.; Jin, X.; Xuan, Y.; Zhou, X.; Chen, W.; Wang, Y.-X.; and Yan, X. 2019. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *Advances in neural information processing systems*, 32.
- Liu, Z.; Xu, W.; Zhang, W.; and Jiang, Q. 2023. An emotion-based personalized music recommendation framework for emotion improvement. *Information Processing & Management*, 60(3): 103256.
- Malik, M.; Adavanne, S.; Drossos, K.; Virtanen, T.; Ticha, D.; and Jarina, R. 2017. Stacked convolutional and recurrent neural networks for music emotion recognition. *arXiv preprint arXiv:1706.02292*.
- Nichol, A.; Achiam, J.; and Schulman, J. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*.
- Orjesek, R.; Jarina, R.; Chmulik, M.; and Kuba, M. 2019. DNN based music emotion recognition from raw audio signal. In *2019 29th International Conference Radioelektronika (RADIOELEKTRONIKA)*, 1–4. IEEE.
- Russell, J. A. 1980. A circumplex model of affect. *Journal of personality and social psychology*, 39(6): 1161.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- Su, D.; and Fung, P. 2012. Personalized music emotion classification via active learning. In *Proceedings of the second international ACM workshop on Music information retrieval with user-centered and multimodal strategies*, 57–62.
- Sung, F.; Yang, Y.; Zhang, L.; Xiang, T.; Torr, P. H.; and Hospedales, T. M. 2018. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1199–1208.
- Thrun, S.; and Pratt, L. 1998. Learning to learn: Introduction and overview. In *Learning to learn*, 3–17. Springer.
- Tran, H.; Le, T.; Do, A.; Vu, T.; Bogaerts, S.; and Howard, B. 2023. Emotion-aware music recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 16087–16095.
- Wang, J.-C.; Yang, Y.-H.; Wang, H.-M.; and Jeng, S.-K. 2012. Personalized music emotion recognition via model adaptation. In *Proceedings of The 2012 Asia Pacific Signal and Information Processing Association Annual Summit and Conference*, 1–7. IEEE.
- Wang, Y.; Yao, Q.; Kwok, J. T.; and Ni, L. M. 2020. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3): 1–34.



Yang, Y.-H.; Su, Y.-F.; Lin, Y.-C.; and Chen, H. H. 2007. Music emotion recognition: The role of individuality. In *Proceedings of the international workshop on Human-centered multimedia*, 13–22.

Zhang, K.; Zhang, H.; Li, S.; Yang, C.; and Sun, L. 2018. The PMEmo dataset for music emotion recognition. In *Proceedings of the 2018 acm on international conference on multimedia retrieval*, 135–142.

Zhang, L.; Yang, X.; Zhang, Y.; and Luo, J. 2023. Dual Attention-Based Multi-Scale Feature Fusion Approach for Dynamic Music Emotion Recognition. In *ISMIR*, 207–214.

Zou, J.; Mei, J.; Ye, G.; Huai, T.; Shen, Q.; and Dong, D. 2023. EMID: An Emotional Aligned Dataset in Audio-Visual Modality. In *Proceedings of the 1st International Workshop on Multimedia Content Generation and Evaluation: New Methods and Practice*, 41–48.