

Received 8 February 2023, accepted 30 March 2023, date of publication 10 April 2023, date of current version 17 April 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3265807

RESEARCH ARTICLE

Embedding-Based Music Emotion Recognition Using Composite Loss

NAOKI TAKASHIMA^{1,4}, FRÉDÉRIC LI², MARCIN GRZEGORZEK^{2,5},
AND KIMIYUKI SHIRAHAMA^{3,6}

¹Graduate School of Science and Engineering, Kindai University, Higashiosaka, Osaka 577-8502, Japan

²Institute of Medical Informatics, University of Lübeck, 23538 Lübeck, Germany

³Faculty of Informatics, Kindai University, Higashiosaka, Osaka 577-8502, Japan

⁴Speee Inc., Minato, Tokyo 106-0032, Japan

⁵Department of Knowledge Engineering, University of Economics in Katowice, 40287 Katowice, Poland

⁶Department of Information Systems Design, Doshisha University, Kyotanabe, Kyoto 610-0394, Japan

Corresponding author: Kimiyuki Shirahama (kshiraha@mail.doshisha.ac.jp)

This work was supported in part by the Japan Society for the Promotion of Science (JSPS) within Grant-in-Aid for Scientific Research (B) under Grant 19H04172.

ABSTRACT Most music emotion recognition approaches perform classification or regression that estimates a general emotional category from a distribution of music samples, but without considering emotional variations (e.g., happiness can be further categorised into much, moderate or little happiness). We propose an embedding-based music emotion recognition approach that associates music samples with emotions in a common embedding space by considering both general emotional categories and fine-grained discrimination within each category. Since the association of music samples with emotions is uncertain due to subjective human perceptions, we compute composite loss-based embeddings obtained to maximise two statistical characteristics, one being the correlation between music samples and emotions based on canonical correlation analysis, and the other being a probabilistic similarity between a music sample and an emotion with KL-divergence. The experiments on two benchmark datasets demonstrate the effectiveness of our embedding-based approach, the composite loss and learned acoustic features. In addition, detailed analysis shows that our approach can accomplish robust bidirectional music emotion recognition that not only identifies music samples matching with a specific emotion but also detects emotions expressed in a certain music sample.

INDEX TERMS Music emotion recognition, embeddings, canonical correlation analysis, Kullback-Leibler divergence, bidirectional retrieval.

I. INTRODUCTION

Music is a powerful means to evoke human emotions. Analysing the interactions between them is thus important in affective computing, and is one of main focuses of *Music Emotion Recognition* (MER) which attempts to automatically identify the emotion matching a specific music [1]. MER is useful for many potential applications such as music recommendation and playlist generation for streaming services, and even music therapy in biomedicine [2].

MER is performed differently in the literature depending on several factors. The first one is how emotions are

modelised. Two main frameworks to define emotions currently co-exist: the *categorical* and the *continuous* ones. The former defines emotions as explicit categories, either directly using the six ‘basic’ emotions highlighted in Ekman’s theory (i.e., happiness, sadness, anger, fear, disgust and surprise) [3] or derivatives from them. The latter decomposes emotions along several axes, among which the most popular axes are arousal (level of energy) and valence (level of pleasantness) based on Russell’s Circumplex model [4]. Both of the categorical and continuous frameworks have their pros and cons. While the former can clearly identify general emotions in music, it is not the most appropriate to take into account the richness and variations of human emotions. For example, there are several degrees of happiness ranging from little to

The associate editor coordinating the review of this manuscript and approving it for publication was Thomas Canhao Xu.

intense happiness, that cannot be distinguished from each other with the categorical models. On the other hand, the continuous approach can express fine-grained human emotions in a vector space defined by the arousal and valence axes. However, it is difficult to identify general emotions because dissimilar emotions such as ‘fear’ and ‘anger’ are located close to each other in the arousal-valence space [5]. Therefore, neither categorical nor continuous approach has become predominant over the other in the literature, despite the benefits of each approach being essential for MER.

The second main difference among MER work lies in how MER is translated into a machine learning problem. The most popular approaches so far have consisted in considering MER either as a *classification* or a *regression* problem, depending on whether the categorical or continuous modelisation of emotions is used [6]. Classification and regression methods are however unidirectional, and mostly investigate *Music to Emotion* (M2E). In this framework, a classification or regression model is trained to respectively output categorical emotion estimations or arousal-valence intensity scores given some music-related input data (e.g., audio records, lyrics transcripts, playlist information, etc.). On the other hand, *Emotion to Music* (E2M) aims to retrieve some relevant music extract given some emotion-related input remains more marginal. More recently, *embedding-based retrieval* approaches have emerged to address this issue. In a retrieval problem, a model is trained to return a list of examples ordered in a descending order of similarities to an input example, also referred to as a *query* [7]. For M2E, the query and retrieved examples are respectively music-related data and emotion, while the reverse is true for E2M. Embedding-based methods on the other hand aim to project examples from various modalities into a common space referred to as an *embedding space*, so that the *embeddings* (i.e. vectorial projections) of two data examples associated to the same concept are close to each other [8]. The combination of embedding and retrieval approaches enables *bidirectional* retrieval, allowing retrieval methods to perform either E2M or M2E in the case of MER, and consequently has led to an increased interest from the research community over the past years [9], [10], [11], [12], [13], [14]. But while audio, image and text modalities are the most popular in the literature to the best of our knowledge, no other work has so far attempted to investigate bidirectional retrieval between audio and emotion ratings, except for our previous study [15].

In this paper, we propose an *Embedding-based Music Emotion Recognition* (EMER) approach that performs bidirectional retrieval based on the continuous model of emotions. Our approach can directly analyse the similarity between music and emotion in the embedding space, where an embedding designates here the vector representation of a music sample or an emotion in the embedding space. Fig. 1 shows a standard EMER approach that projects music samples and emotions into an embedding space, in such a way that associated music and emotions are close to each other in the

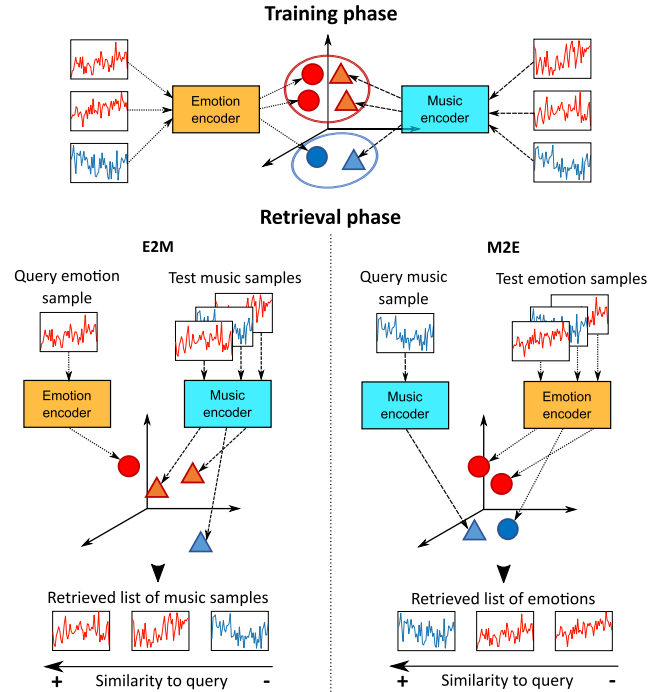


FIGURE 1. Principle of an embedding-based retrieval approach for MER. A music encoder and an emotion encoder are first jointly trained to project associated music samples and emotions close to each other in an embedding space. The trained encoders are then used to obtain embeddings of the query and test samples. The test samples are finally ordered by decreasing similarity to the query in the embedding space.

embedding space. This allows it to identify general emotions because similar emotions are gathered close to each other in terms of their embeddings. It can be noted that by projecting highly associated music samples and specific emotions in proximity, their fine-grained relations are preserved in the embedding space. This way, EMER can treat both of general and specific emotions. Once both music and emotion encoders are trained, they can be used to obtain the embeddings of a query and test samples, and rank the latter by decreasing similarities to the query in the embedding space.

One challenge when working directly with emotional ratings is that they are inherently uncertain, because they are subjectively annotated according to human perceptions which are highly influenced by many factors such as age, personality, cultural background and surrounding conditions [1], [16]. We refer to this phenomenon as *emotional uncertainty*. On the one hand, some existing MER approaches bypass this problem by extracting emotion-related information from sources that are more reliable than individual subjective reports, such as music tags [14], [17] or lyrics [12], [13]. On the other hand, many other past work directly use subjective emotion ratings without considering emotional uncertainty, and assume that the provided annotations are completely correct. This uncertainty could however cause the trained models to be either inaccurate or biased, and thus negatively impact the recognition performances.

To mitigate the impact of emotional uncertainty, we develop an approach, called *EMER using Composite Loss* (EMER-CL) that trains music and emotion embeddings with a compound loss examining two statistical characteristics that are affected in a limited way by possible inaccuracies in the emotional annotations. Firstly, we assume that even if emotional intensities differ from user to user in terms of intensity values, they remain nevertheless correlated when listening to the same music sample. Thus, Canonical Correlation Analysis (CCA) is used to devise a correlation-based loss [18]. This loss enables us to deal with inter-subject variations in emotional intensities by maximising the correlation between music samples and their associated emotions in an embedding space, so as to find their ‘relative’ connection. That is, the embedding space characterises how acoustic features change according to an increase/decrease in arousal/valence intensities and vice versa. Secondly, music samples yielding very different acoustic features can evoke similar emotions. For instance, happiness can be expressed in different genres like rock, blues and jazz. This kind of large intra-class variation in acoustic features for one emotional category makes projecting a music sample or an emotion into a single point (as shown in Fig. 1) suboptimal. Thus, we additionally project each of music samples and emotions as a probability distribution in another embedding space [19], [20]. This idea is implemented by defining a distribution-based loss that measures the Kullback-Leibler (KL) divergence between the probability distribution for a music sample and the one for an emotion in the embedding space.

Our composite loss consisting of the correlation- and distribution-based losses is necessary for managing the aforementioned inter-subject and intra-class variations resulting from the emotional uncertainty. Only using the correlation-based loss cannot cover the large intra-class variation of acoustic features, while the inter-subject variation of emotional intensities cannot be handled only with the distribution-based loss. The experimental results in Section IV-D validate the necessity of combining the correlation- and distribution-based losses.

To sum up, this paper contains the following three main contributions: Firstly, we propose EMER-CL that can work with both general and specific emotions since it uses the continuous model of emotions to obtain embeddings of emotions, and the embedding space maintains not only associated music samples and emotions close to each other but also non-associated ones far away. The embedding space serves as a bridge between music samples and emotions and offers bidirectional MER as a by-product of EMER-CL. Secondly, we propose a new composite loss combining the CCA and KL-divergence losses to take into account the emotional uncertainty. Finally, we perform extensive experiments on two benchmark datasets, MediaEval Database for Emotional Analysis in Music (DEAM) [21] and PMemo [22]. We demonstrate the effectiveness of EMER-CL over regression baselines not relying on embeddings, of the composite loss over other alternatives and of the features learned by

EMER-CL relatively to the state-of-the-art MER methods. In addition, detailed analysis of EMER-CL results reveals that reasonable recognition is robustly attained even in cases of mis-recognition.

This paper is organised as follows: Section II reviews the literature of existing MER approaches grouped into several categories. Section III details our EMER-CL and Section IV reports the experimental results demonstrating its effectiveness. Detailed analysis of recognition results by EMER-CL is conducted in Section V. Finally, Section VI presents the conclusion and our future work. In addition to these main contents, several appendices are provided to show experimental details, such as hyper-parameter tuning for EMER-CL and the comparison approaches involved in the comparative studies in Section IV, additional insights and results for the detailed analysis in Section V, and the computational cost of EMER-CL.

II. RELATED WORK

This section provides a short review of existing MER approaches by dividing them into M2E and E2M. We first review existing M2E approaches by classifying them into three categories: “feature engineering” that hand-crafts emotion-related acoustic features, “feature learning” based on deep learning that automatically learns emotion-related features, and “relation modelling” to extract the relationship between emotional intensities and acoustic features obtained by feature engineering or learning. We then discuss past work dealing with E2M. Through this review, we clarify the novelties of the proposed EMER-CL.

A. FEATURE ENGINEERING FOR M2E

Several libraries like MIRtoolbox [23] and openSMILE [24] are currently available to extract fundamental acoustic features such as Zero-Crossing Rate (ZCR), Root-Mean-Square (RMS) energy, Mel-Frequency Cepstral Coefficients (MFCCs), Short-Time Fourier Transform (STFT), etc. However, acoustic signal analysis alone might not be enough to account for all required acoustic characteristics [5]. As a result, a large focus of M2E approaches has been put on feature engineering in recent years. Panda et al. [25] distinguished several types of emotion-related acoustic features including spectral features (low-level feature), rhythm clarity (perceptual feature) and genre (high-level semantic feature). Mo and Niu [26] presented an acoustic feature extraction technique that combines three signal processing algorithms, the orthogonal matching pursuit, Gabor functions, and the Wigner distribution function, to provide an adaptive time-varying description of music signals with a higher spatial and temporal resolution. Panda et al. [27] proposed algorithms to extract acoustic features related to musical texture and expressive performance techniques (e.g., vibrato, tremolo and glissando).

The aforementioned work mainly focuses on designing acoustic features and feature selection to effectively estimate an emotion from a music sample, but it can be claimed that

feature engineering does not inherently take into account the emotional uncertainty unlike our EMER-CL approach.

B. FEATURE LEARNING FOR M2E

The advantage of feature learning is the ability to capture high-level features from raw data or hand-crafted (low-level) features. Feature learning for M2E has become a fast moving research topic due to the increasing interest in deep learning over the past decade. For this reason, we report only the most recent work (i.e. less than five-year-old) that we found related to this topic. Malik et al. [28] demonstrated the effectiveness of stacking a Convolutional Neural Network (CNN) and Recurrent Neural Network (RNN) to predict arousal/valence from acoustic features exclusively based on log mel-band energy. Dong et al. [29] developed a Bidirectional Convolutional Recurrent Sparse Network (BCRSN) that uses the spectrogram of audio signals and reduces computational complexity by converting the continuous arousal/valence prediction process to multiple binary classification problems. Sarkar et al. [30] applied a CNN taking log-mel spectrogram as input to the four-class classification problem defined by Russell's model quadrants [4]. Hizlisoy et al. [2] proposed a Convolutional Long short term memory Deep Neural Network (CLDNN) for the classification of three quadrants excluding low arousal - high valence from Russell's model quadrants [4]. Choi et al. [31] presented a transfer learning approach where a CNN taking mel-spectrograms as input is firstly trained for a music tagging task, and then transferred and fine-tuned for six other tasks such as music genre classification, speech/music classification, emotion prediction etc. Koh and Dubnov [32] presented a comparison of state-of-the-art deep feature learning architectures including VGGish and L^3 -Net that take audio spectrograms as inputs. The two models outperformed MFCCs features on various M2E datasets for either classification or regression tasks. Orjesek et al. [33] proposed two deep-learning-based architectures to learn features for M2E as a regression problem. The first one is based on a CNN stacked with a bidirectional Gated Recurrent Unit (GRU) and Multi-Layer Perceptron (MLP). The other has the same architecture, except that an autoencoder is inserted between the CNN and bidirectional GRU. The ensemble is trained by adding a reconstruction term to the loss function. He and Ferguson [34] proposed a two-stage approach for the classification of low/high arousal and valence. Log-mel spectrograms obtained from the raw audio signals are first used to train a convolutional autoencoder. The encoder is then used to train two bidirectional Long Short-term Memory (LSTM), one for arousal and the other for valence classification.

The aforementioned approaches do not take into account the emotional uncertainty. On the other hand, our EMER-CL approach uses high-level acoustic features learned by a pre-trained VGGish model [35], and takes advantage of the composite loss to deal with the emotional uncertainty.

C. RELATION MODELLING FOR M2E

Past work has also investigated relationships between music samples and emotions, although the lower popularity of this topic in MER research compared to feature engineering or feature learning means that this field is moving at a slower pace. Yang et al. [36] built a group-wise MER scheme (GWMER) which divides users into various groups based on user information such as generation, gender, occupation and personality, and trains a Support Vector Regression (SVR) for the prediction of arousal/valence for each group. GWMER can this way partially address the problem that continuous emotions are more affected by subjective issue than discrete emotions when annotating. Yang and Chen [37] presented a ranking-based neural network model that ranks a collection of music samples by emotion and determines the emotional intensity of each music sample. Yang and Chen [38] and Chin et al. [39] developed probabilistic approaches to deal with the emotional uncertainty by estimating the distribution of emotional intensities from hand-crafted acoustic features. Markov and Matsui [40] showed that modelling with Gaussian Processes (GP) was more powerful than SVR for arousal/valence regression with hand-crafted acoustic features. Fukayama and Goto [41] evaluated the effectiveness of aggregating multiple GP regressions, each trained with different acoustic features. Wang et al. [42] presented Acoustic Emotion Gaussians (AEGs) that treat the emotional uncertainty by modelling hand-crafted acoustic features as a parametric probability distribution (soft assignment) instead of a single point (hard assignment). Wang et al. [43] proposed a Histogram Density Mixture (HDM) model that quantises the arousal/valence space into cells and extracts latent histograms representing characteristic emotion distributions over cells based on hand-crafted acoustic features. Wang et al. [44] developed a MER system for 34 emotional categories based on Hierarchical Dirichlet Process Mixture Model (HDPMM) that links emotion classes using the property of sharing components in the HDPMM.

To the best of our knowledge, relation modelling approaches have so far exclusively relied on feature engineering, and not yet been used with high-level features obtained by feature learning. On the other hand, our EMER-CL approach uses high-level acoustic features based on VGGish, and projects emotional intensities into an embedding space, which enables us to extract high-level feature representations for emotions.

D. E2M

E2M has not been explored as extensively as M2E, especially not in the recent MER literature. One possible reason for this scarcity is the fact that the existing acoustic features associated with emotions are high dimensional, and thus not easy to predict directly from emotions using traditional machine learning methods like regression. Except for studies older than a decade [45], [46], [47], [48], we could find only a single recent method incorporating E2M elements proposed

by Deng et al. [47]. In this work, a music recommendation method taking into account the emotions of the user is proposed. An M2E model is first trained using a classification framework to predict the emotional state associated with a music sample. The model is then used to predict an emotional state sequence containing the emotions associated with the last songs listened by the user. A model based on Conditional Random Fields is used to predict the user’s current emotion based on this emotional state sequence. Finally, the similarities between the predicted current emotion and the ones associated to songs of the dataset are computed to retrieve relevant songs to be suggested.

In this paper, we propose a new bidirectional MER approach able to perform either M2E and E2M based on projecting music samples and associated emotions in proximity in two embeddings spaces. Unlike Deng et al., our approach can directly perform E2M without relying on a M2E system.

E. EMBEDDING-BASED RECOGNITION

Approaches in this category have attracted much attention as techniques that can perform effective bidirectional recognition between different modalities (e.g., image, text and audio). When it comes to approaches involving audio modality, the most common investigations include embedding-based recognition between audio and image [9], [10], between audio and text (lyrics) [11], [12], or even between all three of image, audio and text [13]. Some attempts have also focused on extracting meaningful embeddings from music meta-data (e.g. genre, instrument, mood/theme) and playlist information [14]. However to the best of our knowledge, no existing work addresses embedding-based recognition between audio and emotion except our previous study [15], where MLPs trained with the CCA loss are used to compute embeddings of music samples and emotions. This paper is an extension of our previous study by adopting RNNs in addition to MLPs, devising a composite loss that combines the CCA and KL-divergence losses, and conducting significantly deeper analysis of experimental results.

III. EMER-CL APPROACH

Fig. 2 shows an overview of our EMER-CL approach. First, a music sample is converted into a sequence of acoustic features $X^{(m)} = x_1^{(m)}, \dots, x_{T_m}^{(m)}$ of length T_m . Here, $x_t^{(m)} \in \mathbb{R}^{D_x^{(m)}}$ is a $D_x^{(m)}$ -dimensional feature vector at time t ($1 \leq t \leq T_m$). In our implementation, VGGish¹ produced by Google [35] is used to segment every music sample recorded with a sampling rate of 44.1kHz into segments of 0.96 seconds, and then $x_t^{(m)}$ is extracted as a 128-dimensional vector (i.e., $D_x^{(m)} = 128$) from each segment. The emotion associated with the music sample is represented as a sequence $X^{(e)} = x_1^{(e)}, \dots, x_{T_e}^{(e)}$ of length T_e where $x_t^{(e)} \in \mathbb{R}^{D_x^{(e)}}$ is

¹<https://github.com/tensorflow/models/tree/master/research/audioset/vggish>

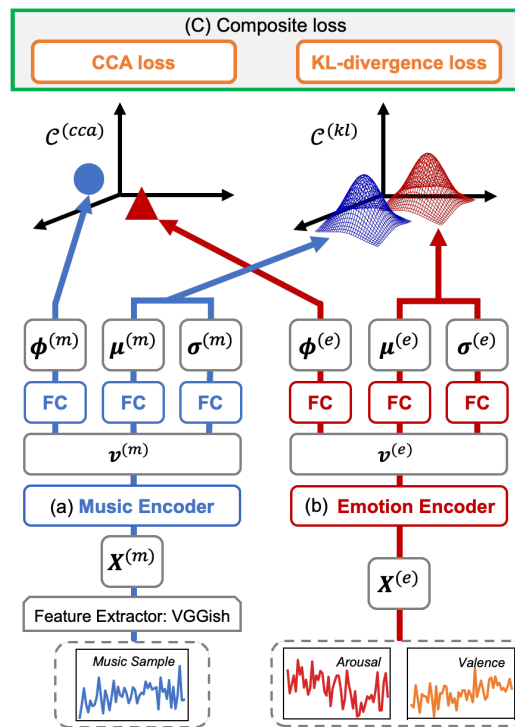


FIGURE 2. An overview of our EMER-CL approach. $X^{(m)}$ and $X^{(e)}$ respectively designate the music and emotion sequences input to the music and emotion encoders. $v^{(m)}$ and $v^{(e)}$ are the vectorial outputs of the music and emotion encoders, respectively. Branches of fully connected (FC) layers project $v^{(*)}$ into a point-based embedding $\phi^{(*)}$ in the space $\mathcal{C}^{(cca)}$, and a probabilistic-based embedding following a multivariate Gaussian distribution $\mathcal{N}(\mu^{(*)}, \sigma^{(*)})$ in the space $\mathcal{C}^{(kl)}$, for $* \in \{m, e\}$. Both music and emotion models are jointly trained to minimise a composite loss computing the CCA loss between $\phi^{(m)}$ and $\phi^{(e)}$, and a KL-divergence loss between $\mathcal{N}(\mu^{(m)}, \sigma^{(m)})$ and $\mathcal{N}(\mu^{(e)}, \sigma^{(e)})$.

a $D_x^{(e)}$ -dimensional vector containing characteristics of the emotion at time t ($1 \leq t \leq T_e$). In our setting, $x_t^{(e)}$ is defined as a two-dimensional vector (i.e., $D_x^{(e)} = 2$) indicating arousal and valence intensities recorded with a sampling rate of 2Hz. Unlike $X^{(m)}$, we do not perform feature extraction on the raw arousal/valence intensities and use the latter directly as $X^{(e)}$. This is because we consider that only two types of intensities recorded with a low sampling rate have relatively simple characteristics. Note that for simplicity $X^{(e)}$ is called an *arousal/valence sequence* in the following discussions.

It is cumbersome to directly project $X^{(m)}$ and $X^{(e)}$ into an embedding space because they are sequences of different lengths T_m and T_e . Thus, as shown in Fig. 2 (a) and (b), music and emotion encoders are used to transform $X^{(m)}$ and $X^{(e)}$ into vectors $v^{(m)} \in \mathbb{R}^{D_v^{(m)}}$ and $v^{(e)} \in \mathbb{R}^{D_v^{(e)}}$, respectively. Each of $v^{(m)}$ and $v^{(e)}$ is a high-level feature that effectively summarises the features in $X^{(m)}$ or $X^{(e)}$ and their temporal relations. We use either an MLP or RNN based on bidirectional Gated Recurrent Unit (GRU) [49], [50] as music and emotion encoders as described in Section III-A.

Then, $v^{(m)}$ and $v^{(e)}$ are projected into an embedding space $\mathcal{C}^{(cca)}$ of dimensionality $D^{(cca)}$ using different Fully Connected (FC) layers with linear activation. The embeddings

for $\mathbf{v}^{(m)}$ and $\mathbf{v}^{(e)}$ in $\mathcal{C}^{(cca)}$ are denoted by $\boldsymbol{\phi}^{(m)} \in \mathbb{R}^{D^{(cca)}}$ and $\boldsymbol{\phi}^{(e)} \in \mathbb{R}^{D^{(cca)}}$, respectively. In addition, two branches of FC layers are used to transform $\mathbf{v}^{(m)}$ into a mean vector $\boldsymbol{\mu}^{(m)} \in \mathbb{R}^{D^{(kl)}}$ and a covariance matrix $\boldsymbol{\Sigma}^{(m)} \in \mathbb{R}^{D^{(kl)} \times D^{(kl)}}$. This defines an additional embedding for $\mathbf{v}^{(m)}$ as a multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}^{(m)}, \boldsymbol{\Sigma}^{(m)})$ in another embedding space $\mathcal{C}^{(kl)}$ of dimensionality $D^{(kl)}$. Considering the expensive computational cost to process multivariate Gaussian distributions, we assume that each dimension is independent based on the standard practice of the literature [51], [52]. Thus, $\mathcal{N}(\boldsymbol{\mu}^{(m)}, \boldsymbol{\Sigma}^{(m)})$ is reduced to $\mathcal{N}(\boldsymbol{\mu}^{(m)}, \boldsymbol{\sigma}^{(m)})$ by replacing $\boldsymbol{\Sigma}^{(m)}$ with the variance vector $\boldsymbol{\sigma}^{(m)} \in \mathbb{R}^{D^{(kl)}}$ representing the variance in each dimension. Similarly, $\mathbf{v}^{(e)}$ is converted into $\mathcal{N}(\boldsymbol{\mu}^{(e)}, \boldsymbol{\sigma}^{(e)})$ in $\mathcal{C}^{(kl)}$ using two branches of FC layers. Under the above-mentioned setting, EMER-CL trains the music and emotion encoders and the six branches of FC layers by jointly minimising the CCA loss between $\boldsymbol{\phi}^{(m)}$ and $\boldsymbol{\phi}^{(e)}$ and the KL-divergence loss between $\mathcal{N}(\boldsymbol{\mu}^{(m)}, \boldsymbol{\sigma}^{(m)})$ and $\mathcal{N}(\boldsymbol{\mu}^{(e)}, \boldsymbol{\sigma}^{(e)})$, as illustrated in Fig. 2 (c).

The following sections describe encoding of music samples and emotions, and more details of the training process. The dimensionalities $D_v^{(m)}$, $D_v^{(e)}$, $D^{(cca)}$ and $D^{(kl)}$ are hyper-parameters of EMER-CL whose specific values are provided in Section IV-C.

A. MUSIC AND EMOTION ENCODERS

For the music encoder, two types of neural networks were tested: an MLP and an RNN using bidirectional GRU [49], [50]. When the former is used, a mean feature vector $\bar{\mathbf{x}}^{(m)}$ is computed by averaging $\mathbf{x}_1^{(m)}, \dots, \mathbf{x}_{T_m}^{(m)}$ in $\mathbf{X}^{(m)}$, and fed into the MLP which performs several non-linear transformations on $\bar{\mathbf{x}}^{(m)}$ to output $\mathbf{v}^{(m)}$. The RNN using bidirectional GRU computes two types of $D_h^{(m)}$ -dimensional hidden states $\vec{\mathbf{h}}_t^{(m)} \in \mathbb{R}^{D_h^{(m)}}$ and $\overleftarrow{\mathbf{h}}_t^{(m)} \in \mathbb{R}^{D_h^{(m)}}$ that represent temporal characteristics of $\mathbf{X}^{(m)}$ in the forward and backward directions, respectively. Roughly speaking, $\vec{\mathbf{h}}_t^{(m)}$ is computed by recursively aggregating $\mathbf{x}_t^{(m)}$ at the current time t and $\vec{\mathbf{h}}_{t-1}^{(m)}$ obtained at the previous time $t - 1$. Letting f be a function for the recursive aggregation, $\vec{\mathbf{h}}_t^{(m)}$ is described as $\vec{\mathbf{h}}_t^{(m)} = f(\mathbf{x}_t^{(m)}, \vec{\mathbf{h}}_{t-1}^{(m)})$ [49]. In contrast, $\overleftarrow{\mathbf{h}}_t^{(m)}$ in the backward direction is computed by aggregating $\mathbf{x}_t^{(m)}$ and $\overleftarrow{\mathbf{h}}_{t+1}^{(m)}$ at the next time $t + 1$, that is, $\overleftarrow{\mathbf{h}}_t^{(m)} = f(\mathbf{x}_t^{(m)}, \overleftarrow{\mathbf{h}}_{t+1}^{(m)})$. The vector obtained by concatenating $\vec{\mathbf{h}}_{T_m}^{(m)}$ and $\overleftarrow{\mathbf{h}}_1^{(m)}$ is a high-level feature expressing bidirectional temporal characteristics in $\mathbf{X}^{(m)}$, and is fed into the FC layers to produce a higher-level feature as $\mathbf{v}^{(m)}$.

Our preliminary experiments showed the effectiveness of an RNN using bidirectional GRU as the emotion encoder regardless of datasets. We hypothesise that this is due the fact that bidirectional GRUs can capture the best the variations in arousal and valence levels during the listening of a music sample, and therefore take this information into account to produce meaningful embeddings. Therefore, a feature vector $\mathbf{v}^{(e)}$ for an arousal/valence sequence $\mathbf{X}^{(e)}$ is extracted the same

way as $\mathbf{v}^{(m)}$ when an RNN is used as the music encoder. Finally, the specific values of hyper-parameters like $D_h^{(m)}$, $D_h^{(e)}$ for the RNN-based emotion encoder, and the configuration of FC layers are provided in Section IV-C.

B. TRAINING WITH THE COMPOSITE LOSS

Let $\mathcal{X} = \{(\mathbf{X}_n^{(m)}, \mathbf{X}_n^{(e)})\}_{n=1}^N$ be a batch consisting of N pairs of an acoustic feature sequence and an arousal/valence sequence for associated music samples and emotions. And, $\mathcal{V} = \{(\mathbf{v}_n^{(m)}, \mathbf{v}_n^{(e)})\}_{n=1}^N$ is a set of feature pairs obtained by feeding $(\mathbf{X}_n^{(m)}, \mathbf{X}_n^{(e)}) \in \mathcal{X}$ to the music and emotion encoders. The CCA loss $CCA(\mathcal{F})$ [18] is computed by converting \mathcal{V} into a set of embeddings $\mathcal{F} = \{(\boldsymbol{\phi}_n^{(m)}, \boldsymbol{\phi}_n^{(e)})\}_{n=1}^N$. In addition, the KL-divergence loss $KL(\mathcal{G})$ [19] is calculated by transforming \mathcal{V} into a set of pairs of multivariate Gaussian distributions $\mathcal{G} = \{(\mathcal{N}(\boldsymbol{\mu}_n^{(m)}, \boldsymbol{\sigma}_n^{(m)}), \mathcal{N}(\boldsymbol{\mu}_n^{(e)}, \boldsymbol{\sigma}_n^{(e)}))\}_{n=1}^N$. Our composite loss $CL(\mathcal{F}, \mathcal{G})$ combines $CCA(\mathcal{F})$ and $KL(\mathcal{G})$ as follows:

$$CL(\mathcal{F}, \mathcal{G}) = \lambda \cdot CCA(\mathcal{F}) + (1 - \lambda) \cdot KL(\mathcal{G}) \quad (1)$$

where $\lambda \in [0, 1]$ is a weight parameter to balance $CCA(\mathcal{F})$ and $KL(\mathcal{G})$. Details of $CCA(\mathcal{F})$ and $KL(\mathcal{G})$ are described in the following Sections III-B1 and III-B2.

1) CORRELATION-BASED EMBEDDING WITH CCA

The CCA loss is employed to construct an embedding space $\mathcal{C}^{(cca)}$ where $\{\boldsymbol{\phi}_n^{(m)}\}_{n=1}^N$ and $\{\boldsymbol{\phi}_n^{(e)}\}_{n=1}^N$ are strongly correlated. More specifically, for each dimension of $\mathcal{C}^{(cca)}$, music and emotion embeddings are linearly correlated with each other regardless of their actual values. This linear correlation indicates what change in acoustic features (or emotions) would be associated to the corresponding change in emotions (or acoustic features) for each of these dimensions. This allows us to characterise the relationship between music samples and emotions in a way that is independent from the actual emotion intensities attributed by individuals, which is useful for addressing the inter-subject variations described in Section I. To extract embeddings capturing complex correlations between music samples and emotions, FC layers are firstly used to refine $\mathbf{v}_n^{(m)}$ and $\mathbf{v}_n^{(e)}$ into a $D_z^{(m)}$ -dimensional vector $\mathbf{z}_n^{(m)} \in \mathbb{R}^{D_z^{(m)}}$ and a $D_z^{(e)}$ -dimensional vector $\mathbf{z}_n^{(e)} \in \mathbb{R}^{D_z^{(e)}}$, respectively. The CCA loss is computed on $\mathbf{z}_n^{(m)}$ and $\mathbf{z}_n^{(e)}$.

Let $\mathbf{z}^{(m)}$ be a random vector that is sampled from the probability distribution estimated using a set of N samples $\{\mathbf{z}_n^{(m)}\}_{n=1}^N$, and $\mathbf{z}^{(e)}$ be a random vector from the probability distribution estimated using $\{\mathbf{z}_n^{(e)}\}_{n=1}^N$. In addition, let us assume that $\mathbf{w}^{(m)} \in \mathbb{R}^{D_z^{(m)}}$ and $\mathbf{w}^{(e)} \in \mathbb{R}^{D_z^{(e)}}$ are $D_z^{(m)}$ - and $D_z^{(e)}$ -dimensional weight vectors to project $\mathbf{z}^{(m)}$ and $\mathbf{z}^{(e)}$ into scalars, respectively. CCA optimises $\mathbf{w}^{(m)}$ and $\mathbf{w}^{(e)}$ so as to maximise the following correlation between $\mathbf{w}^{(m)T} \mathbf{z}^{(m)}$ and $\mathbf{w}^{(e)T} \mathbf{z}^{(e)}$ [18]:

$$\begin{aligned} & \rho(\mathbf{w}^{(m)T} \mathbf{z}^{(m)}, \mathbf{w}^{(e)T} \mathbf{z}^{(e)}) \\ &= \frac{\mathbf{w}^{(m)T} \boldsymbol{\Sigma}^{(me)} \mathbf{w}^{(e)}}{\sqrt{\mathbf{w}^{(m)T} \boldsymbol{\Sigma}^{(mm)} \mathbf{w}^{(m)}} \sqrt{\mathbf{w}^{(e)T} \boldsymbol{\Sigma}^{(ee)} \mathbf{w}^{(e)}}} \quad (2) \end{aligned}$$

where $\Sigma^{(me)} \in \mathbb{R}^{D_z^{(m)} \times D_z^{(e)}}$ is the cross-covariance matrix computed from $\{(z_n^{(m)}, z_n^{(e)})\}_{n=1}^N$ and $\Sigma^{(mm)} \in \mathbb{R}^{D_z^{(m)} \times D_z^{(m)}}$ and $\Sigma^{(ee)} \in \mathbb{R}^{D_z^{(e)} \times D_z^{(e)}}$ are the covariance matrices for $\{z_n^{(m)}\}_{n=1}^N$ and $\{z_n^{(e)}\}_{n=1}^N$, respectively. In Eq. (2) the quantity to maximise is invariant in scaling of $w^{(m)}$ and $w^{(e)}$, so it is possible to focus on the problem where the denominator is equal to 1. In other words, the objective of CCA is to maximise the numerator in Eq. (2) subject to the constraints $w^{(m)T} \Sigma^{(mm)} w^{(m)} = 1$ and $w^{(e)T} \Sigma^{(ee)} w^{(e)} = 1$.

The CCA approach described above can be re-applied independently on each dimension of $C^{(cca)}$. For this, $D^{(cca)}$ pairs of weight vectors $\{(w_d^{(m)}, w_d^{(e)})\}_{d=1}^{D^{(cca)}}$ are found to maximise the sum of correlations between $w_d^{(m)T} z^{(m)}$ and $w_d^{(e)T} z^{(e)}$. In other words, letting $W^{(m)} \in \mathbb{R}^{D_z^{(m)} \times D^{(cca)}}$ be a matrix where each column is $w_d^{(m)}$, $W^{(m)T} z^{(m)}$ forms a $D^{(cca)}$ -dimensional embedding $\phi^{(m)}$. Similarly, $W^{(e)} \in \mathbb{R}^{D_z^{(e)} \times D^{(cca)}}$ where each column is $w_d^{(e)}$ is defined to create $\phi^{(e)} = W^{(e)T} z^{(e)}$. From this perspective, the general CCA maximises the sum of correlations each computed for one dimension of $\phi^{(m)}$ and $\phi^{(e)}$. Past work has shown that the batch optimisation of $\{(w_d^{(m)}, w_d^{(e)})\}_{d=1}^{D^{(cca)}}$ can be done by solving the following constrained optimisation problem [18]:

$$\begin{aligned} & \text{minimise: } CCA(\mathcal{F}) = -\text{tr}(W^{(m)T} \Sigma^{(me)} W^{(e)}) \\ & \text{subject to: } W^{(m)T} \Sigma^{(mm)} W^{(m)} = W^{(e)T} \Sigma^{(ee)} W^{(e)} = \mathbf{I} \quad (3) \end{aligned}$$

where the trace operation (tr) is used to sum up the correlations on each dimension of $\phi^{(m)}$ and $\phi^{(e)}$. After obtaining the optimal $W^{(m)*}$ and $W^{(e)*}$, the correlation-based embedding for $z_n^{(m)}$ of a music sample and the one for $z_n^{(e)}$ of an emotion are computed as $\phi_n^{(m)} = W^{(m)*T} z_n^{(m)}$ and $\phi_n^{(e)} = W^{(e)*T} z_n^{(e)}$, respectively.

2) DISTRIBUTION-BASED EMBEDDING WITH KL-DIVERGENCE

The CCA loss analyses only the relation (correlation) between music samples and their associated emotions, but neither the relation of a music sample to non-associated emotions nor the relation of an emotion to non-associated music samples. In addition, $\phi_n^{(m)}$ and $\phi_n^{(e)}$ are points in $C^{(cca)}$, which is unsuitable for managing the large intra-class variation of acoustic features, as discussed in Section I. To address these issues, the KL-divergence loss $KL(\mathcal{G})$ is used to build an embedding space $C^{(kl)}$ that attempts to fulfil the following conditions: 1) a music sample and its associated emotion are projected as multivariate Gaussian distributions $\mathcal{N}(\mu_n^{(m)}, \sigma_n^{(m)})$ and $\mathcal{N}(\mu_n^{(e)}, \sigma_n^{(e)})$ which are similar to each other; 2) a music sample and its non-associated emotion are projected as dissimilar distributions $\mathcal{N}(\mu_n^{(m)}, \sigma_n^{(m)})$ and $\mathcal{N}(\mu_{n'}^{(e)}, \sigma_{n'}^{(e)})$ ($n \neq n'$). Similarly, an emotion and its non-associated music sample are transformed into dissimilar distributions $\mathcal{N}(\mu_n^{(e)}, \sigma_n^{(e)})$ and $\mathcal{N}(\mu_{n'}^{(m)}, \sigma_{n'}^{(m)})$.

For simplicity, $\mathcal{N}(\mu_n^{(m)}, \sigma_n^{(m)})$ and $\mathcal{N}(\mu_n^{(e)}, \sigma_n^{(e)})$ are abbreviated into $\mathcal{N}_n^{(m)}$ and $\mathcal{N}_n^{(e)}$, respectively. In addition, we use the term *positive pair* to indicate a pair of $\mathcal{N}_n^{(m)}$ and $\mathcal{N}_n^{(e)}$

obtained for a music sample and its associated emotion. On the other hand, a *negative pair* expresses a pair of $\mathcal{N}_n^{(m)}$ and $\mathcal{N}_{n'}^{(e)}$ or a pair of $\mathcal{N}_{n'}^{(m)}$ and $\mathcal{N}_n^{(e)}$, consisting of non-associated music sample and emotion. Note that by following the standard of embedding-based retrieval [19], [53], we consider that a music sample and an emotion whose indices are the same form a positive pair, and any other pair is a negative pair.

The aforementioned conditions for $C^{(kl)}$ can be formulated using a triplet $(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)}, \mathcal{N}_{n'}^{(e)})$:

$$\psi(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)}) \leq \alpha + \psi(\mathcal{N}_n^{(m)}, \mathcal{N}_{n'}^{(e)}), \quad (4)$$

where $\psi(\cdot, \cdot)$ represents a distance between two distributions. Additionally, $\alpha > 0$ is a margin hyper-parameter which determines how far the difference between the distance for the positive pair $(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)})$ and the one for the negative pair $(\mathcal{N}_n^{(m)}, \mathcal{N}_{n'}^{(e)})$ is allowed to be. Eq. (4) uses $\mathcal{N}_n^{(m)}$ as an anchor and checks whether its distance to $\mathcal{N}_n^{(e)}$ is sufficiently smaller than its distance to $\mathcal{N}_{n'}^{(e)}$. Similarly, another triplet $(\mathcal{N}_n^{(e)}, \mathcal{N}_n^{(m)}, \mathcal{N}_{n'}^{(m)})$ can define the distance condition using $\mathcal{N}_n^{(e)}$ as an anchor:

$$\psi(\mathcal{N}_n^{(e)}, \mathcal{N}_n^{(m)}) \leq \alpha + \psi(\mathcal{N}_n^{(e)}, \mathcal{N}_{n'}^{(m)}). \quad (5)$$

For each positive pair $(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)})$ in \mathcal{G} , $KL(\mathcal{G})$ examines the distance conditions in Eqs. 4 and 5. Specifically, the following ranking loss $r(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)})$ is defined by combining two hinge losses as follows:

$$\begin{aligned} r(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)}) &= \sum_{n' \neq n} \max \left\{ 0, \psi(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)}) - \alpha - \psi(\mathcal{N}_n^{(m)}, \mathcal{N}_{n'}^{(e)}) \right\} \\ &+ \sum_{n' \neq n} \max \left\{ 0, \psi(\mathcal{N}_n^{(e)}, \mathcal{N}_n^{(m)}) - \alpha - \psi(\mathcal{N}_n^{(e)}, \mathcal{N}_{n'}^{(m)}) \right\}, \end{aligned} \quad (6)$$

where the first term becomes zero if all the negative pairs defined using $\mathcal{N}_n^{(m)}$ as an anchor lead to distances that are greater than the distance between the positive pair by more than α . The second term also checks a similar distance condition using $\mathcal{N}_n^{(e)}$ as an anchor. This way $r(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)})$ indicates how small $\psi(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)})$ is relatively to $\psi(\mathcal{N}_n^{(m)}, \mathcal{N}_{n'}^{(e)})$ and $\psi(\mathcal{N}_{n'}^{(m)}, \mathcal{N}_n^{(e)})$ defined for all the negative pairs.

To compute $r(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)})$, the KL-divergence is employed as a distance $\psi(\cdot, \cdot)$ between two multivariate Gaussian distributions $\mathcal{N}(\mu_1, \sigma_1)$ and $\mathcal{N}(\mu_2, \sigma_2)$ in $C^{(kl)}$ of dimensionality $D^{(kl)}$, and is computed as follows

$$\begin{aligned} \psi(\mathcal{N}(\mu_1, \sigma_1), \mathcal{N}(\mu_2, \sigma_2)) &= \frac{1}{2} \left\{ \sum_{d=1}^{D^{(kl)}} \frac{\sigma_{1,d}}{\sigma_{2,d}} - D^{(kl)} + \ln \frac{\prod_{d=1}^{D^{(kl)}} \sigma_{2,d}}{\prod_{d=1}^{D^{(kl)}} \sigma_{1,d}} \right. \\ &\quad \left. + \sum_{d=1}^{D^{(kl)}} \frac{(\mu_{2,d} - \mu_{1,d})^2}{\sigma_{2,d}} \right\}, \end{aligned} \quad (7)$$

where μ_1 and σ_1 are expanded as $(\mu_{1,1}, \dots, \mu_{1,D^{(kl)}})^T$ and $(\sigma_{1,1}, \dots, \sigma_{1,D^{(kl)}})^T$, respectively. Similarly, μ_2 and σ_2 are also expanded.

Finally, $KL(\mathcal{G})$ is defined as the sum of $r(\mathcal{N}_n^{(m)}, \mathcal{N}_n^{(e)})$ for all the positive pairs in \mathcal{G} . The minimisation of $KL(\mathcal{G})$ can therefore lead both music and emotion encoders and FC layers to learn parameters so that the KL-divergence between each positive pair is minimised, while maximising the KL-divergence between each negative pair.

C. TESTING EMER-CL IN M2E AND E2M

We evaluate EMER-CL in the framework of M2E and E2M that are formulated as a retrieval task. In the following paragraphs, q designates the index of a query, with $1 \leq q \leq Q$ where Q is the number of examples in the test set. In M2E, the query is $X_q^{(m)}$ which represents a music sample that is associated with an emotion rating $X_q^{(e)}$. The trained music model (consisting in the music encoder and three branches of FC layers) is used to encode the query music sample $X_q^{(m)}$ into its correlation-based embedding $\phi_q^{(m)}$ and distribution-based embedding $\mathcal{N}_q^{(m)}$. Similarly, the trained emotion model (consisting in the emotion encoder and three branches of FC layers) is used to convert the j th test emotion $X_j^{(e)}$ (for $1 \leq j \leq Q$) into its correlation-based embedding $\phi_j^{(e)}$ and distribution-based embedding $\mathcal{N}_j^{(e)}$. The similarity $s(X_q^{(m)}, X_j^{(e)})$ between the query music sample and the j th test emotion is computed as follows:

$$s(X_q^{(m)}, X_j^{(e)}) = \lambda \cdot \Gamma(\phi_q^{(m)}, \phi_j^{(e)}) + (1 - \lambda) \cdot \psi(\mathcal{N}_q^{(m)}, \mathcal{N}_j^{(e)}), \quad (8)$$

where $\lambda \in [0, 1]$ is the same weighting parameter used for the loss in Eq. (1), ψ designates the negative KL-divergence, and Γ is a correlation-based similarity between $\phi_q^{(m)} = (\phi_{q,1}^{(m)}, \dots, \phi_{q,D^{(cca)}}^{(m)})^T$ and $\phi_j^{(e)} = (\phi_{j,1}^{(e)}, \dots, \phi_{j,D^{(cca)}}^{(e)})^T$.

To determine a suitable correlation-based similarity Γ , we proceeded as follows. As a reminder, the CCA loss is designed to maximise the correlation on each of $D^{(cca)}$ dimensions independently. The higher the correlation ρ_d on the d -th dimension ($1 \leq d \leq D^{(cca)}$) is, the more linearly aligned embedding values for music samples and emotions in training data (i.e., $\{\phi_{n,d}^{(m)} = \mathbf{w}_d^{(m)T} \mathbf{z}_n^{(m)}\}_{n=1}^N$ and $\{\phi_{n,d}^{(e)} = \mathbf{w}_d^{(e)T} \mathbf{z}_n^{(e)}\}_{n=1}^N$) are. Based on this, linear regression is performed to extract the approximation line $g_d(\cdot)$ that takes as input $\phi_{q,d}$ being the embedding value for the query music sample on the d th dimension and outputs $g_d(\phi_{q,d}^{(m)})$ being an approximate embedding value for the associated emotion. Thus, the negative of the absolute difference between $g_d(\phi_{q,d}^{(m)})$ and the embedding value $\phi_{j,d}^{(e)}$ for the test emotion is defined as its similarity to the query music sample on the d th dimension. By summing up such similarities on all the $D^{(cca)}$ dimensions,

$\Gamma(\phi_q^{(m)}, \phi_j^{(e)})$ is computed as follows:

$$\Gamma(\phi_q^{(m)}, \phi_j^{(e)}) = \sum_{d \in \{d' | 1 \leq d' \leq D^{(cca)}, \rho_{d'} \geq P\}} -\rho_d \left| g_d(\phi_{q,d}^{(m)}) - \phi_{j,d}^{(e)} \right| \quad (9)$$

where each dimension d' is filtered out or weighted by the correlation $\rho_{d'}$. If $\rho_{d'}$ is lower than the threshold P , the approximation on the d' th dimension is regarded as inaccurate and the similarity on this dimension is not counted. In contrast, as $\rho_{d'}$ becomes higher, the approximation is regarded as more accurate and the similarity is more prioritised by weighting it with $\rho_{d'}$.

The similarities between the query music and all test emotions $s(X_q^{(m)}, X_j^{(e)})$ are computed for all $1 \leq j \leq Q$ using Eq. (8), and then sorted by decreasing similarities. The performance of M2E is evaluated by examining whether the test emotion associated with the query music sample $X_q^{(e)}$ is ranked at a high position in the sorted output list.

Similarly to M2E, E2M is performed by encoding a query emotion $X_q^{(e)}$ and test music samples $X_j^{(m)}$ with the trained emotion and music models, respectively. Then, the Q test music samples are sorted by computing their similarities to the query emotion $s(X_q^{(e)}, X_j^{(m)})$ for $1 \leq j \leq Q$ according to Eq. (8). The rank of the music sample associated with the query emotion is checked to measure the performance of E2M.

IV. EXPERIMENTS

In this section, we evaluate EMER-CL on two datasets: MediaEval Database for Emotional Analysis in Music (DEAM) [21] and PMemo [22]. We first present an overview and pre-processing on each dataset, the evaluation metrics, and the implementation details. Then, we present the results of three experiments. The first is an ablation study to validate the composite loss, the second compares EMER-CL to regression baselines not relying on embeddings, and the last examines the generality of EMER-CL by comparing it to the state-of-the-art MER methods.

A. DATASETS

DEAM [21]² provides 1802 music samples that are free audio source records, and their corresponding arousal/valence sequences where arousal and valence intensities lie in $[-10, 10]$. Each music sample was annotated with arousal and valence intensities every 0.5 seconds by at least 5 subjects recruited using Amazon's crowdsourcing platform *Mechanical Turk*. These intensities were projected into the range $[-1, 1]$ for each subject. DEAM contains 1744 45-second-long music samples and 58 samples that have durations longer than 45 seconds. The authors of DEAM decided to discard the first 15 seconds of annotations after observing high instabilities due to a high variance in how music samples start. Because of this and the fact that most music samples

²<https://cvml.unige.ch/databases/DEAM>

last only 45 seconds, each music sample is normalised to have a length of 30 seconds by taking the segment starting at 15 seconds and ending at 45 seconds. In order to make our system robust for the average music listener, an ‘‘average sequence’’ is created for each of arousal and valence by computing the average value over all subjects at each timestamp. The average sequences for arousal and valence are then concatenated into an arousal/valence sequence $\mathbf{X}^{(e)}$. Finally, the 30-second segment corresponding to the paired music sample is extracted.

PMemo [22] (and more specifically the updated dataset PMemo2019³) contains 794 music samples which are the chorus parts of high quality popular pop-songs gathered from the Billboard Hot 100, the iTunes Top 100 Songs (USA) and the UK Top 40 Singles Chart. 457 subjects including 366 Chinese university students, 44 Chinese music students and 47 English speaking individuals were recruited for the annotation process. Each music sample is annotated with arousal and valence intensities between 1 (low) and 9 (high) every 0.5 seconds, and then projected into the range [0, 1]. Similarly to DEAM, the first 15 seconds of annotations were discarded by taking into account a large variance in beginnings of music samples. Unlike DEAM, music samples and associated arousal and valence sequences in PMemo have variable lengths ranging from 0.08 to 73.24 seconds. We decided to select music samples with a total length of at least 7.0 seconds to evaluate in total 701 samples. Similarly to DEAM, an arousal/valence sequence $\mathbf{X}^{(e)}$ was associated to each music sample by averaging arousal and valence intensities over all subjects who annotated the sample at each timestamp.

B. EVALUATION METRICS

Each dataset is split into training and test partitions with a proportion of 8 : 2. Specifically, DEAM is split into training and test partitions, respectively containing 1441 and $Q = 361$ pairs of a music sample and an emotion. The training and test partitions of PMemo include 560 and $Q = 141$ pairs, respectively. A model trained on a training partition is evaluated on the corresponding test partition in the framework of M2E and E2M. On both datasets, M2E is run Q times using each of the Q music samples $\mathbf{X}_j^{(m)}$ as a query. For each query, the Q test emotions $\mathbf{X}_j^{(e)}$ are sorted in decreasing order of their similarities to the query $s(\mathbf{X}_q^{(m)}, \mathbf{X}_j^{(e)})$, and the performance is evaluated by checking the rank of the emotion associated with the query music sample. Similarly, E2M is executed Q times by adopting each of the Q emotions $\mathbf{X}_j^{(e)}$ as a query and examining the rank of its associated music sample $\mathbf{X}_j^{(m)}$.

In this framework, only one sample (i.e., emotion or music sample) is associated with a query (i.e., query music sample or query emotion). We use the *Mean Reciprocal Rank* (MRR) as the main evaluation metric as it is commonly employed in retrieval studies [11], [12], [54]. The MRR is calculated based

on r_q that is the rank of the sample associated with a query q in the sorted list of samples ($1 \leq r_q \leq Q$). The MRR is defined as the average of reciprocals of all r_q over Q queries, that is, $\text{MRR} = \frac{1}{|Q|} \sum_{q=1}^{|Q|} \frac{1}{r_q}$. However, the MRR is biased in the sense that it puts much higher priorities on samples ranked at high positions than those at low positions. This can lead to low MRR values even if the rank of the sample associated to the query is low. For example, $r_q = 1$ leads to a reciprocal of 1 while it is close to 0 (0.05) for $r_q = 20$, even though $r_q = 20$ might still be a very good retrieval result. In our case, samples associated with queries are rarely ranked at the very top positions because each query in either the DEAM or PMemo dataset has many ‘close neighbours’ (i.e., samples annotated with similar emotions, or from the same music style) that may easily be ranked above the sample associated with the query. The MRR values might therefore be non-intuitive and not trivially interpretable when checking the performances of our system. Thus, we additionally compute the *Average Rank* (AR) that is the average of all r_q over Q queries, namely $\text{AR} = \frac{1}{|Q|} \sum_{q=1}^{|Q|} r_q$. Using the AR, samples can be equally evaluated regardless of their ranks. Although the median of all r_q is one popular evaluation measure for embedding-based retrieval [53], we use their average to be consistent with the calculation of the MRR. To sum up, our evaluation is based on the MRR and AR that respectively become higher and lower as a better performance is obtained. Finally, all models in each configuration are run 10 times. In each of them, all the parameters in EMER-CL (i.e., parameters of music and emotion encoders and six branches of FC layers in Fig. 2) are randomly initialised, and a dataset is randomly split into training and test partitions with a proportion of 8 : 2. The mean and standard-deviation of MRRs and ARs obtained in 10 runs are reported.

C. IMPLEMENTATION DETAILS

We tested various combinations of an MLP and RNN for music and emotion encoders on DEAM and PMemo. To simplify the selection of encoders, MLPs (or RNNs) with the same architecture were used regardless of encoder types. We found that for the music encoder, an MLP and RNN performed the best on DEAM and PMemo, respectively. For the emotion encoder, RNNs are the best on both datasets.

The numbers of layers and units per layer of the MLP and RNNs were chosen by grid search. The MLP consists of five FC layers, each of which performs a non-linear transformation based on units using softplus $\sigma(x) = \log(1 + e^x)$ as their activation function. The number of units in each layer is 256 for the first layer, 512 for the second and third layers, and 1024 units for the fourth and fifth layers. That is, $D_v^{(m)} = 1024$ when the encoder is an MLP. A dropout layer with a dropout rate of 0.5 is inserted between two consecutive layers.

Each RNN based on bidirectional GRU has a single layer with a 512-dimensional hidden state (i.e., $D_h^{(m)} = D_h^{(e)} = 512$), and finally outputs a 1024-dimensional vector by

³<https://github.com/HuiZhangDB/PMemo>

concatenating the hidden states obtained in the forward and backward directions. This vector is subsequently passed to an MLP consisting of five FC layers where units use softplus as their activation function. The number of units per FC layer was chosen as 512 for the first three ones, and 1024 for the two last ones (i.e., $D_v^{(m)} = D_v^{(e)} = 1024$). A dropout layer with a dropout rate of 0.5 is also added behind all layers except the RNN layer and the output layer.

Regarding the embedding spaces based on the CCA and KL-divergence losses, their dimensionalities are set to $D^{(cca)} = D^{(kl)} = 1024$. Based on our experiments, we recommend to set the threshold for the correlation-based similarity P , the margin in the KL-divergence loss α and the combination weight for the composite loss λ to default values of $(P, \alpha, \lambda) = (0.4, 1.0, 0.5)$ when testing our approach on a new dataset. It is nevertheless possible to optimise these values on a specific dataset by following the optimisation strategy detailed in Appendix A. The results in the next subsections were obtained using the optimal hyper-parameters $(P, \alpha, \lambda) = (0.5, 1.0, 0.6)$ on DEAM and $(P, \alpha, \lambda) = (0.7, 2.0, 0.5)$ on PMemo.

Our EMER-CL model is trained using Adam [55] as the optimiser with an initial learning rate of $1e^{-5}$. The model was trained for 5001 and 10001 epochs on DEAM and PMemo, respectively. We implemented all the codes using TensorFlow library (version 1.15) on a machine equipped with Intel i9-9900K CPU, 64GB RAM, NVIDIA RTX 2080Ti GPU and CUDA version 10.0.

D. EVALUATION OF THE COMPOSITE LOSS

To evaluate the effectiveness of our proposed composite loss (*Composite*), we compare its performance to the ones individually obtained only using the CCA loss (*CCA-Loss*) or KL-divergence loss (*KL-Loss*). In addition, to examine the effectiveness of projecting music samples and emotions as probability distributions, we implement the most popular embedding approach that projects them as points based on their cosine similarities in an embedding space [53]. For this approach, all the configurations of our EMER-CL model are the same except that $v^{(m)}$ and $v^{(e)}$ from the music and emotion encoders are projected into vectors instead of multivariate Gaussian distributions, and $\psi(\cdot, \cdot)$ in Eq. (6) is replaced with the negative of their cosine similarity. We report both the performance only using the loss based on cosine similarities (*Cos-Loss*) and the one obtained by the composite loss combining *CCA-Loss* and *Cos-Loss* (*Composite-C*). Following the optimisation strategy described in Appendix B, the hyper-parameters of *Cos-Loss* and *Composite-C* were also optimised to $(P, \alpha, \lambda) = (0.9, 0.3, 0.7)$ on DEAM and $(P, \alpha, \lambda) = (0.9, 0.1, 0.1)$ on PMemo.

Table 1 shows the results obtained with the five losses described previously. For both DEAM and PMemo, the MRRs and ARs using *KL-Loss* are better than those using *CCA-Loss*. This may be due to the fact that *CCA-Loss* alone only considers the correlation between music samples and emotions, and does not necessarily make sure that

TABLE 1. Comparison of MRRs and ARs using five different losses. The average and standard deviation of each metric over 10 runs are reported on both datasets.

LossType	DEAM			
	M2E		E2M	
	MRR	AR	MRR	AR
<i>CCA-Loss</i>	0.048 ± 0.007	85.4 ± 5.1	0.046 ± 0.006	86.9 ± 4.3
<i>KL-Loss</i>	0.071 ± 0.007	69.9 ± 5.2	0.069 ± 0.006	69.9 ± 5.2
<i>Cos-Loss</i>	0.060 ± 0.006	70.7 ± 4.4	0.069 ± 0.006	70.7 ± 4.0
<i>Composite</i>	0.128 ± 0.011	56.0 ± 1.4	0.114 ± 0.011	55.3 ± 1.8
<i>Composite-C</i>	0.063 ± 0.010	70.5 ± 5.4	0.078 ± 0.010	68.4 ± 2.5
LossType	PMemo			
	M2E		E2M	
	MRR	AR	MRR	AR
<i>CCA-Loss</i>	0.063 ± 0.008	44.3 ± 2.3	0.071 ± 0.011	43.8 ± 2.3
<i>KL-Loss</i>	0.448 ± 0.112	5.0 ± 2.4	0.442 ± 0.113	5.0 ± 2.4
<i>Cos-Loss</i>	0.099 ± 0.013	37.6 ± 2.8	0.093 ± 0.010	37.2 ± 2.6
<i>Composite</i>	0.544 ± 0.056	3.8 ± 0.6	0.542 ± 0.051	3.8 ± 0.5
<i>Composite-C</i>	0.365 ± 0.054	9.5 ± 1.7	0.361 ± 0.056	9.6 ± 1.7

music samples and emotions in positive pairs are placed close to each other in the embedding space. On the other hand, *KL-Loss* assigns music samples and emotions in positive pairs to similar multivariate Gaussian distributions while distinguishing the ones in negative pairs by dissimilar distributions. Furthermore, the performances are significantly improved when using *Composite* compared to using only *KL-Loss* or *CCA-Loss*. This verifies the effectiveness of *Composite* that simultaneously considers *CCA-Loss* and *KL-Loss*. Also, the fact that *KL-Loss* outperforms *Cos-Loss* and *Composite* outperforms *Composite-C* on both DEAM and PMemo datasets indicates the effectiveness of distribution-based embeddings compared to point-based ones.

Finally, it can be noted that the performances on PMemo are significantly better than those on DEAM. This could be attributed to the fact that music samples in PMemo are more standardised, for instance by including only chorus parts of pop songs, which leads the music encoder to find more specialised feature. In other words, the higher diversity in music samples on DEAM makes M2E and E2M on this dataset likely to be more difficult.

E. COMPARISON WITH THE BASELINE MODELS

To the best of our knowledge, no existing EMER method that can be directly compared to EMER-CL has been proposed yet. In addition, all the existing approaches using continuous emotion modelling only perform M2E based on a regression approach to predict real-valued characteristics of the arousal/valence sequence (e.g., the average arousal or valence) for a given music sample [28], [29], [36], [37], [38], [40], [41], [42], [43], [56], [57], [58]. Moreover, no existing method can handle E2M to predict acoustic features of the music sample for a given arousal/valence sequence. Considering the aforementioned state of the current MER research,

we define the following regression-based baselines to show the effectiveness of EMER-CL.

1) M2E BASELINES

Two M2E baselines *RegBiGRU-M2E* and *RegMLP-M2E* train a regression model that analyses a query music sample and outputs a two-dimensional emotion vector $\bar{\mathbf{x}}^{(e)}$ representing the arousal and valence averaged over time for this query music sample. In particular, *RegBiGRU-M2E* predicts $\bar{\mathbf{x}}^{(e)}$ by applying an RNN based on bidirectional GRU to a sequence of acoustic features $\mathbf{X}^{(m)}$, while *RegMLP-M2E* employs an MLP that uses the mean $\bar{\mathbf{x}}^{(m)}$ of features in $\mathbf{X}^{(m)}$ over time to compute $\bar{\mathbf{x}}^{(e)}$. Both baselines are trained to minimise the Mean Absolute Error (MAE) between $\bar{\mathbf{x}}^{(e)}$ and the ground-truth mean emotion $\bar{\mathbf{x}}^{(e)}$ computed from the actual arousal/valence sequence $\mathbf{X}^{(e)}$.

In the evaluation, given the q th test music sample as a query, we evaluate the trained model by predicting its mean emotion $\bar{\mathbf{x}}_q^{(e)}$ and checking whether $\bar{\mathbf{x}}_q^{(e)}$ is similar to the ground-truth mean emotion $\bar{\mathbf{x}}_q^{(e)}$. To this end, we compute the similarities of $\bar{\mathbf{x}}_q^{(e)}$ to the ground-truth mean emotions $\{\bar{\mathbf{x}}_t^{(e)}\}_{t=1}^Q$ for all the Q test music samples. The Absolute Error (AE) between $\bar{\mathbf{x}}_q^{(e)}$ and $\bar{\mathbf{x}}_t^{(e)}$ is used as their dissimilarity. Then, the ground-truth mean emotions are sorted in ascending order of their AEs to get the rank r_q of $\bar{\mathbf{x}}_q^{(e)}$. Finally, r_q is used to calculate an MRR and AR.

2) E2M BASELINES

Similarly to the M2E baselines, an RNN based on bidirectional GRU (*RegBiGRU-E2M*) and an MLP model (*RegMLP-E2M*) are used as E2M baselines to predict a mean acoustic feature $\bar{\mathbf{x}}^{(m)}$. *RegBiGRU-E2M* and *RegMLP-E2M* take as input an arousal/valence sequence $\mathbf{X}^{(e)}$ and the mean vector $\bar{\mathbf{x}}^{(e)}$ of $\mathbf{X}^{(e)}$, respectively. *RegBiGRU-E2M* and *RegMLP-E2M* are trained to minimise the MAE between $\bar{\mathbf{x}}^{(m)}$ and the ground-truth mean acoustic feature $\bar{\mathbf{x}}^{(m)}$ computed from a sequence of acoustic features $\mathbf{X}^{(m)}$. Like M2E, an MRR and AR is computed by predicting $\bar{\mathbf{x}}_q^{(m)}$ for the q th test emotion, measuring the AEs of $\bar{\mathbf{x}}_q^{(m)}$ to the ground-truth mean acoustic features $\{\bar{\mathbf{x}}_t^{(m)}\}_{t=1}^Q$ for all the Q music samples, and check the rank r_q of $\bar{\mathbf{x}}_q^{(m)}$.

The baseline models were tuned by grid search to find the hyper-parameters leading to the best performances on each dataset. Details about this hyper-parameter tuning can be found in Appendix C. Table 2 shows the comparison between the above-mentioned baselines and EMER-CL (referred to as *Composite* in Table 1). EMER-CL significantly outperforms the baselines based on one-way regression of emotion or acoustic features. This highlights the superiority of our embedding-based approach over traditional regression methods not using embeddings.

F. COMPARISON TO THE MER STATE-OF-THE-ART

To evaluate the features learnt by EMER-CL, we performed a comparison of the latter against the state-of-the-art on DEAM

TABLE 2. Comparison between the baselines and EMER-CL (Composite). The average and standard deviation of each metric over 10 runs are reported on both datasets.

ModelType	DEAM			
	M2E		E2M	
	MRR	AR	MRR	AR
<i>RegMLP</i>	0.066 ± 0.002	82.5 ± 0.3	0.039 ± 0.0003	136.5 ± 0.1
<i>RegBiGRU</i>	0.057 ± 0.005	83.2 ± 2.7	0.046 ± 0.001	131.2 ± 0.3
<i>Composite</i>	0.128 ± 0.011	56.0 ± 1.4	0.114 ± 0.011	55.3 ± 1.8
ModelType	PMEmo			
	M2E		E2M	
	MRR	AR	MRR	AR
<i>RegMLP</i>	0.082 ± 0.010	38.4 ± 1.5	0.050 ± 0.002	67.3 ± 0.5
<i>RegBiGRU</i>	0.095 ± 0.011	38.2 ± 0.8	0.067 ± 0.003	57.3 ± 0.2
<i>Composite</i>	0.544 ± 0.056	3.8 ± 0.6	0.542 ± 0.051	3.8 ± 0.5

and PMEmo. To the best of our knowledge, the EMER problem remains still unexplored in the literature, which makes it difficult to find a past study with which our results can be directly compared. Therefore, we carry out the evaluation on the significantly more popular task of MER.

The MER literature is fairly scattered, with each study carrying out experiments on different datasets, choosing different evaluation metrics and strategies. Our experiments on both DEAM and PMEmo were carried out based on the most commonly used setting, that is, a K-fold cross validation evaluated either using the Root Mean Squared Error (RMSE), the Pearson correlation coefficient R or the coefficient of determination R^2 .

Our EMER-CL model is trained for $K = 10$ folds using the default values of $\alpha = 1$ and $\lambda = 0.5$ (P is only required for a retrieval problem, not a regression one). After training the whole of our EMER-CL model, three vectors obtained from the music model (i.e., $\phi^{(m)}$, $\mu^{(m)}$ and $\sigma^{(m)}$ in Fig. 2) are used to train a soft-margin Support Vector Regression model (C-SVR) with Radial Basis Function (RBF) kernel that attempts to predict the arousal and valence intensities associated with a music sample that is input to the music model. Three different configurations for the input to C-SVR are tested: $\phi^{(m)}$ used alone (also referred to as *EMER-CL (cca)*), $\mu^{(m)}$ and $\sigma^{(m)}$ concatenated together (*EMER-CL (kl)*), and all the three vectors concatenated together (*EMER-CL (all)*). On DEAM, the target to predict was chosen as $\bar{\mathbf{x}}^{(e)}$, the average of an arousal/valence sequence $\mathbf{X}^{(e)}$ over time. On PMEmo, the last arousal and valence values of the emotion sequence $\mathbf{X}^{(e)}$ were predicted instead. The hyper-parameters of the C-SVR (soft-margin and kernel parameters) were optimised by maximising the average R^2 after grid search.

Table 3 shows the results obtained for arousal and valence prediction on DEAM and PMEmo. The learnt features $\phi^{(m)}$, $\mu^{(m)}$ and $\sigma^{(m)}$ can compete with the state-of-the-art for MER. This indicates that our EMER-CL model can still yield proper

TABLE 3. Comparison between EMER-CL and the state-of-the-art for M2E. The metrics for EMER-CL are provided as the average obtained on 10 folds.

DEAM						
Approach	R^2	Arousal	R^2	Valence		
Hult et al. [59]	0.35			0.34		
Cheuk et al. [60]	0.672			0.367		
<hr/>						
<i>EMER-CL (cca)</i>	0.484			0.661		
<i>EMER-CL (kl)</i>	0.491			0.667		
<i>EMER-CL (all)</i>	0.489			0.664		
<hr/>						
PMEmo						
Approach	Arousal			Valence		
	RMSE	R	R^2	RMSE	R	R^2
Zhang et al. [22]	0.102	0.764	-	0.124	0.638	-
Hult et al. [59]	-	-	0.64	-	-	0.42
Chapaneri et al. [61]	0.064	-	0.618	0.093	-	0.376
de Berardinis et al. [62]	0.232	-	0.600	0.232	-	0.481
<hr/>						
<i>EMER-CL (cca)</i>	0.115	0.701	0.493	0.111	0.792	0.628
<i>EMER-CL (kl)</i>	0.115	0.702	0.494	0.111	0.793	0.629
<i>EMER-CL (all)</i>	0.115	0.702	0.495	0.111	0.793	0.630

MER performances. Curiously, $\phi^{(m)}$, $\mu^{(m)}$ and $\sigma^{(m)}$ yield average results for arousal prediction, and notably good ones for valence prediction which is commonly considered as the most difficult of the two problems. All the three tested combinations of $\phi^{(m)}$, $\mu^{(m)}$ and $\sigma^{(m)}$ return fairly similar performances.

V. DETAILED ANALYSIS

MRRs and ARs are global metrics that only depend on the rank r_q of the music sample or emotion associated with a query. It is also desirable to check the relevance of the top-ranked music samples (or emotions) to the query. For this, we compute for M2E an *average cosine similarity* that averages the cosine similarities between the mean acoustic feature of the query music sample $\bar{x}_q^{(m)}$ and the ones associated with the top 5% emotions retrieved by EMER-CL, i.e., $\{\bar{x}_j^{(m)} | 1 \leq j \leq Q, r_j \leq 0.05 \times Q\}$. For E2M, this average cosine similarity is computed in a likewise way by taking the mean of the cosine similarities computed between the query emotion averaged over time $\bar{x}_q^{(e)}$ and the ones associated with the top 5% retrieved music samples in E2M, i.e., $\{\bar{x}_j^{(e)} | 1 \leq j \leq Q, r_j \leq 0.05 \times Q\}$. A high average cosine similarity means that EMER-CL can recognise music samples that express emotions similar to a query emotion, or emotions expressed in music samples which are acoustically similar to a query music sample.

In what follows, we present the analysis for E2M since the mean emotion for each music sample is two-dimensional and can be interpreted easily. It should be noted that arousal/valence intensities are in $[-1, 1]$ and $[0, 1]$ for DEAM and PMEmo respectively, meaning that average cosine similarities range between -1 and 1 on DEAM, and 0 and 1 on PMEmo. Fig. 3 shows the average cosine similarities for DEAM and PMEmo. In the bar graphs in the left side

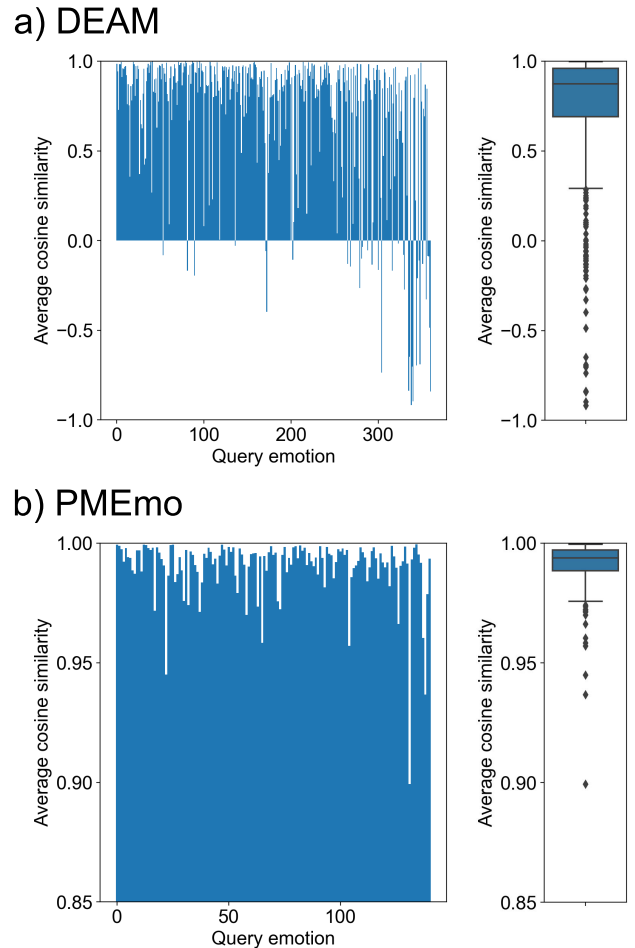


FIGURE 3. Bar graphs showing average cosine similarities and their box plots on DEAM and PMEmo.

of Fig. 3 (a) and (b), each query emotion on the horizontal axis is sorted in increasing order of the rank r_q of its associated music sample. That is, the more to the left a query emotion is, the higher its ground truth music sample was ranked for E2M, meaning that the music sample associated with the query emotion was well recognised.

As shown in the bar graphs of Fig. 3 (a) and (b), the average cosine similarities obtained on both DEAM and PMEmo are fairly high (close to 1), showing that the top 5% recognised music samples are relevant to most query emotions regardless of the ranks of their ground-truth music samples. More specifically, the median of these average cosine similarities is 0.873 for DEAM and 0.994 for PMEmo (the reason for the very high average cosine similarities on PMEmo is provided in Appendix E). In addition, the box plots in the right side of Fig. 3 (a) and (b) show the variations in the average cosine similarities. Here, at least 75 percent of all the average cosine similarities are higher than the 25th percentile (first quartile). The fact that the 25th percentile for DEAM and PMEmo are 0.692 and 0.989 respectively, indicates that EMER-CL is able to robustly recognise music samples associated to highly similar emotions to a query emotion. In other words, even if

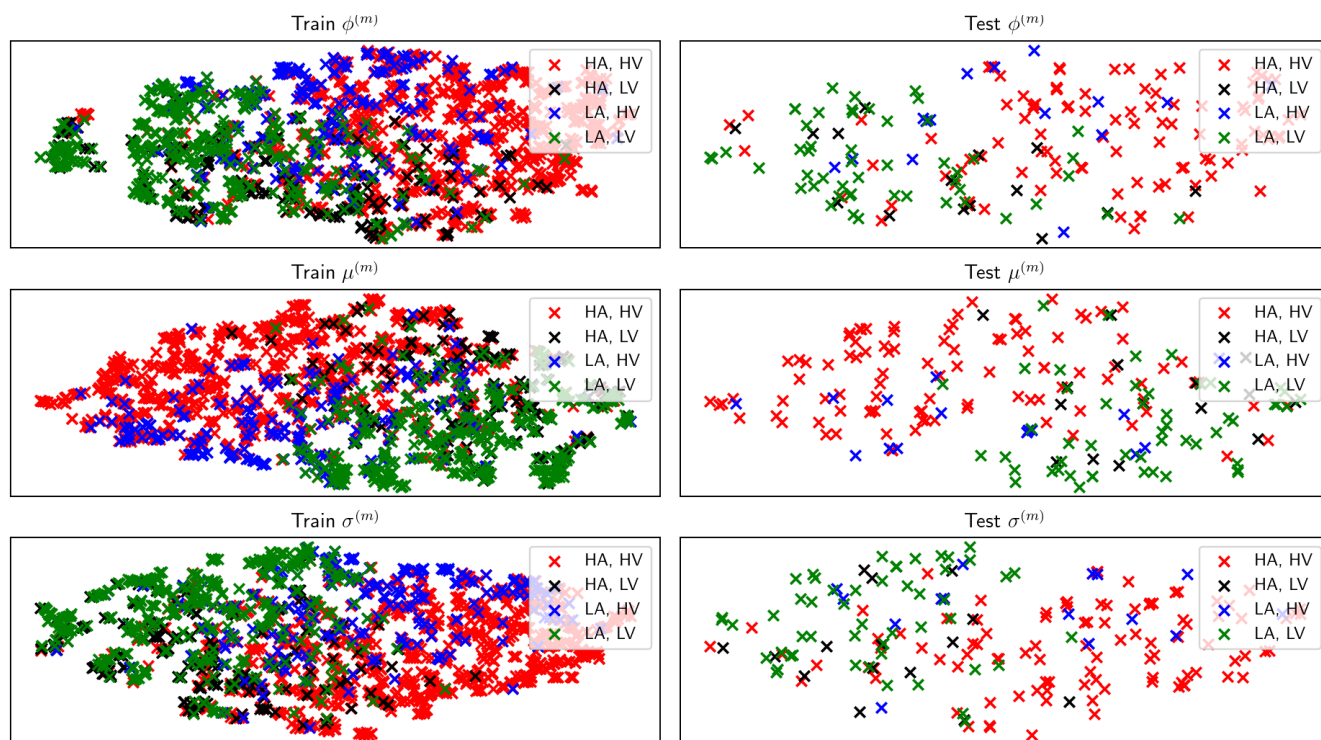


FIGURE 4. t-SNE projections of the music embeddings $\phi^{(m)}$, $\mu^{(m)}$ and $\sigma^{(m)}$ (respectively first, second and third rows) on the DEAM training (left column) and testing (right column) sets. The t-SNE projections are labelled with their associated emotional quadrants, i.e., high arousal/high valence (HA, HV), high arousal/low valence (HA, LV), low arousal/low valence (LA, LV) and low arousal/high valence (LA, HV).

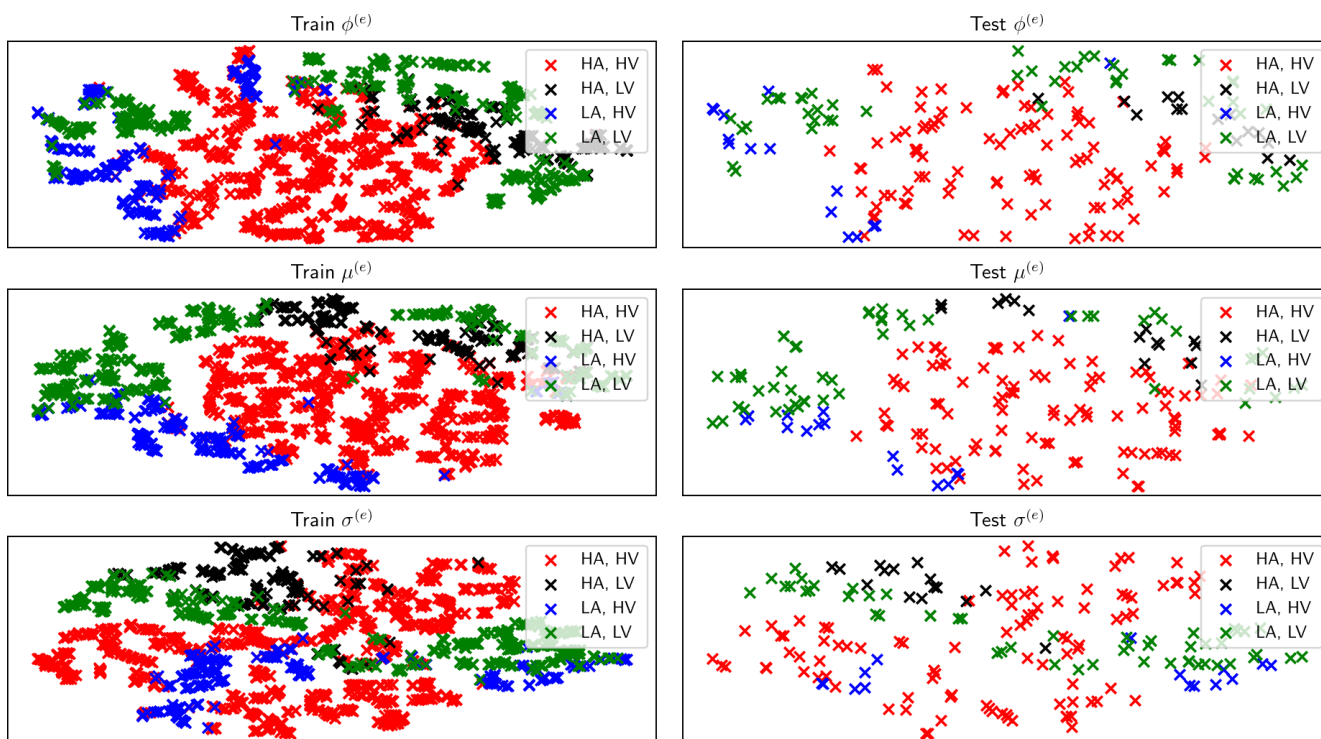


FIGURE 5. t-SNE projections of the emotion embeddings $\phi^{(e)}$, $\mu^{(e)}$ and $\sigma^{(e)}$ (respectively first, second and third rows) on the DEAM training (left column) and testing (right column) sets. The t-SNE projections are labelled with their associated emotional quadrants, i.e., high arousal/high valence (HA, HV), high arousal/low valence (HA, LV), low arousal/low valence (LA, LV) and low arousal/high valence (LA, HV).

the music sample associated to a query emotion was ranked at a low position, the top 5% music samples recognised by

EMER-CL still exhibit emotions close to the query emotion. Nevertheless, average cosine similarities for some query

emotions in DEAM are low, indicating room for improvement in the future.

The same experiment for M2E that computes the average cosine similarity between the acoustic feature of a query music sample and those of music samples associated with the top 5% emotions in M2E showed similarly good performances. Figures showing such average cosine similarities can be found in Appendix D, and the medians of average cosine similarities on DEAM and PMemo are 0.753 and 0.837, respectively.

Finally as a last check of the validity of our approach, we also visualise the embeddings learnt by our music and emotion models on the training and testing sets of both DEAM and PMemo using t-SNE [63]. We plotted the correlation-based and probabilistic-based embeddings produced by the music model ($\phi^{(m)}$, $\mu^{(m)}$ and $\sigma^{(m)}$) and the emotion model ($\phi^{(e)}$, $\mu^{(e)}$ and $\sigma^{(e)}$). The t-SNE projections were labelled with the emotion-related information available on both DEAM and PMemo datasets as follows: each emotion sequence $X_j^{(e)}$ in either the training or testing set was first averaged over time to obtain the two-dimensional emotion vector $\bar{x}_j^{(e)}$, and then associated to one of the four quadrants of the arousal/valence space: high arousal/high valence (HA/HV), high arousal/low valence (HA/LV), low arousal/low valence (LA/LV) and low arousal/high valence (LA/HV). The t-SNE projections of the embeddings obtained either from $X_j^{(e)}$ or its associated music sample $X_j^{(m)}$ were then annotated with one of these four labels. The t-SNE plots for the DEAM dataset are provided in Figs. 4 and 5 for the music and emotion embeddings, respectively. Since DEAM emotion annotations lie in the range $[-1, 1]$, the cut-off value for the definition of the quadrants was set to 0 for both arousal and valence.

It can be seen from both Figs. 4 and 5 that the DEAM music and emotion embeddings associated with opposite emotional quadrants (e.g. HA/HV and LA/LV, or HA/LV and LA/HV) are very well separated in their respective embedding spaces for both the training and testing examples. This indicates that our EMER-CL approach was successful in learning to project music samples associated to similar emotions close to each other in the embedding space, while simultaneously maximising the distance between embeddings of dissimilar emotions. Similar plots can be obtained for the PMemo dataset and are provided in Appendix G.

VI. CONCLUSION AND FUTURE WORK

In this paper, we introduced an Embedding-based Music Emotion Recognition using Composite Loss (EMER-CL) approach that projects music samples and emotions into embedding spaces, in order to consider both general emotional categories and fine-grained discrimination within each category. In particular, to deal with the emotional uncertainty, EMER-CL uses the composite loss consisting of the CCA loss to maximise the correlation between music samples and their associated emotions in an embedding space, and the KL-divergence loss to project them as similar

multivariate Gaussian distributions in another embedding space. The experimental results on DEAM and PMemo validate the effectiveness of the composite loss, embedding-based approach and features learned by EMER-CL. In addition, a detailed analysis of EMER-CL's results demonstrates that it can robustly recognise reasonable music samples (or emotions) even when failing to identify the ground-truth ones.

To further improve the performance of EMER-CL, we aim to extend the music and emotion encoders by pre-training them with self-supervised learning [64] which can learn underlying feature representations using unlabelled data. We also plan to adopt a self-attention layer [65] which can capture long-term dependencies of features, and have led to promising performances when jointly used with bidirectional LSTM and GRU, in particular for sentiment analysis [66] or sensor-based emotion recognition [67]. Because emotions are also strongly dependent on cultural background, we also plan to use MER datasets that provide detailed background information about their raters in future work. The generality of our approach across cultures could be then be demonstrated.

Finally, the codes (and the instruction of data usage) used in this paper are available on our GitLab repository⁴, in order for other researchers to reproduce the results and extend the current EMER-CL more easily.

APPENDIX A HYPER-PARAMETER TUNING FOR EMER-CL

This section presents how to tune EMER-CL's hyper-parameters, especially $P \in [0, 1]$ used in the correlation-based similarity to filter out useless dimensions characterised by weak correlations between music samples and their associated emotions, $\alpha \in \mathbb{R}^+$ used in the KL-divergence loss to handle the margin between associated music-emotion pairs and non-associated ones, and $\lambda \in [0, 1]$ to control the combination weights of the CCA and KL-divergence losses. Only using either of these losses, an embedding space can be constructed to perform M2E and E2M. Thus, P is firstly tuned based on the performances of M2E and E2M only using the CCA loss. Similarly, α is tuned by carrying out M2E and E2M only with the KL-divergence loss. Finally, λ is tuned based on M2E and E2M by combining the CCA and KL-divergence losses that are configured by the separately optimised P and α , respectively.

A. TUNING P

Fig. 6 shows the various performances obtained only using the CCA loss configured by different values of P . As previously described in Section IV, a performance is measured by an MRR and AR. A good performance is indicated by a high MRR and a low AR. Fig. 6 shows the box plots of the MRR and AR performances obtained for each value of P over 10 runs. They were computed as follows: Since building an embedding space based on the CCA loss is independent of the choice of P , 10 spaces are firstly constructed by randomly

⁴https://mu-lab.info/naoki_takashima/emer-cl

initialising all parameters in EMER-CL (i.e., the parameters of both music and emotion encoders and the ones of six branches of FC layers in Fig. 2⁵), and randomly splitting a dataset into training and test partitions with a proportion fixed to 8 : 2. Then, every value of P is used to filter out useless dimensions in each of these 10 embedding spaces to get 10 performances.

To determine a range of values to be tested for P , we check the maximum correlations among the 1024 dimensions of each embedding space. In particular, among the 10 embedding spaces constructed for each dataset, the maximum correlation is 0.783-0.833 and 0.719-0.774 for DEAM and PMemo, respectively.⁶ We test values between 0 and the maximum correlation with increments of 0.1 for P .

In each graph of Fig. 6, the larger P is, the smaller the number of dimensions used in the correlation-based similarity is. In other words, if P is large, only a small number of dimensions characterised by correlations higher than it are used to compute correlation-based similarities. For each dataset, the optimal value of P is selected as the one that yields the best ‘overall’ performance by considering both M2E and E2M performances. We provide an example of how to select the optimal P value on DEAM by referring to Fig. 6. As seen from this figure, although $P = 0.2$ yields the highest median of MRRs in M2E, $P = 0.5$ and 0.6 both lead to the highest median of MRRs in E2M and the lowest medians of ARs in both M2E and E2M. Thus, $P = 0.5$ or 0.6 can be considered as optimal on DEAM. In this case in particular, $P = 0.5$ is selected after observing that its neighbouring value $P = 0.4$ yields higher performances than $P = 0.7$ neighbouring $P = 0.6$. For PMemo, $P = 0.7$ is chosen because of its significantly higher performances on E2M compared to the other P values. It can be noted that the performances of CCA-Loss in the comparative study in Section IV-D are nothing but the ones that are obtained only using the CCA loss based on the aforementioned optimal P values.

B. TUNING α

Unlike P that is bounded, the margin α can theoretically take any positive value. We decided to test values between 0 and 1.5 with increments of 0.1, and powers of 2 between 2 and 128. Using the same box plot format as Fig. 6, Fig. 7 illustrates the performances obtained using only the KL-divergence loss, which is configured by different values of α . Using a similar strategy as for P , we select the optimal α value in Fig. 7 as the one that leads to the best overall

⁵Since the KL-divergence loss is always zero in this setting, the four branches of FC layers to produce mean and variance vectors of multivariate Gaussian distributions are not trained.

⁶When the composite loss is used, the maximum correlations for 10 CCA-based embedding spaces increase to 0.856-0.857 and 0.9995-0.9997 for DEAM and PMemo, respectively. This is possibly due to more generalised music and emotion encoders being trained by exploiting both the CCA and KL-divergence losses, which leads to higher-quality embedding spaces. Thus, better EMER-CL’s performances than those reported in this paper might be obtained by carrying out grid search on P , α and λ . But, due to its expensive computational cost, we opt to separately tune these hyper-parameters in this paper.

performance. As it can be seen from Fig. 7, the performances are similar for all tested values, with the exception of small values 0.0, 0.1 and 0.2 for which the performances are significantly worse. We decide to select $\alpha = 1.0$ and 2.0 as the optimal values for DEAM and PMemo respectively, on the basis that the overall performances with these P values are slightly higher than those with the others. It can be noted that the performances acquired by the aforementioned optimal α values are reported as the ones of KL-Loss in the comparative study in Section IV-D.

C. TUNING λ

We tested values of λ between 0 and 1 with increments of 0.1. In the same manner as Figs. 6 and 7, Fig. 8 illustrates EMER-CL’s performances obtained for different values of λ . For each of them, the CCA and KL-divergence losses that are configured by the optimal P and α values (found from Figs. 6 and 7 respectively) are used to compute the composite loss. The higher λ is, the higher the weight of the CCA loss is. In particular, $\lambda = 0$ means only using the KL-divergence loss while only the CCA loss is used with $\lambda = 1$. Following the same strategy as the ones employed for choosing the optimal P and λ values, $\lambda = 0.6$ is chosen as the optimal value yielding the best overall performance on DEAM, and similarly $\lambda = 0.5$ is regarded as optimal for PMemo.

To summarise the whole EMER-CL hyper-parameter selection process, the optimal values found are $P = 0.5$, $\alpha = 1.0$ and $\lambda = 0.6$ on DEAM, and $P = 0.7$, $\alpha = 2.0$ and $\lambda = 0.5$ on PMemo. The performances of EMER-CL using these optimal values are reported as the ones of *Composite* in Section IV-D. In addition, it should be noted that these optimal hyper-parameter values only marginally differ from the default ones (i.e., $P = 0.4$, $\alpha = 1.0$ and $\lambda = 0.5$). Table 4 shows the performances of EMER-CL with optimised parameters - referred to as *Composite* - and the ones obtained with the default parameters - referred to as *Composite-d*. As it can be seen from this table, the performances of the latter are relatively similar to the ones of the former. The marginal difference in hyper-parameter values between *Composite* and *Composite-d* and their similar performances validate the relevance of the default values.

APPENDIX B HYPER-PARAMETER TUNING FOR COMPOSITE-C

In this section, we tune hyper-parameters of the two methods *Cos-loss* and *Composite-C* used in the comparative study in Section IV-D. *Cos-Loss* produces point-based embeddings using the loss based on cosine similarities between music samples and emotions [53], and is used to examine the effectiveness of our proposed KL-Loss implementing distribution-based embeddings. Both of *Cos-Loss* and *KL-Loss* construct an embedding space in the same ranking loss framework that involves α to control the margin between associated music-emotion pairs and non-associated ones. Thus, α for *Cos-Loss* is tuned in the same manner as α for *KL-Loss* in Section A-B.

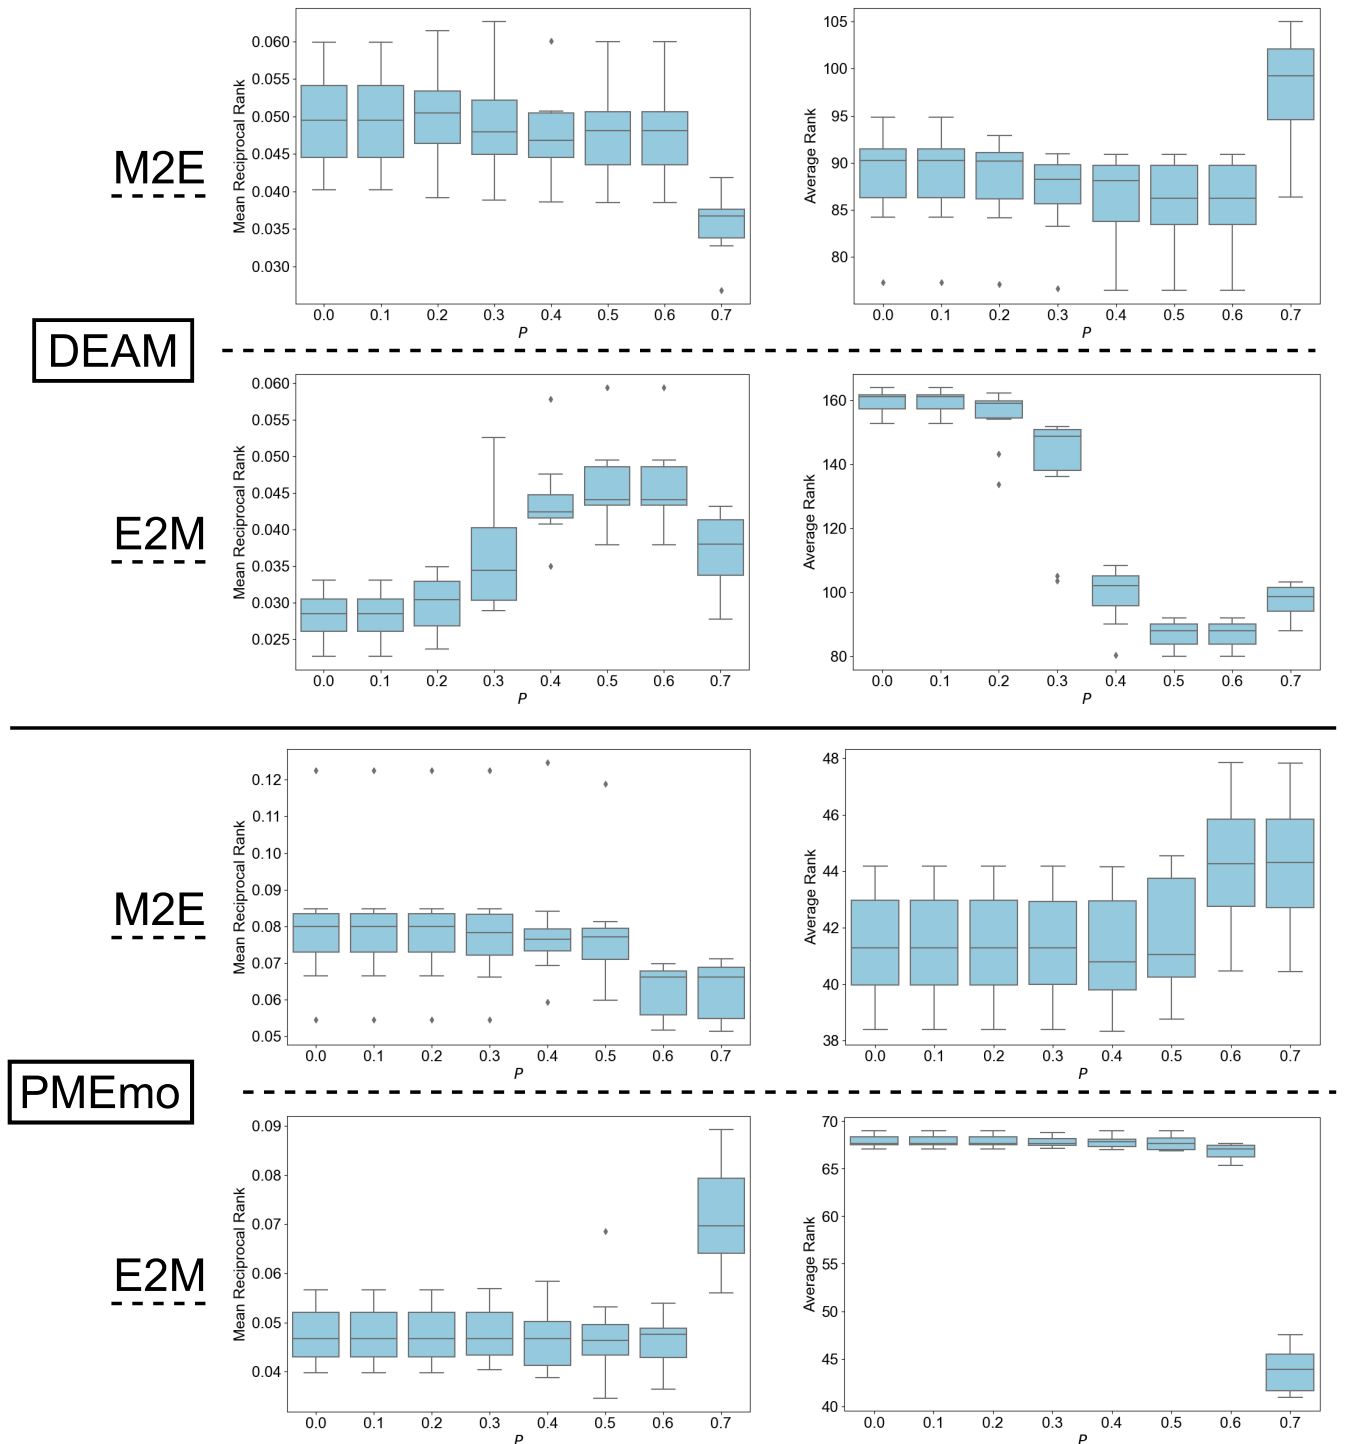


FIGURE 6. Transitions of EMER-CL performances (MRRs and ARs) obtained only using the CCA loss that is configured by different values of P . The optimal P value is selected as 0.5 and 0.7 for DEAM and PMEmo, respectively.

Similarly to *Composite*, *Composite-C* is characterised by λ that handles the combination weights of *CCA-Loss* and *Cos-Loss*. However, while *Cos-Loss* is based on normalised cosine similarities ranging from -1 to 1 , the KL-divergences used in *KL-Loss* are contained in a much wider range, like on average about 5.1 for DEAM and 423.7 for PMEmo. It should be noted that P impacts the range of correlation-based

similarities because it determines the number of dimension-wise similarities counted to compute an overall similarity, as seen from Eq. (9). Thus, both P and λ can be tuned to balance the combination of *CCA-Loss* and *Cos-Loss*, and the P values found in Fig. 6 for *Composite* may not be suitable for *Composite-C*. Consequently, after selecting the optimal α for *Cos-Loss*, grid search on P and λ is carried out to avoid

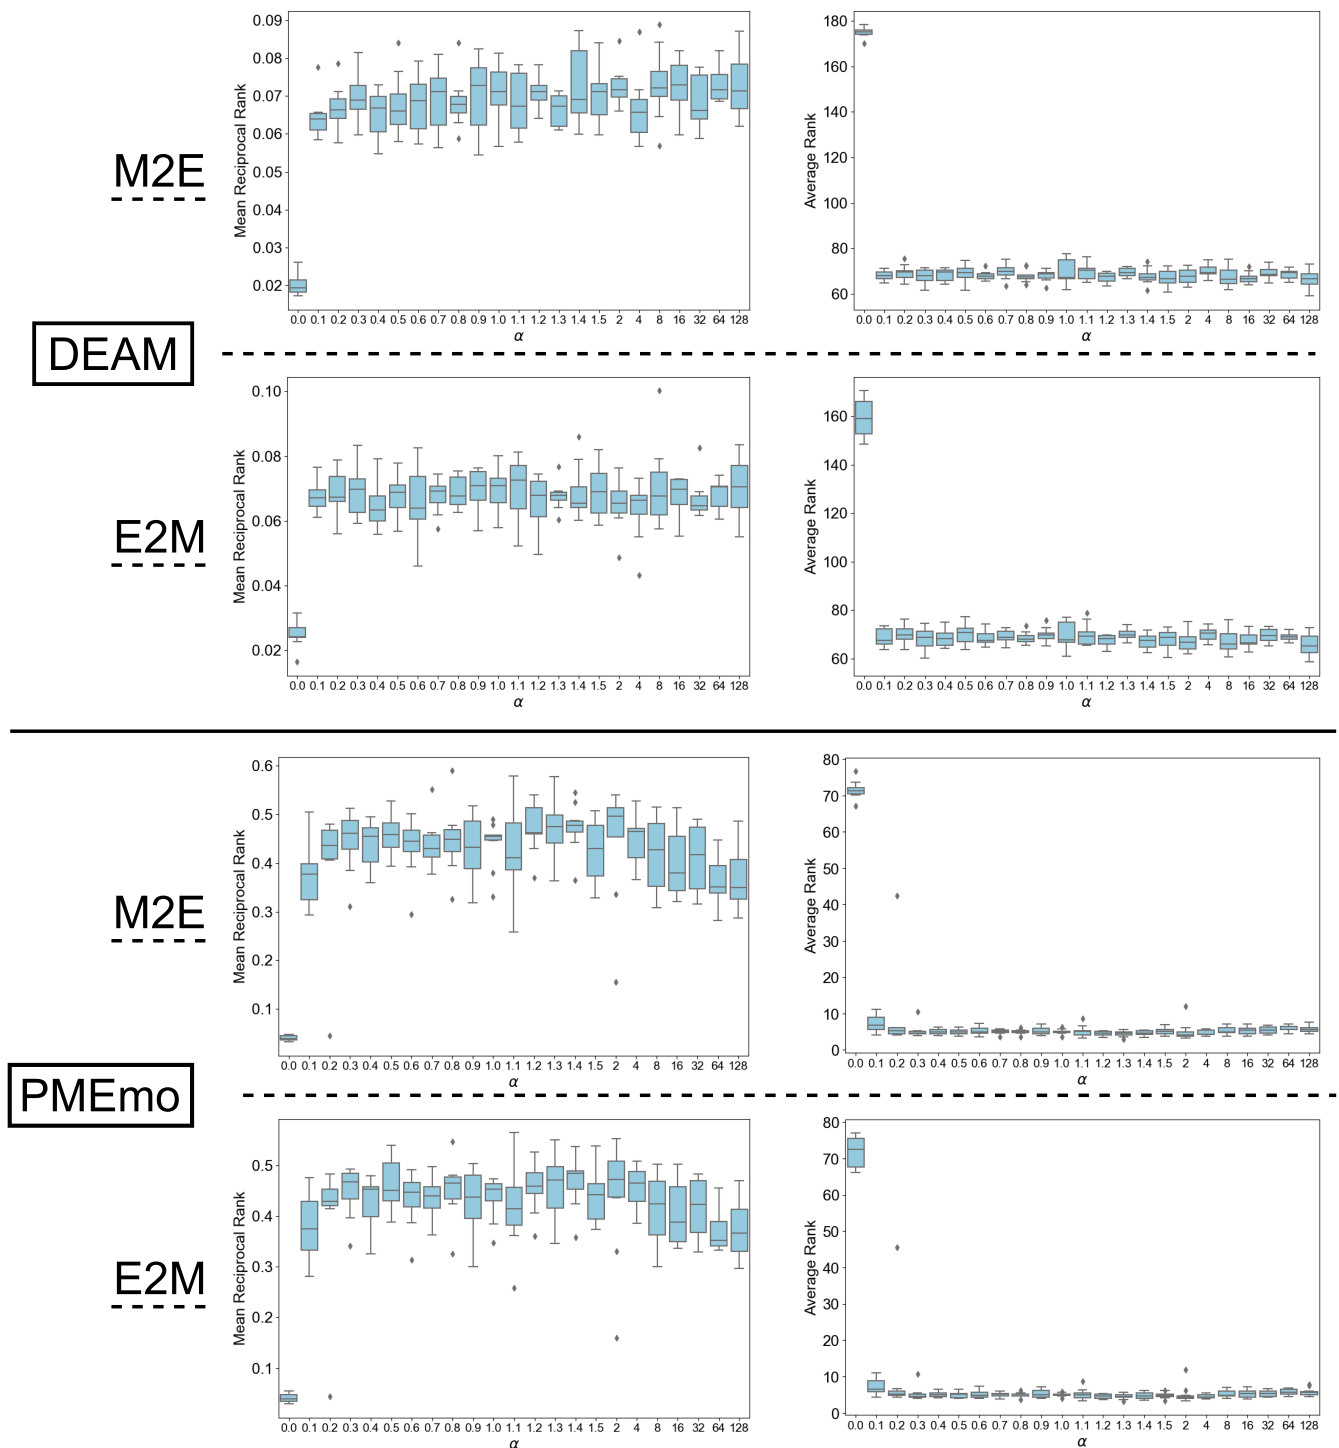


FIGURE 7. Transitions of EMER-CL performances (MRRs and ARs) obtained only using the KL-divergence loss that is configured by different values of α . The optimal α value is selected as 1.0 and 2.0 for DEAM and PMEmo, respectively.

missing an effective combination of *CCA-Loss* and *Cos-Loss* for *Composite-C*. This grid search for *Composite-C* is more exhaustive and favourable than the separate optimisation of P and λ employed for *Composite*.

A. TUNING α OF COS-LOSS

The same values as in Section A-B were tested to tune α for *Cos-Loss*, i.e., values between 0 and 1.5 with increments

of 0.1, and powers of 2 between 2 and 128. Using the same box plot format as Figs. 6, 7 and 8, Fig. 9 displays various performances on DEAM and PMEmo only using *Cos-Loss* configured by different values of α . By following the aforementioned criteria addressing the best overall performance on M2E and E2M, $\alpha = 0.3$ for DEAM and $\alpha = 0.1$ for PMEmo are chosen as the optimal values. The performances of *Cos-Loss* configured

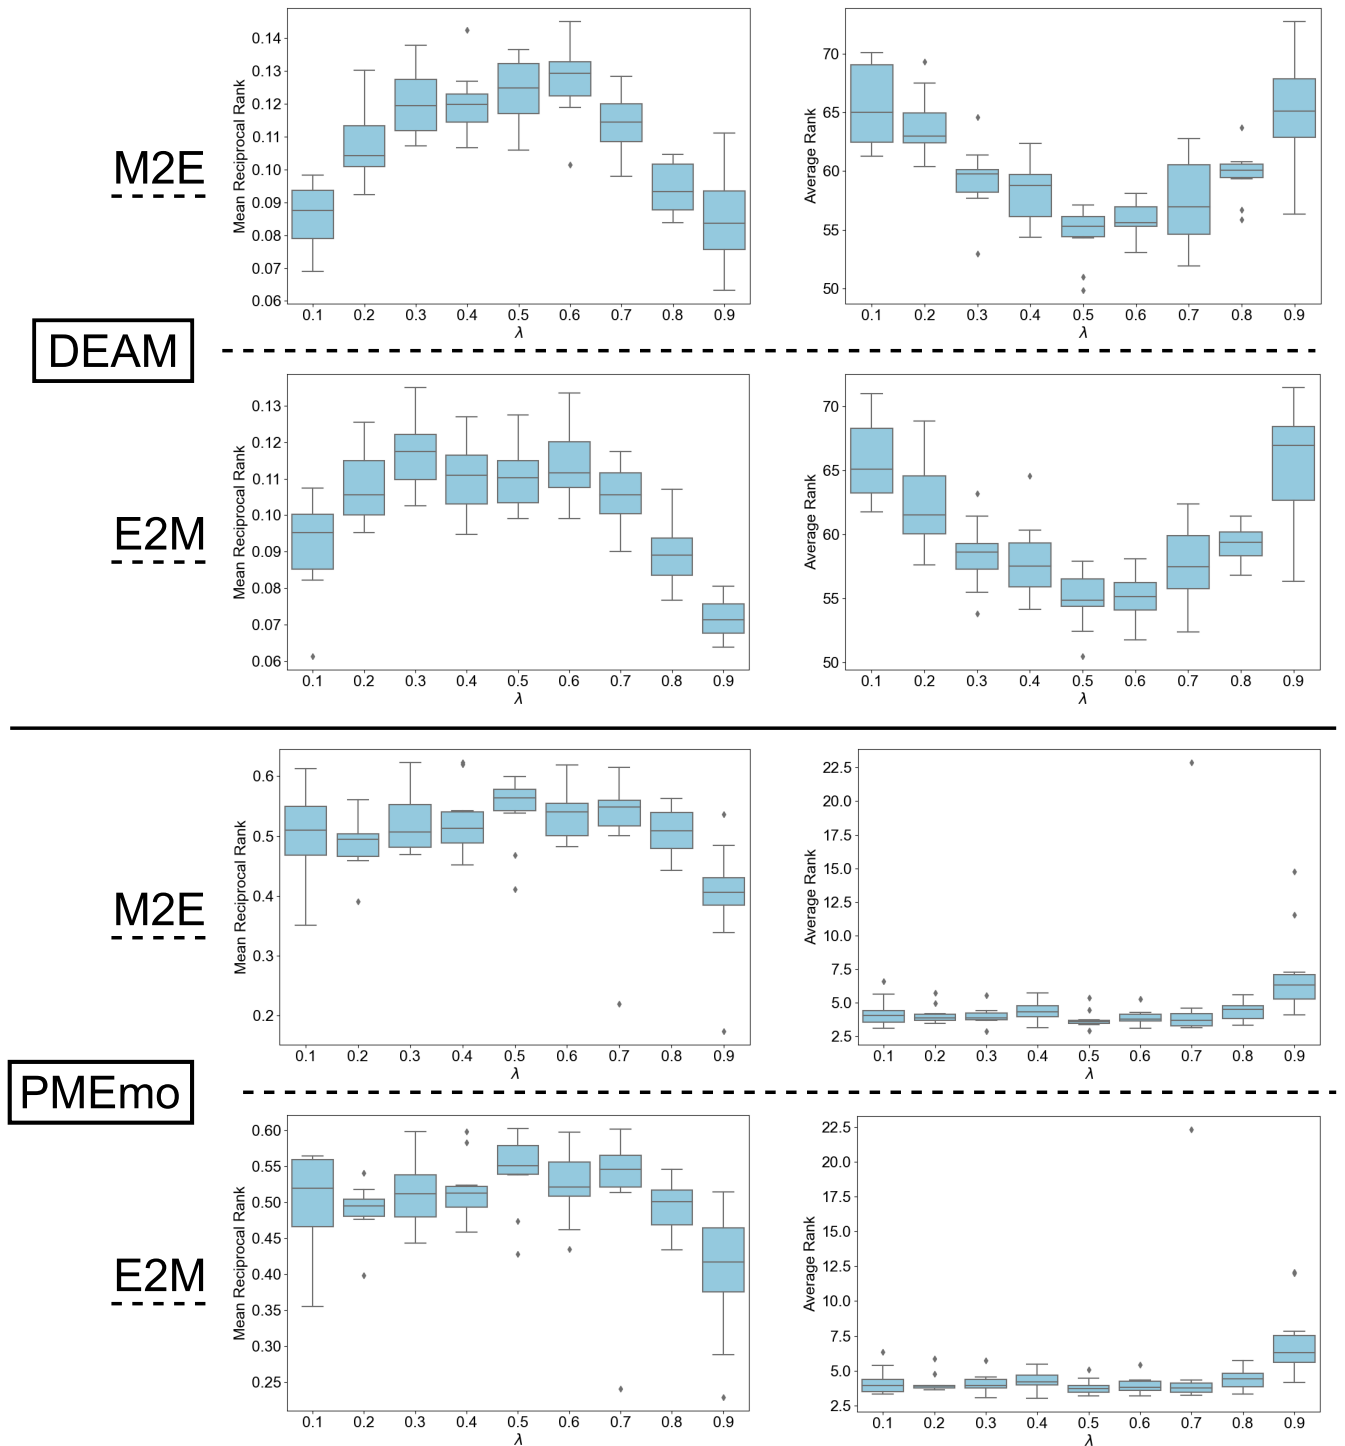


FIGURE 8. Transitions of EMER-CL performances (MRRs and ARs) obtained by different λ values, each of which is used to combine the CCA and KL-divergence losses configured by the optimal P and λ values found from Figs. 6 and 7, respectively. The optimal λ value is selected as 0.6 and 0.5 for DEAM and PMEmo, respectively.

with these α values are used in the comparative study in Section IV-D.

B. TUNING P AND λ OF COMPOSITE-C BY GRID SEARCH

Fig. 10 presents grid search results for different pairs of P and λ values, where P ranges between 0 and the maximum

correlation on the considered dataset with an increment of 0.1 and $\lambda \in \{0.1, \dots, 0.9\}$. Each bar in a three-dimensional bar graph indicates the mean of 10 performances (i.e., MRRs or ARs) obtained using a pair of P and λ values. Like for the hyper-parameter tuning procedure previously described, these 10 performances are acquired by randomly

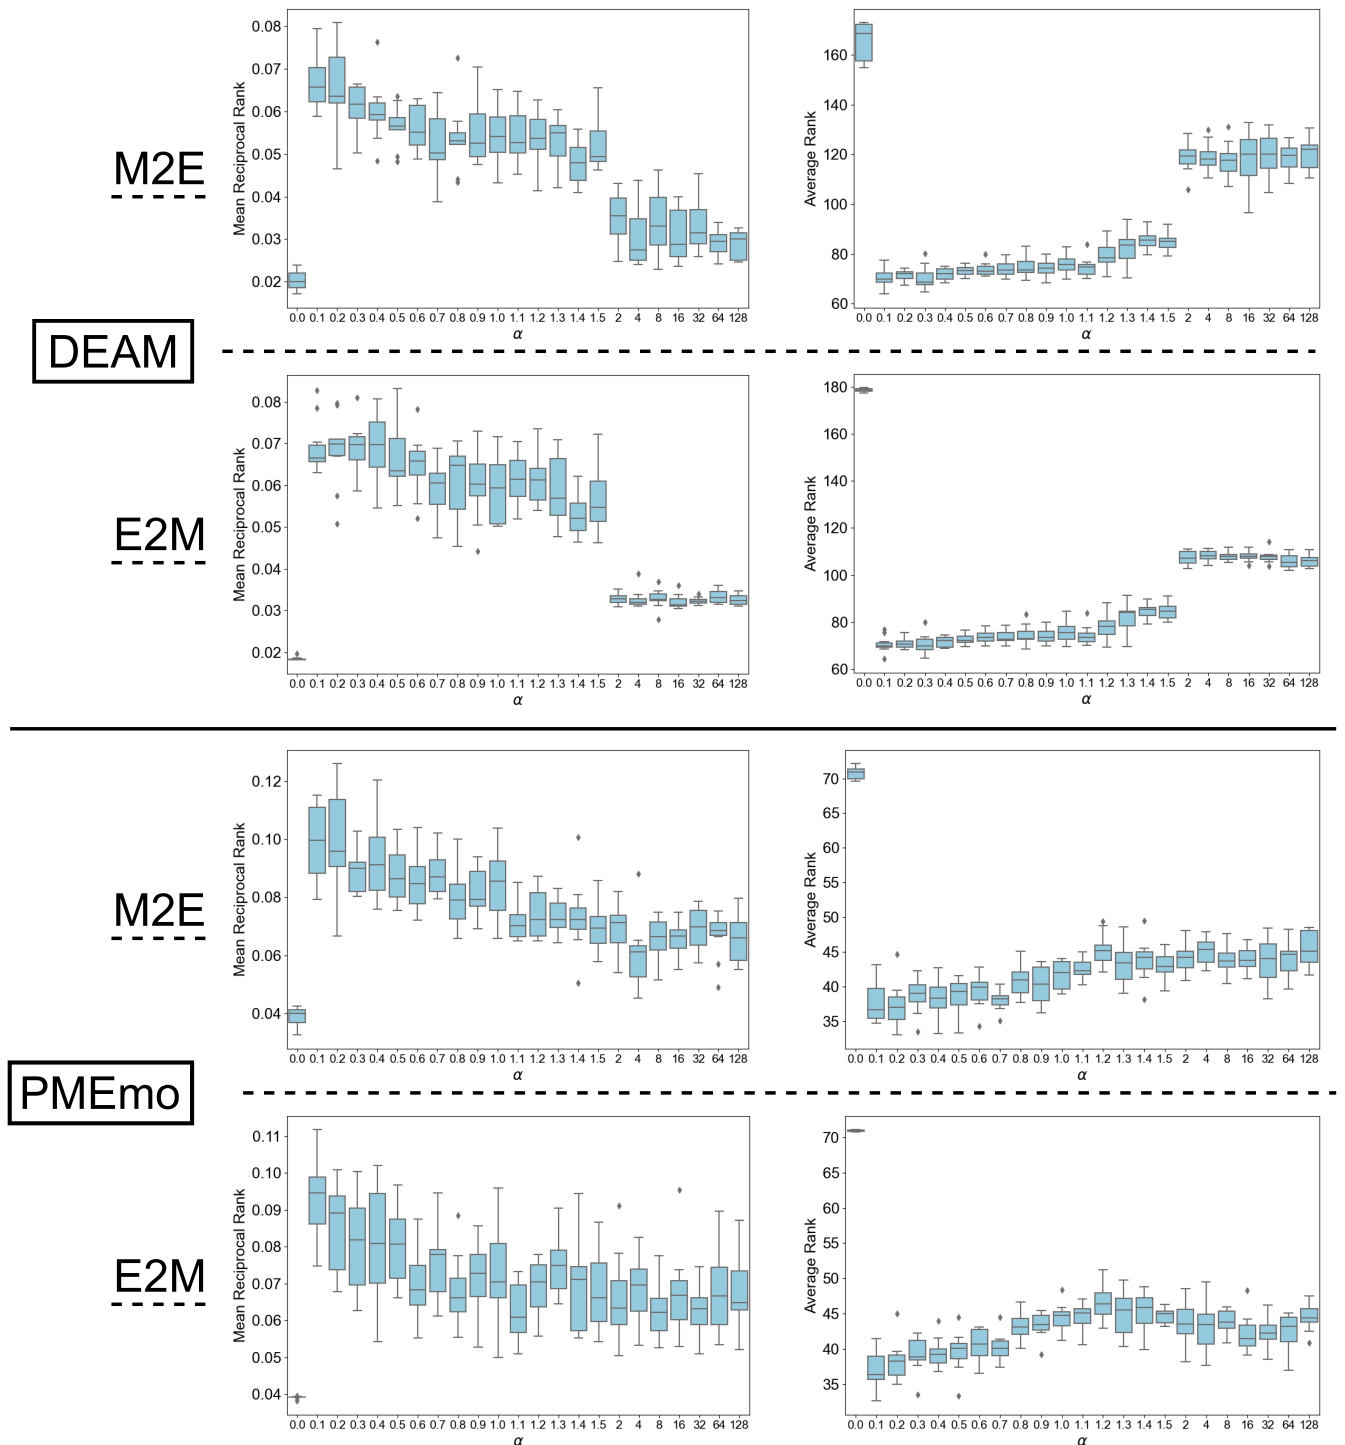


FIGURE 9. Transition of *Composite-C* performances (MRRs and ARs) obtained using *Cos-Loss* configured by different values of α . The optimal α value is selected as 0.3 and 0.1 for DEAM and PMEmo, respectively.

initialising all parameters in *Composite-C* and randomly splitting a dataset into training and test partitions with a ratio of 8 : 2. In addition, the computational cost of grid search can be reduced by considering that P is only related to the test process as described in Section A-A. More specifically, after training 10 *Composite-C* models (i.e., music and emotion

encoders, and embedding spaces based on *CCA-Loss* and *Cos-Loss*) using the α value optimised in the previous section and a specific λ value, their test processes are repeatedly run to obtain 10 performances for each of the different P values. Note that to make visual interpretation of the results easier, the axes of λ and P are depicted in the horizontal and depth

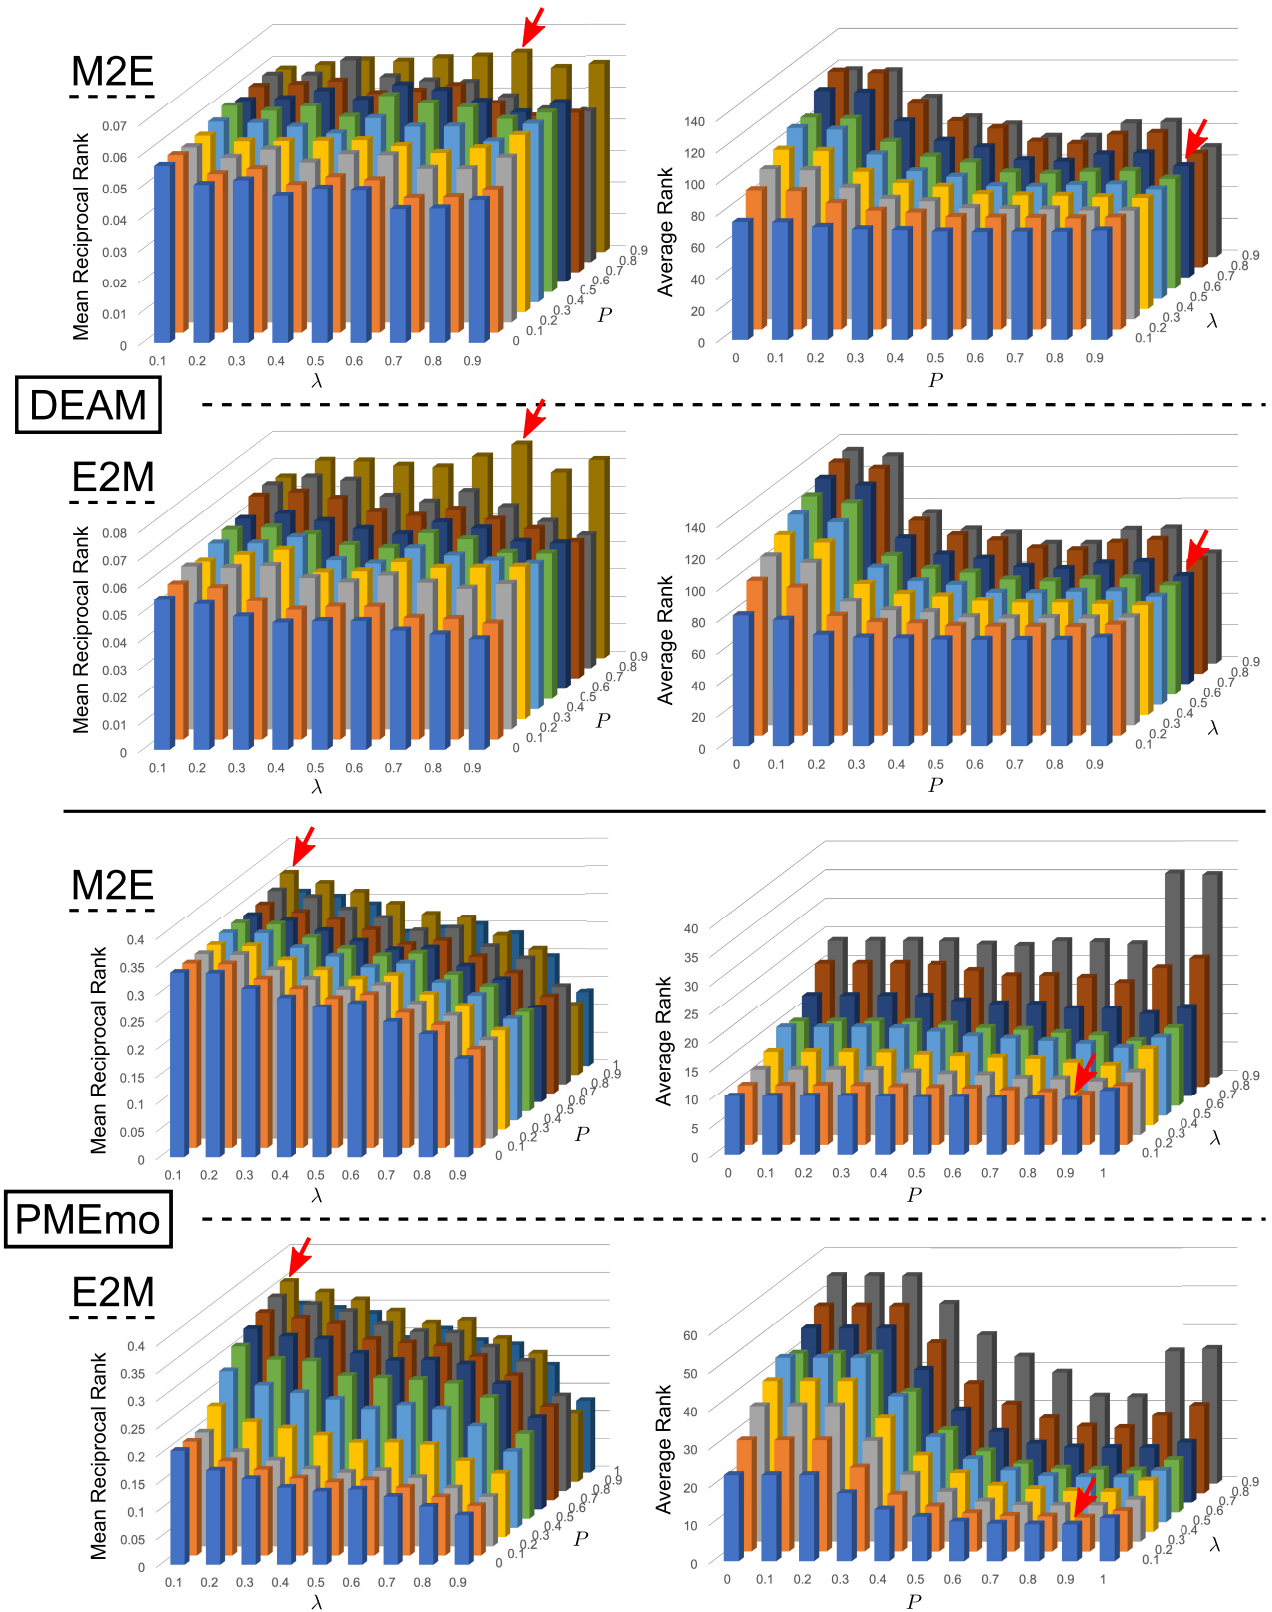


FIGURE 10. Transitions of *Composite-C* performances (MRRs and ARs) obtained by different pairs of P and λ values. Here, α of *Cos-Loss* is set to the optimal values found from Fig. 9. As indicated by the red arrows, the optimal pair of P and λ values is selected as ($P = 0.9, \lambda = 0.7$) and ($P = 0.9, \lambda = 0.1$) for DEAM and PMEmo, respectively.

TABLE 4. Performance comparison between EMER-CL using the optimal hyper-parameter values (*Composite*) and EMER-CL using the default ones (*Composite-d*).

DEAM				
ModelType	M2E		E2M	
	MRR	AR	MRR	AR
<i>Composite</i>	0.128 ± 0.011	56.0 ± 1.4	0.114 ± 0.011	55.3 ± 1.8
<i>Composite-d</i>	0.112 ± 0.001	57.4 ± 3.2	0.108 ± 0.006	56.5 ± 3.6

PMemo				
ModelType	M2E		E2M	
	MRR	AR	MRR	AR
<i>Composite</i>	0.544 ± 0.056	3.8 ± 0.6	0.542 ± 0.051	3.8 ± 0.5
<i>Composite-d</i>	0.538 ± 0.037	4.0 ± 0.5	0.497 ± 0.037	4.5 ± 0.6

directions for the MRR histograms, while the directions of these axes are swapped for the AR histograms in Fig. 10.

As indicated by the red arrows in Fig. 10, the best overall performance is attained using $P = 0.9$ and $\lambda = 0.7$ for DEAM and $P = 0.9$ and $\lambda = 0.1$ for PMemo. The performances obtained using these optimal P and λ are reported as the ones of *Composite-C* in the comparative study in Section IV-D. The fact that the *Composite-C* performances are significantly worse than the ones of *Composite* verifies the effectiveness of our proposed EMER-CL based on distribution-based embeddings.

APPENDIX C TUNING BASELINES

We explain how to tune the hyper-parameters of the regression-based baselines, *RegMLP-M2E*, *RegMLP-E2M*, *RegBiGRU-M2E* and *RegBiGRU-E2M*, used in Section IV-E. *RegMLP-M2E* and *RegMLP-E2M* employ an MLP, and *RegBiGRU-M2E* and *RegBiGRU-E2M* adopt an RNN based on bidirectional GRU. The inputs of *RegBiGRU-M2E* and *RegMLP-M2E* are a sequence $X^{(m)}$ of 128-dimensional acoustic features extracted by VGGish, and their mean $\bar{x}^{(m)}$, respectively. The outputs of these baselines are a 2-dimensional emotion vector $\bar{x}^{(e)}$ that is a prediction of the average arousal and valence for the music sample corresponding to $X^{(m)}$ or $\bar{x}^{(m)}$. In contrast, *RegBiGRU-E2M* and *RegMLP-E2M* take as input an arousal/valence sequence $X^{(e)}$ of two-dimensional emotion vectors, and their mean $\bar{x}^{(e)}$, respectively. These baselines output a 128-dimensional acoustic feature $\bar{x}^{(m)}$ as a prediction of the mean acoustic feature for the music sample associated with $X^{(e)}$ or $\bar{x}^{(e)}$. Below, we describe hyper-parameter tuning for hidden layers used between the above-mentioned inputs and outputs.

Using the same music and emotion encoder architectures as the ones optimised for EMER-CL (as reported in Section IV-C) did not yield satisfactory performances for the

regression-based baselines. More specifically, the M2E baseline obtained an MRR of 0.035 ± 0.001 and AR of 109.9 ± 0.3 on DEAM, and an MRR of 0.071 ± 0.004 and AR of 46.5 ± 0.7 on PMemo. The E2M baseline yielded an MRR of $0.018 \pm 2e^{-6}$ and AR of 181.0 ± 0.006 on DEAM, and an MRR of 0.057 ± 0.003 and AR of 65.9 ± 1.2 on PMemo. These performances are orders of magnitude worse than the ones of EMER-CL reported in our paper. For this reason, we attempted to improve the performances of the baselines by further tuning their hyper-parameters.

We focus especially on the number of hidden layers and the number of units per hidden layer, and carry out grid search on them. More specifically, we tested configurations involving a number of layers between one and five, and a number of units per layer in $\{16, 32, 64, 128, 256, 512\}$. That is, when using one, two, three, four and five hidden layers, the numbers of possible architectures are 6^1 , 6^2 , 6^3 , 6^4 and 6^5 , respectively. Thus, our grid search examines in total 9330 architectures. Note that for *RegBiGRU-M2E* and *RegBiGRU-E2M*, the first hidden layer is defined as a bidirectional GRU layer while the others are defined as FC layers. For *RegMLP-M2E* and *RegMLP-E2M*, all the hidden layers are defined as FC layers. The activation function of units in FC layers is always soft-plus, and a dropout rate of 0.5 is applied to FC layers that have 256 or 512 units. For a bidirectional GRU layer, the activation functions are used as specified in [49], and no dropout is employed. Furthermore, the MRR of each architecture is monitored every 100 epochs until 10000 epochs, and the model trained at the epoch that yielded the highest MRR is selected for this architecture. One exception is that this performance monitoring is continued until 50000 epochs unless the training converges after 10000 epochs. It should be noted that this setting to train the baselines is very favourable, considering that the performance of EMER-CL (i.e., *Composite*) is evaluated only after the last epoch (5001 and 10001 epochs for DEAM and PMemo, respectively).

Finally, the architecture with the highest MRR after grid search is always used for the comparison with EMER-CL for all baselines, with one exception. In some configurations, an architecture might yield a slightly lower MRR than the highest one, but also a significantly better (lower) AR. It can be considered that the architecture with the highest MRR is relatively subject to overfitting, in which for some queries (i.e., query music samples or query emotions), their associated samples (i.e., emotions or music samples) are ranked at very high positions, while the ranks of samples associated with other queries are very low. In these configurations, we manually select the architecture with the notably better AR among the ones returning the five highest MRRs.

Table 5 shows the selected architecture for each of *RegMLP-M2E*, *RegMLP-E2M*, *RegBiGRU-M2E* and *RegBiGRU-E2M* on DEAM and PMemo. As seen from Table 2, these baselines are significantly outperformed by EMER-CL despite the comprehensive and careful hyper-parameter tuning strategy. This verifies the effectiveness of EMER-CL's embedding approach.

TABLE 5. Hyper-parameters of the baselines after grid search.

DEAM		
Model name	Epochs	Number of units per hidden layer
RegMLP-M2E	8300	[128, 64, 16, 128, 512]
RegMLP-E2M	19200	[32]
RegBiGRU-M2E	2000	[512, 64, 32, 16, 32]
RegBiGRU-E2M	9200	[128, 128, 32, 128, 128]

PMemo		
Model name	Epochs	Number of units per hidden layer
RegMLP-M2E	5500	[16]
RegMLP-E2M	9900	[32, 128, 128]
RegBiGRU-M2E	7500	[512, 32, 32]
RegBiGRU-E2M	3400	[512]

APPENDIX D
DETAILED ANALYSIS RESULTS FOR M2E

Fig. 11 shows the average cosine similarities computed for M2E on DEAM and PMemo. Cosine similarities are computed between the mean acoustic feature of a query music sample, and the one of the music sample associated to each emotion ranked in the top 5% by EMER-CL. A high average cosine similarity means that EMER-CL recognises emotions expressed in music samples that are acoustically similar to a query music sample. In other words, even if EMER-CL fails to properly recognise the ground-truth emotion associated with the query music sample, emotions reasonably relevant to the query are still recognised because they are associated with acoustically similar music samples to the query.

Just like Fig. 3, the bar graphs in the left part of Fig. 11 (a) and (b) are drawn by sorting query music samples in ascending order of the ranks r_q of their associated ground-truth emotions. The more to the left a query music sample is located, the higher its associated emotion is ranked by EMER-CL. It should be noted that the median of pairwise cosine similarities among acoustic features of music samples is 0.674 and 0.800 on DEAM and PMemo, respectively. Each of these numbers can be interpreted as the cosine similarity between the acoustic features of two randomly selected music samples. The bar graphs and box plots in Fig. 11 show that the median of average cosine similarities is 0.753 and 0.837 on DEAM and PMemo, respectively. These medians are significantly higher than the median of pairwise cosine similarities, which validates the meaningfulness of EMER-CL’s M2E results. In addition, as shown by the box plots in Fig. 11, the 25th percentile (first quartile) of average cosine similarities on DEAM and PMemo are 0.709 and 0.810, respectively. The fact that even these 25th percentiles are higher than the medians of pairwise cosine similarities, indicates that in most cases EMER-CL’s M2E works better than randomly selecting an emotion for a query music sample.

APPENDIX E
ABOUT THE DETAILED ANALYSIS FOR E2M ON PMemo

Fig. 3 (b) shows that the average cosine similarities in E2M on PMemo are very high, with a median of 0.994 close to

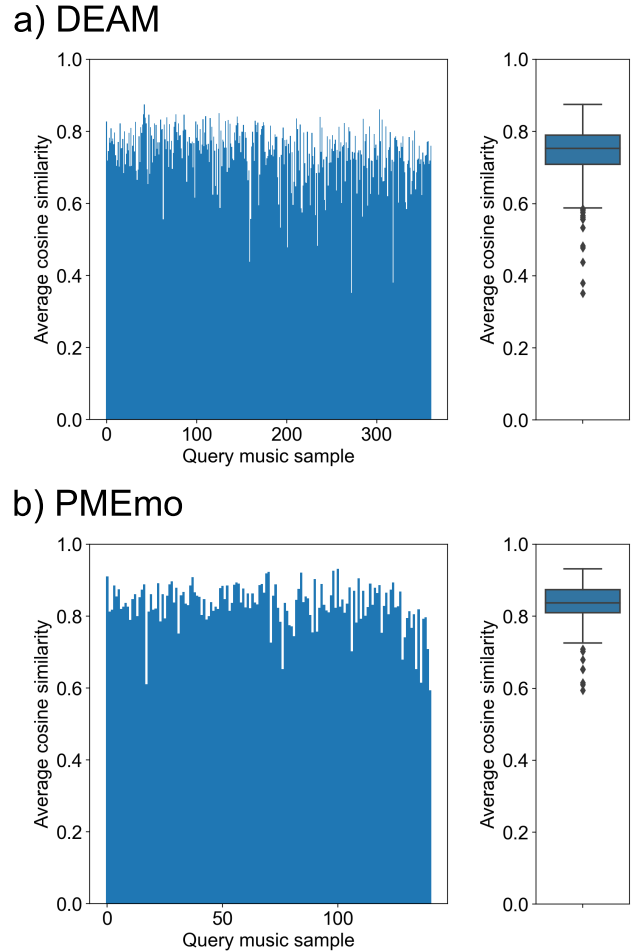


FIGURE 11. Bar graphs showing average cosine similarities and their box plots for M2E on DEAM and PMemo.

the maximum value of 1. The reasons behind this are discussed below. Fig. 12 depicts the distribution of 141 test emotions, each of which is represented by the two-dimensional mean emotion vector $\bar{x}^{(e)}$ of an arousal/valence sequence $X^{(e)}$. Here, arousal and valence intensities in PMemo are in [0, 1], so all the emotions are distributed only in the first quadrant. Furthermore, the variance of this distribution is small in particular, as illustrated in Fig. 12. As a result, the average of pairwise cosine similarities among these emotions is 0.9884 ± 0.0196 , and even the smallest pairwise cosine similarity characterised by the emotions marked by the crosses in Fig. 12 is 0.7846.

It should be noted that if the emotion ratings in Fig. 12 could be translated so that their origin is (0, 0), the range of average cosine similarities would be much wider and a detailed analysis based on them would be much clearer. However, it is difficult to precisely locate the origin of the arousal/valence space in Fig. 12. For example, z-normalisation can be carried out so that emotions have zero mean and unit variance, but there is no guarantee that the resulting zero vector corresponds to the origin. Cosine




FIGURE 12. Distribution of PMEmo test emotions in the valence/arousal space.

similarities significantly rely on the location of the origin because they measure similarities between the angles of two vectors. In addition, considering Russell’s circumplex model where emotions are circularly located in the arousal/valence space [4], an angle-based similarity measure like cosine similarity is preferred to other types of measures like Euclidean distance. Thus, reliable analysis is impossible when the exact location of the origin is unknown. For this reason, our detailed analysis on PMEmo is performed using the original arousal and valence intensities without any modification.

APPENDIX F
COMPUTATIONAL COSTS OF EMER-CL

We discuss about the computational costs of EMER-CL by referring to Table 6. Each number in this table indicates the average runtime of 10 runs to train or test an EMER-CL model with a random initialisation of all parameters and a random 8 : 2 split of a dataset into training and test partitions. The runtime for testing the EMER-CL model is the average elapsed time to get an M2E (or E2M) result given a query music sample (or query emotion). The hyper-parameters of the EMER-CL model are set to the optimal values found in Section A. Furthermore, the runtimes in Table 6 are measured using a computer equipped with Intel i9-7900X CPU, 128GB RAM and NVIDIA RTX 2080Ti GPU. The codes are written using the TensorFlow library (version 1.15) based on CUDA version 11.2.⁷

As it can be seen from the first column in Table 6, training an EMER-CL model on PMEmo takes significantly longer time than training it on DEAM, even though the training partition of PMEmo only contains 561 music-emotion pairs. One main reason is that both of the music and emotion encoders for PMEmo are defined as RNNs that need to process the acoustic feature or emotion vector sequentially in time. To improve the scalability of EMER-CL, we plan to define the music and emotion encoders as self-attention

TABLE 6. Runtimes of EMER-CL.

DEAM		
Training	Test (M2E)	Test (E2M)
2118.45 sec	0.0046 sec	0.0045 sec
PMEmo		
Training	Test (M2E)	Test (E2M)
9472.44 sec	0.0025 sec	0.0019 sec

models that can perform batch processing of acoustic features or emotion vectors at all times [65].

The second and last columns in Table 6 show the very short runtimes of EMER-CL’s test process. This is due to the fact that the test partitions of DEAM and PMEmo only contain 361 and 141 samples (i.e., music samples for E2M or emotions for M2E), respectively. Nevertheless, since the runtime of EMER-CL’s test process scales linearly with the number of samples, the test process would still be expected to finish within seconds even with a number of samples three orders of magnitude higher than the currently tested numbers.

APPENDIX G
T-SNE PLOTS OF EMBEDDINGS ON THE PMEmo DATASET

In a similar way as for the DEAM dataset, we plotted the t-SNE projections of the correlation-based and probabilistic-based embeddings produced by both the music ($\phi^{(m)}$, $\mu^{(m)}$ and $\sigma^{(m)}$) and emotion models ($\phi^{(e)}$, $\mu^{(e)}$ and $\sigma^{(e)}$) trained on the PMEmo dataset. Plots were obtained for examples in both the training and testing sets. For each example of either dataset, we first computed the average emotion vector over time $\bar{x}_j^{(e)}$ and used it to label the t-SNE projections in terms of which quadrant in the arousal/valence space (HA/HV, HA/LV, LA/LV or LA/HV) each music sample $X_j^{(m)}$ or mean emotion sequence $x_j^{(e)}$ was associated with.

Unlike for the DEAM dataset, the definition of the emotional quadrants on PMEmo is not trivial. This is due to the fact that both arousal and valence ratings are projected into the [0, 1] interval, and that the dataset includes exclusively pop songs which skew it strongly towards the quadrant high arousal/high valence (HA/HV). Taking the centre of the emotional space (0.5, 0.5) as cut-off point for the definition of the quadrants leads to a very imbalanced data repartition with 476, 53, 106 and 66 songs out of 701 associated to the HA/HV, HA/LV, LA/LV and LA/HV quadrants, respectively. To mitigate the effects of this imbalance, we instead take the barycentre of all 701 PMEmo music samples in the arousal/valence space (0.61, 0.63) as mid-point to define the quadrants. This leads to a more balanced repartition of 306, 77, 84 and 234 songs associated to the HA/HV, HA/LV, LA/LV and LA/HV quadrants, respectively. The t-SNE plots of embeddings obtained with these emotion annotations are

⁷The codes are available at https://mu-lab.info/naoki_takashima/emmer-cl

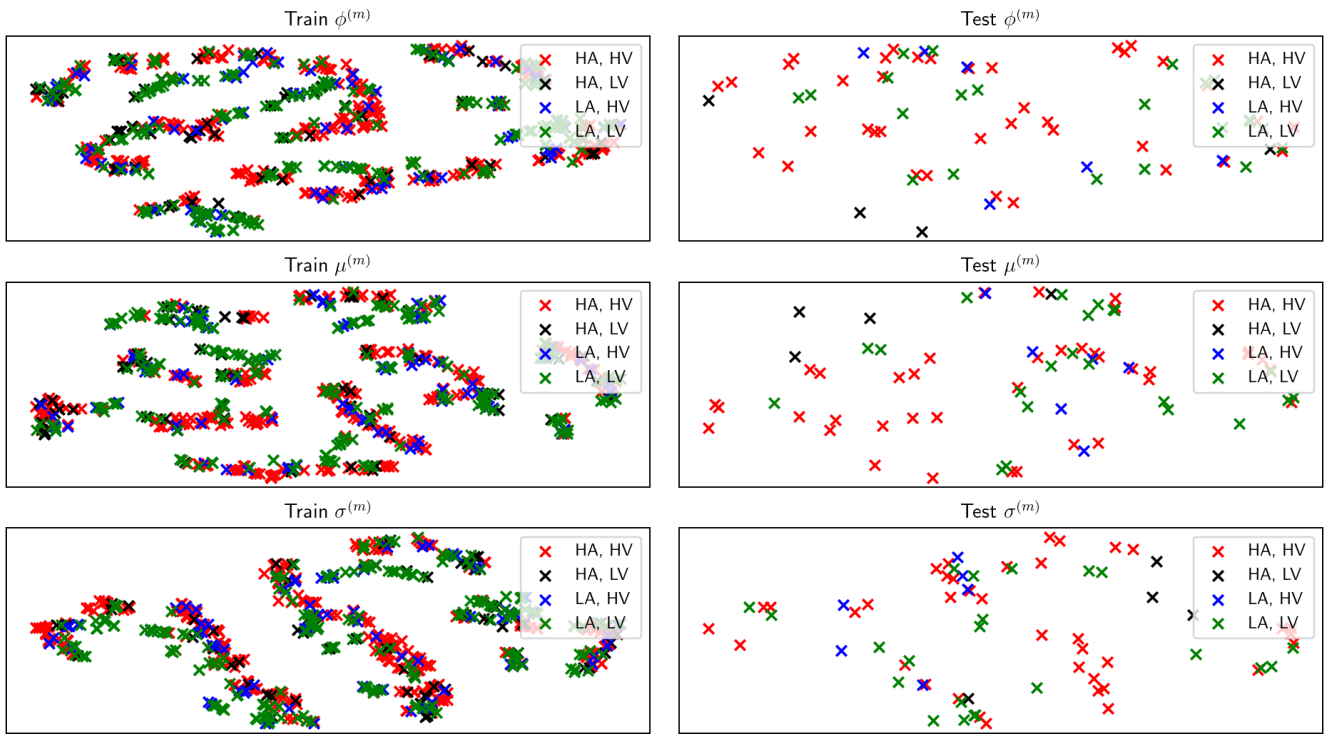


FIGURE 13. t-SNE projections of the music embeddings $\phi^{(m)}$, $\mu^{(m)}$ and $\sigma^{(m)}$ (respectively first, second and third rows) on the PMEmo training (left column) and testing (right column) sets. The t-SNE projections are labelled with their associated emotional quadrants, i.e., high arousal/high valence (HA, HV), high arousal/low valence (HA, LV), low arousal/low valence (LA, LV) and low arousal/high valence (LA, HV).

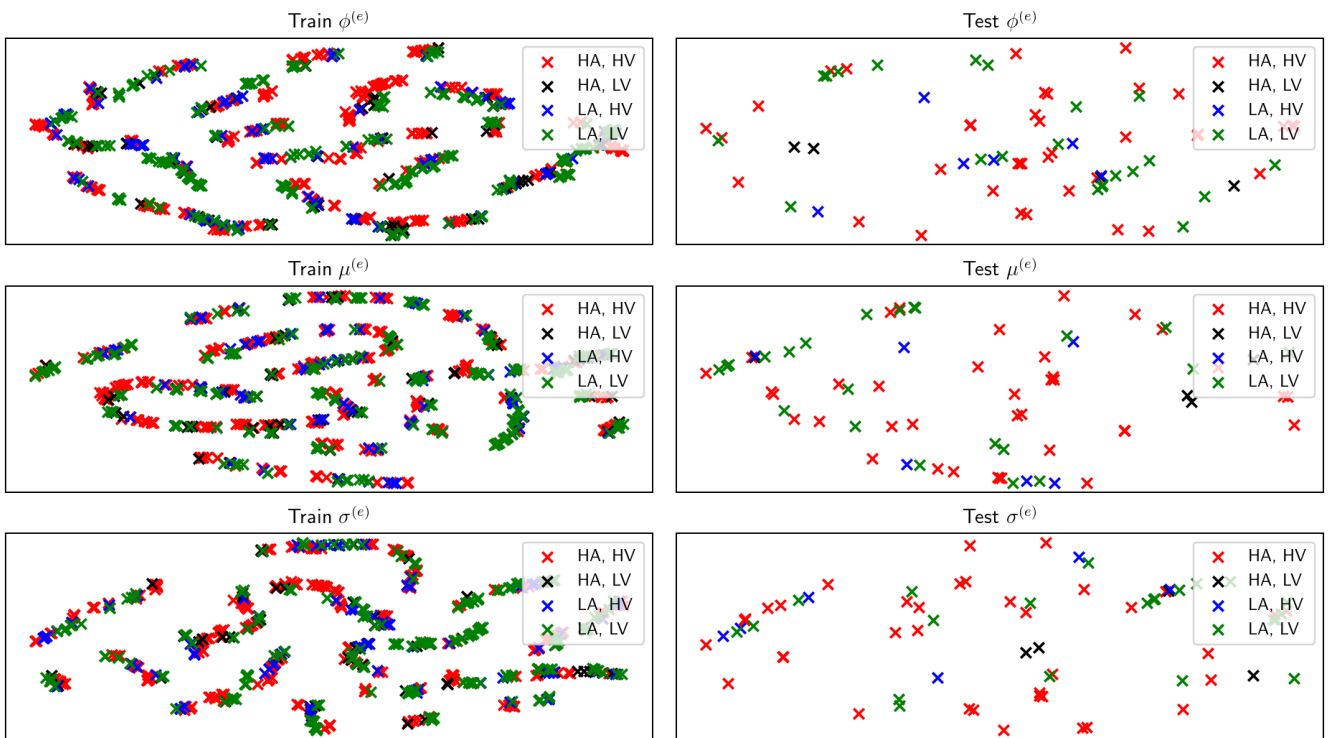


FIGURE 14. t-SNE projections of the emotion embeddings $\phi^{(e)}$, $\mu^{(e)}$ and $\sigma^{(e)}$ (respectively first, second and third rows) on the PMEmo training (left column) and testing (right column) sets. The t-SNE projections are labelled with their associated emotional quadrants, i.e., high arousal/high valence (HA, HV), high arousal/low valence (HA, LV), low arousal/low valence (LA, LV) and low arousal/high valence (LA, HV).

provided in Figs. 13 and 14 for the music and emotion embeddings, respectively.


These figures show that embeddings tend to be grouped by their associated emotional quadrants on both the training

and testing sets, although this trend on the PMemo dataset is less clear than the one for the DEAM dataset, as some groups mixing embeddings associated with different quadrants can also be seen. We hypothesise that this is due to the fact that the PMemo dataset aggregates music samples belong to a single genre (pop songs) that are quite similar to each other in terms of audio features and elicited emotions. Despite this, music samples or emotions associated to the same quadrants still tend to be located in the same neighbourhoods of the embedding space.

REFERENCES

- [1] D. Han, Y. Kong, J. Han, and G. Wang, "A survey of music emotion recognition," *Frontiers Comput. Sci.*, vol. 16, no. 6, Dec. 2022, Art. no. 166335.
- [2] S. Hizlisoy, S. Yildirim, and Z. Tufekci, "Music emotion recognition using convolutional long short term memory deep neural networks," *Eng. Sci. Technol., Int. J.*, vol. 24, no. 3, pp. 760–767, Jun. 2021.
- [3] P. Ekman, "An argument for basic emotions," *Cognition Emotion*, vol. 6, nos. 3–4, pp. 169–200, May 1992.
- [4] J. A. Russell, "A circumplex model of affect," *J. Personality Social Psychol.*, vol. 39, no. 6, p. 1161, Dec. 1980.
- [5] I. Dufour and G. Tzanetakis, "Using circular models to improve music emotion recognition," *IEEE Trans. Affect. Comput.*, vol. 12, no. 3, pp. 666–681, Jul. 2021.
- [6] J. S. Gomez-Canon, E. Cano, T. Eerola, P. Herrera, X. Hu, Y.-H. Yang, and E. Gomez, "Music emotion recognition: Toward new, robust standards in personalized and context-sensitive applications," *IEEE Signal Process. Mag.*, vol. 38, no. 6, pp. 106–114, Nov. 2021.
- [7] W. Chen, Y. Liu, W. Wang, E. M. Bakker, T. Georgiou, P. Fieguth, L. Liu, and M. S. Lew, "Deep learning for instance retrieval: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Nov. 1, 2022, doi: 10.1109/TPAMI.2022.3218591.
- [8] S. Chun, S. J. Oh, R. S. de Rezende, Y. Kalantidis, and D. Larlus, "Probabilistic embeddings for cross-modal retrieval," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 8411–8420.
- [9] D. Zeng, Y. Yu, and K. Oyama, "Audio-visual embedding for cross-modal music video retrieval through supervised deep CCA," in *Proc. IEEE Int. Symp. Multimedia (ISM)*, Dec. 2018, pp. 143–150.
- [10] D. Zeng, Y. Yu, and K. Oyama, "Deep triplet neural networks with cluster-CCA for audio-visual cross-modal retrieval," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 16, no. 3, pp. 1–23, Aug. 2020.
- [11] Y. Yu, S. Tang, F. Raposo, and L. Chen, "Deep cross-modal correlation learning for audio and lyrics in music retrieval," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 15, no. 1, pp. 1–16, Oct. 2019.
- [12] M. Won, J. Salamon, N. J. Bryan, G. J. Mysore, and X. Serra, "Emotion embedding spaces for matching music to stories," 2021, *arXiv:2111.13468*.
- [13] L. Shang, D. Zhang, J. Shen, E. L. Marmion, and D. Wang, "CCMR: A classic-enriched connotation-aware music retrieval system on social media with visual inputs," *Social Netw. Anal. Mining*, vol. 11, no. 1, p. 119, Dec. 2021.
- [14] A. Ferraro, X. Favory, K. Drossos, Y. Kim, and D. Bogdanov, "Enriched music representations with multiple cross-modal contrastive learning," *IEEE Signal Process. Lett.*, vol. 28, pp. 733–737, 2021.
- [15] N. Takashima, F. Li, M. Grzegorzec, and K. Shirahama, "Cross-modal music-emotion retrieval using DeepCCA," in *Proc. ITIB*. Cham, Switzerland: Springer, 2021, pp. 133–145.
- [16] Y.-H. Yang, C.-C. Liu, and H. H. Chen, "Music emotion classification: A fuzzy approach," in *Proc. ACM Multimedia*, 2006, pp. 81–84.
- [17] M. Won, S. Oramas, O. Nieto, F. Gouyon, and X. Serra, "Multimodal metric learning for tag-based music retrieval," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Jun. 2021, pp. 1–5.
- [18] G. Andrew, R. Arora, J. Bilmes, and K. Livescu, "Deep canonical correlation analysis," in *Proc. ICML*, 2013, pp. 1247–1255.
- [19] K. Hama, T. Matsubara, and K. Uehara, "Image-caption retrieval by embedding to Gaussian distribution," *IEICE Tech. Rep.*, vol. 118, no. 220, pp. 17–20, 2018.
- [20] Z. Ren, H. Jin, Z. Lin, C. Fang, and A. Yuille, "Joint image-text representation by Gaussian visual-semantic embedding," in *Proc. 24th ACM Int. Conf. Multimedia*, Oct. 2016, pp. 207–211.
- [21] A. Aljanaki, Y.-H. Yang, and M. Soleymani, "Developing a benchmark for emotional analysis of music," *PLoS ONE*, vol. 12, no. 3, Mar. 2017, Art. no. e0173392.
- [22] K. Zhang, H. Zhang, S. Li, C. Yang, and L. Sun, "The PMemo dataset for music emotion recognition," in *Proc. ICMR*, 2018, pp. 135–142.
- [23] O. Lartillot and P. Toivainen, "A MATLAB toolbox for musical feature extraction from audio," in *Proc. DAFX*, 2007, pp. 1–8.
- [24] F. Eyben, M. Wöllmer, and B. Schuller, "Opensmile: The Munich versatile and fast open-source audio feature extractor," in *Proc. 18th ACM Int. Conf. Multimedia*, Oct. 2010, pp. 1459–1462.
- [25] R. Panda, R. Malheiro, and R. P. Paiva, "Audio features for music emotion recognition: A survey," *IEEE Trans. Affect. Comput.*, vol. 14, no. 1, pp. 68–88, Jan. 2023.
- [26] S. Mo and J. Niu, "A novel method based on OMPGW method for feature extraction in automatic music mood classification," *IEEE Trans. Affect. Comput.*, vol. 10, no. 3, pp. 313–324, Jul. 2017.
- [27] R. Panda, R. Malheiro, and R. P. Paiva, "Novel audio features for music emotion recognition," *IEEE Trans. Affect. Comput.*, vol. 11, no. 4, pp. 614–626, Oct. 2020.
- [28] M. Malik, S. Adavanne, K. Drossos, T. Virtanen, D. Ticha, and R. Jarina, "Stacked convolutional and recurrent neural networks for music emotion recognition," 2017, *arXiv:1706.02292*.
- [29] Y. Dong, X. Yang, X. Zhao, and J. Li, "Bidirectional convolutional recurrent sparse network (BCRSN): An efficient model for music emotion recognition," *IEEE Trans. Multimedia*, vol. 21, no. 12, pp. 3150–3163, Dec. 2019.
- [30] R. Sarkar, S. Choudhury, S. Dutta, A. Roy, and S. K. Saha, "Recognition of emotion in music based on deep convolutional neural network," *Multimedia Tools Appl.*, vol. 79, nos. 1–2, pp. 765–783, Jan. 2020.
- [31] K. Choi, G. Fazekas, M. Sandler, and K. Cho, "Transfer learning for music classification and regression tasks," 2017, *arXiv:1703.09179*.
- [32] E. Koh and S. Dubnov, "Comparison and analysis of deep audio embeddings for music emotion recognition," 2021, *arXiv:2104.06517*.
- [33] R. Orjesek, R. Jarina, and M. Chmulik, "End-to-end music emotion variation detection using iteratively reconstructed deep features," *Multimedia Tools Appl.*, vol. 81, no. 4, pp. 5017–5031, Feb. 2022.
- [34] N. He and S. Ferguson, "Music emotion recognition based on segment-level two-stage learning," *Int. J. Multimedia Inf. Retr.*, vol. 11, no. 3, pp. 383–394, Sep. 2022.
- [35] S. Hershey, "CNN architectures for large-scale audio classification," in *Proc. ICASSP*, Mar. 2017, pp. 131–135.
- [36] Y.-H. Yang, Y.-C. Lin, Y.-F. Su, and H. H. Chen, "A regression approach to music emotion recognition," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 16, no. 2, pp. 448–457, Feb. 2008.
- [37] Y.-H. Yang and H. H. Chen, "Ranking-based emotion recognition for music organization and retrieval," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 4, pp. 762–774, May 2011.
- [38] Y.-H. Yang and H. H. Chen, "Prediction of the distribution of perceived music emotions using discrete samples," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 7, pp. 2184–2196, Sep. 2011.
- [39] Y.-H. Chin, J.-C. Wang, J.-C. Wang, and Y.-H. Yang, "Predicting the probability density function of music emotion using emotion space mapping," *IEEE Trans. Affect. Comput.*, vol. 9, no. 4, pp. 541–549, Oct. 2018.
- [40] K. Markov and T. Matsui, "Music genre and emotion recognition using Gaussian processes," *IEEE Access*, vol. 2, pp. 688–697, 2014.
- [41] S. Fukayama and M. Goto, "Music emotion recognition with adaptive aggregation of Gaussian process regressors," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2016, pp. 71–75.
- [42] J. C. Wang, Y. H. Yang, H. M. Wang, and S. K. Jeng, "Modeling the affective content of music with a Gaussian mixture model," *IEEE Trans. Affective Computing*, vol. 6, no. 1, pp. 56–68, Jan. 2015.
- [43] J.-C. Wang, H.-M. Wang, and G. Lanckriet, "A histogram density modeling approach to music emotion recognition," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Apr. 2015, pp. 698–702.
- [44] J.-C. Wang, Y.-S. Lee, Y.-H. Chin, Y.-R. Chen, and W.-C. Hsieh, "Hierarchical Dirichlet process mixture model for music emotion recognition," *IEEE Trans. Affect. Comput.*, vol. 6, no. 3, pp. 261–271, Jul. 2015.
- [45] F.-F. Kuo, M.-F. Chiang, M.-K. Shan, and S.-Y. Lee, "Emotion-based music recommendation by association discovery from film music," in *Proc. 13th Annu. ACM Int. Conf. Multimedia*, Nov. 2005, pp. 507–510.
- [46] M. M. Ruxanda, B. Y. Chua, A. Nanopoulos, and C. S. Jensen, "Emotion-based music retrieval on a well-reduced audio feature space," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Apr. 2009, pp. 181–184.

- [47] J. J. Deng, C. H. C. Leung, A. Milani, and L. Chen, "Emotional states associated with music: Classification, prediction of changes, and consideration in recommendation," *ACM Trans. Interact. Intell. Syst.*, vol. 5, no. 1, pp. 1–36, Mar. 2015.
- [48] Y. H. Yang, Y. C. Lin, H. T. Cheng, and H. H. Chen, "Mr. Emo: Music retrieval in the emotion plane," in *Proc. 16th ACM Int. Conf. Multimedia*, Oct. 2008, pp. 1003–1004.
- [49] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," 2014, *arXiv:1412.3555*.
- [50] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Trans. Signal Process.*, vol. 45, no. 11, pp. 2673–2681, Nov. 1997.
- [51] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*.
- [52] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek, "Image classification with the Fisher vector: Theory and practice," *Int. J. Comput. Vis.*, vol. 105, no. 3, pp. 222–245, 2013.
- [53] R. Kiros, R. Salakhutdinov, and R. S. Zemel, "Unifying visual-semantic embeddings with multimodal neural language models," 2014, *arXiv:1411.2539*.
- [54] W. Xin Zhao, J. Liu, R. Ren, and J.-R. Wen, "Dense text retrieval based on pretrained language models: A survey," 2022, *arXiv:2211.14876*.
- [55] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [56] E. M. Schmidt, J. J. Scott, and Y. E. Kim, "Feature learning in dynamic environments: Modeling the acoustic structure of musical emotion," in *Proc. ISMIR*, 2012, pp. 325–330.
- [57] F. Weninger, F. Eyben, and B. Schuller, "On-line continuous-time music mood regression with deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2014, pp. 5412–5416.
- [58] X. Li, J. Tian, M. Xu, Y. Ning, and L. Cai, "DBLSTM-based multi-scale fusion for dynamic emotion prediction in music," in *Proc. IEEE Int. Conf. Multimedia Expo. (ICME)*, Jul. 2016, pp. 1–6.
- [59] S. Hult, L. B. Kreiberg, S. S. Brandt, and B. P. Jónsson, "Analysis of the effect of dataset construction methodology on transferability of music emotion recognition models," in *Proc. Int. Conf. Multimedia Retr.*, Jun. 2020, pp. 316–320.
- [60] K. W. Cheuk, Y.-J. Luo, B. T. Balamurali, G. Roig, and D. Herremans, "Regression-based music emotion prediction using triplet neural networks," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2020, pp. 1–7.
- [61] S. Chapaneri and D. Jayaswal, "Deep Gaussian processes for music mood estimation and retrieval with locally aggregated acoustic Fisher vector," *Sādhanā*, vol. 45, p. 73, Mar. 2020.
- [62] J. D. Berardinis, A. Cangelosi, and E. Coutinho, "The multiple voices of musical emotions: Source separation for improving music emotion recognition models and their interpretability," in *Proc. ISMIR*, 2020, pp. 310–317.
- [63] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, Nov. 2008.
- [64] A. Jaiswal, A. R. Babu, M. Z. Zadeh, D. Banerjee, and F. Makedon, "A survey on contrastive self-supervised learning," *Technologies*, vol. 9, no. 1, p. 2, Dec. 2020.
- [65] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NIPS*, 2017, pp. 5998–6008.
- [66] M. E. Basiri, S. Nemati, M. Abdar, E. Cambria, and U. R. Acharya, "ABCDM: An attention-based bidirectional CNN-RNN deep model for sentiment analysis," *Future Gener. Comput. Syst.*, vol. 115, pp. 279–294, Feb. 2021.
- [67] W. Tao, C. Li, R. Song, J. Cheng, Y. Liu, F. Wan, and X. Chen, "EEG-based emotion recognition via channel-wise attention and self attention," *IEEE Trans. Affect. Comput.*, vol. 14, no. 1, pp. 382–393, Jan. 2023.



FRÉDÉRIC LI received the engineering degree (equivalent of the M.Sc. degree) with a specialization in robotics and embedded systems from ENSTA ParisTech, in 2015, and the Ph.D. degree from the University of Lübeck, Germany, in 2021. From 2016 to 2019, he was a Research Assistant with the Pattern Recognition Group, University of Siegen, Germany, under the supervision of Prof. Dr.-Ing. habil. M. Grzegorzec. Since 2019, he has been a Research Assistant with the Institute of Medical Informatics, University of Lübeck. His research interests include ubiquitous computing, time-series classification, feature extraction, deep learning, and transfer learning.



MARCIN GRZEGORZEC received the master's degree in computer science from the Silesian University of Technology, Gliwice, in 2002, the D.Eng. degree (Hons.) from the University of Erlangen-Nürnberg, in 2007, and the Habilitation degree from the AGH University of Science and Technology, Kraków, in 2014. He is a Full Professor of medical informatics with the University of Lübeck, where he is leading the Medical Data Science Laboratory. He and his team conceptualize, implement, and evaluate new algorithms for automated analysis of human-related data (e.g., wearable sensor data) and demonstrate their applicability in real-world scenarios. He has been a project leader in more than ten third-party funded research projects and acted as a supervisor in nine completed doctoral procedures. He has published more than 100 scientific peer-reviewed articles. His research interests include pattern recognition, machine learning, data science, and sensor data analysis for health-related applications.



KIMIAKI SHIRAHAMA received the B.E., M.E., and D.E. degrees in engineering from Kobe University, Japan, in 2003, 2005, and 2011, respectively. After working as an Assistant Professor with the Muroran Institute of Technology, Japan, he was a Postdoctoral Researcher with the Pattern Recognition Group, University of Siegen, Germany, from 2013 to 2018. He was an Associate Professor with Kindai University, Japan, from 2018 to 2023. He is currently an Associate Professor with Doshisha University, Japan. His research interests include multimedia data processing, machine learning, data mining, and sensor-based human activity recognition. He is a member of ACM SIGKDD, ACM SIGMM, the Institute of Image Information and Television Engineers in Japan (ITE), the Information Processing Society of Japan (IPSI), and the Institute of Electronics, Information and Communication Engineering in Japan (IEICE).



NAOKI TAKASHIMA received the B.E. and M.E. degrees in engineering from Kindai University, Osaka, Japan, in 2020 and 2022, respectively. During his M.E. degree, he did all of the work needed to prepare this article. He is currently with Speee Inc. His research interests include affective computing, music emotion recognition, and deep learning.