

JTAV: Jointly Learning Social Media Content Representation by Fusing Textual, Acoustic, and Visual Features

Hongru Liang¹, Haozheng Wang¹, Jun Wang², Shaodi You³,
Zhe Sun⁴, Jin-Mao Wei¹, Zhenglu Yang^{1*}

¹ CCCE, Nankai University, Tianjin, China

² College of Mathematics and Statistics Science, Ludong University, China

³ Data61-CSIRO, Australian National University, Canberra, Australia

⁴ RIKEN Head Office for Information Systems and Cybersecurity
Computational Engineering Applications Unit, Japan

{lianghr, hzwang, junwang}@mail.nankai.edu.cn, shaodi.you@data61.csiro.au,
zhe.sun.vk@riken.jp, {weijm, yangzl}@nankai.edu.cn

Abstract

Learning social media content is the basis of many real-world applications, including information retrieval and recommendation systems, among others. In contrast with previous works that focus mainly on single modal or bi-modal learning, we propose to learn social media content by fusing jointly textual, acoustic, and visual information (JTAV). Effective strategies are proposed to extract fine-grained features of each modality, that is, attBiGRU and DCRNN. We also introduce cross-modal fusion and attentive pooling techniques to integrate multi-modal information comprehensively. Extensive experimental evaluation conducted on real-world datasets demonstrates our proposed model outperforms the state-of-the-art approaches by a large margin.

1 Introduction

The popularity of the social media (e.g., Twitter, Weibo) over the last two decades has led to increasing demands for learning the content of social media, which may be beneficial to many real-world applications, such as sentiment analysis (Wang et al., 2015), information retrieval (Li et al., 2017), and recommendation systems (Wu et al., 2017). However, the task is non-trivial and challenging, because social media content is commonly multi-modal and involves several types of data, including text, audio, and image (an example is illustrated in Fig. 1). Each representation of the individual modality encodes specific knowledge and complementary, which can be explored to facilitate understanding of the entire meaning of the content. Therefore, gaining comprehensive knowledge from learning arising out of social media requires deliberate exploration of single-modal information and joint learning of the intrinsic correlation among various-modalities.




Figure 1: An example of multi-modal social media content, organized by a textual passage, a clip of the song “Muhammad Ali”, and a picture of Muhammad Ali’s famous victory.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

*Corresponding author

Previous works on learning social media content have focused mainly on single modality, e.g., textual information (Wu et al., 2017) and image information (Garimella et al., 2016). To our best knowledge, no studies have focused on mere acoustic information analysis for social media. These previous works cope with data in the early stage of social media, because users are limited to posting single-modal content. In contrast, today’s social media content usually includes textual, acoustic, and visual ingredients simultaneously, and requires multi-modal learning strategies to integrate rich information from all modalities.

Many researchers have studied the bi-modal utilization to address the issue of incomplete information exploration. This utilization can be classified into two kinds of approaches. The first is bi-modal retrieval, which utilizes one modality to find similar content in the other modality, such as finding the most relevant texts to a given image (Li et al., 2017). This kind of approach is similar to machine translation tasks to some extent, for instance, the input is fine-grained features of one modality, the objective of which is to minimize the distance between the output vectors with fine-grained features of the target modality. The second is bi-modal unification (Park and Im, 2016), which seeks to integrate individual information of the two modalities together. The process of this kind of approach can be divided into two stages: (a) feature extraction, in which the two modalities are encoded into their corresponding vectors or matrices via embedding approaches; and (b) bi-modal feature aggregation, in which the extracted features are mapped into a shared latent space through multilayer perceptrons or directly into the target spaces of tasks through classifiers (Chen et al., 2017).

Learning representations of multi-modal data, which contains features of more than two modalities, is a more challenging task. The main reason is that the fine-grained features of the modalities can be difficult to obtain because of the lack of annotated labels for every modality. Previous studies lack the appropriate strategies to aggregate various kinds of information effectively. It’s easy to imagine that bi-modal aggregation models might be extended to multi-modal learning with linear operations, e.g., the concatenation operation. However, this technique may generate inappropriate integral representation of the multiple features and lacks further exploration in the previous work.

We address the abovementioned issues by introducing effective strategies to represent the individual features for the three modalities. A general unified framework is proposed to integrate multi-modalities seamlessly. Specifically, we begin by encoding single modal parts into dense vectors. For textual content, we design an attention-based network (i.e., attBiGRU), to incorporate various textual information. For acoustic content, we introduce an effective DRCNN approach to embed temporal audio clips locally and globally. For visual content, we fine-tune a state-of-the-art general framework, DenseNet (Huang et al., 2017). We propose a novel fusion framework, which involves the cross-modal fusion and the attentive pooling strategies, to aggregate various modal information. Extensive experimental evaluations conducted on real-world datasets demonstrate that our proposed model outperforms state-of-the-art approaches by a large margin. The contributions of this paper are presented as follows:

- We introduce effective strategies to extract representative features of the three modalities, including the following. (1) For textual information, we design an attention based network, named as attBiGRU, to integrate various textual parts. The independent textual parts are considered as two separate roles, namely, the protagonist and the supporting players. The supporting players are encoded into an attention weight vector on the protagonist. (2) For acoustic information, we propose the DCRNN strategy based on the densely connected convolutional and recurrent networks. The densely connected convolutional networks are used to learn the acoustic features both locally and globally. The recurrent networks are designed to learn the temporal information. (3) For visual information, we introduce state-of-the-art strategies from (Simonyan and Zisserman, 2015; He et al., 2016; Huang et al., 2017).
- We propose a general multi-modal feature aggregation framework, JTAV, to learn information jointly. The framework considers the textual, acoustic, and visual contents to generate a unified representation of social media content with the help of the optimized feature fusion network, CMF-AP, which can learn inner and outer cross-modal information simultaneously.
- We conduct comprehensive experiments on two kinds of tasks, sentiment analysis and information

retrieval, to evaluate the effectiveness of the proposed JTAV framework. The experiments demonstrate that JTAV outperforms state-of-the-art approaches remarkably, that is, about 2% higher than the state-of-the-art approaches on all metrics in the sentiment analysis experiment and more than ten times better than the baseline on main metrics in music information retrieval experiment.

2 Related Work

Over the past decades, social media learning has gained considerable attention in related research literature. Social media content is utilized as raw material for recommendation (Wang et al., 2013), information retrieval (Agichtein et al., 2008), and so forth.

The dominance of the single-modal content in social media in the early stages, however, has caused most previous works to position themselves in single-modal exploration, e.g., natural language processing and image processing. (Han et al., 2013) focused on the short text messages in social media, and normalized lexical variants to their canonical forms. In (Hutto and Gilbert, 2014), combined lexical features of social media text were proposed for expressing and emphasizing sentiment intensity. An example of image processing was (Garimella et al., 2016), social media images, particularly geo-tagged images, were studied in predicting county-level health statistics.

In recent years, users have come to prefer presenting more attractive ingredients, such as image and audio in addition to natural words. Therefore, literature on social media analysis has been shifting its focus from single-modal to multi-modal learning. For example, (Wang et al., 2015) assume that visual and textual content have a shared sentiment label space, and the optimization problem can be converted to the minimization between the derivatives of both modalities. In (Li et al., 2017), the images and lyrics are mapped into a shared tag space. The lyrics, which have the smallest mean squared error with a given image, emerged as the matching object. (Chen et al., 2017) present a new end-to-end framework for visual and textual sentiment analysis; they began with the co-appearing image and text pairs that provide semantic information for each other, and use the concatenation of the vectors generated from the original image and text, in the shared sentiment space. However, the drawback of these works is that they utilized only partial information.

An intuitive idea to gain a comprehensive learning of the entire meaning of social media is to integrate more modalities effectively. The reason is that each representation of the individual modality encodes specific knowledge and is complementary, an aspect that can be explored to facilitate understanding on the entire meaning of the content. However, this task could be extremely challenging because we need to explore single-modal information deliberately and jointly learn the intrinsic correlation among various modalities. This work is the first study that focuses on the seamless integration of the three modalities (i.e., text, audio, and image), into a general framework to jointly learn the social media content.

3 Model Description

The architecture of our proposed framework is illustrated in Fig. 2. It has two main modules, the feature extraction and the feature aggregation. Let t , a , and v denote the textual, acoustic, and visual features, which are encoded from the feature extraction module, respectively. Let u denote the target unified representation of the multi-modal content. The feature aggregation module is designed to generate u from t , a , and v effectively and comprehensively. The example in Fig. 1 is utilized to ground our discussion and to illustrate the intuitive idea of this work.

3.1 Feature Extraction

Fine-grained features, which are related to the qualities of the input of the aggregation processing directly, are the foundation of multi-modal representation learning. The proposed strategies for generating appropriate features for text, audio, and image will be presented in the following subsections.

3.1.1 Text modeling

The purpose of text modeling is to encode raw textual parts into a latent representation t . We design a bidirectional gated recurrent neural network (BiGRU) based attentive network (i.e., attBiGRU), to extract sufficient textual features carrying the long-range dependencies, as illustrated in Fig. 3.

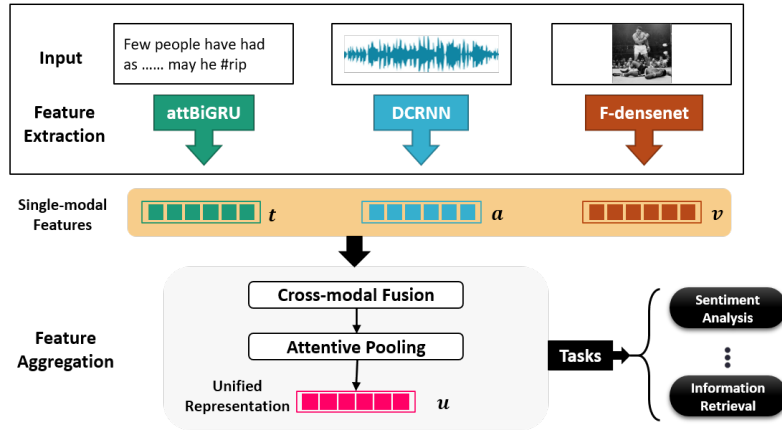


Figure 2: The architecture of the proposed JTAV framework




Figure 3: The attBiGRU strategy for modeling textual content

Generally, social media content contains more than one textual part, such as the lyrics and reviews of a song. It is reasonable to assume that a protagonist and supplying players exist among these textual parts. Intuitively and for sake of simplicity, we can regard the part with the maximum number of words as the protagonist and the other parts as the supplying players. Let T_{large} denote the protagonist and T_{small} denote the supplying roles. Inspired by (Li et al., 2017) to reduce the gap between image and text, we use a pre-trained word embedding model, e.g., FastText (Bojanowski et al., 2017), to encode T_{small} into a word matrix (i.e., \mathcal{T}_{small}). \mathcal{T}_{small} is computed as $w_{small}W_e$, where w_{small} is the words in T_{small} and W_e is a pre-trained embedding matrix. t_{att} is an attention vector generated from \mathcal{T}_{small} with an average pooling layer, and works on every word vector of \mathcal{T}_{large} through the element-wise product. We obtain an attended word matrix, denoted as \mathcal{T} .

The BiGRU encoder is utilized to encode every word vector in the sequence into a latent representation and receive the aid of the other textual parts. Subsequently, we feed the attended word matrix \mathcal{T} into a BiGRU network. The output of the BiGRU encoder is a $2M \times N$ matrix, where M is the number of hidden units of a layer, and N is the number of words in \mathcal{T}_{large} . We denote the output matrix as \mathcal{H} .

The attention with context part is a simplified version of (Yang et al., 2016) without the sentence attention part. It is introduced to balance the importance among all words and is defined as follows:

$$\begin{aligned}
 \hat{h}_i &= \tanh(W_u h_i + b_u), \\
 \alpha_i &= \frac{\exp(\hat{h}_i^\top \hat{h}_c)}{\sum_{j=0}^{N-1} \exp(\hat{h}_j^\top \hat{h}_c)}, \\
 t &= \sum_{j=0}^{N-1} \alpha_j h_j,
 \end{aligned} \tag{1}$$

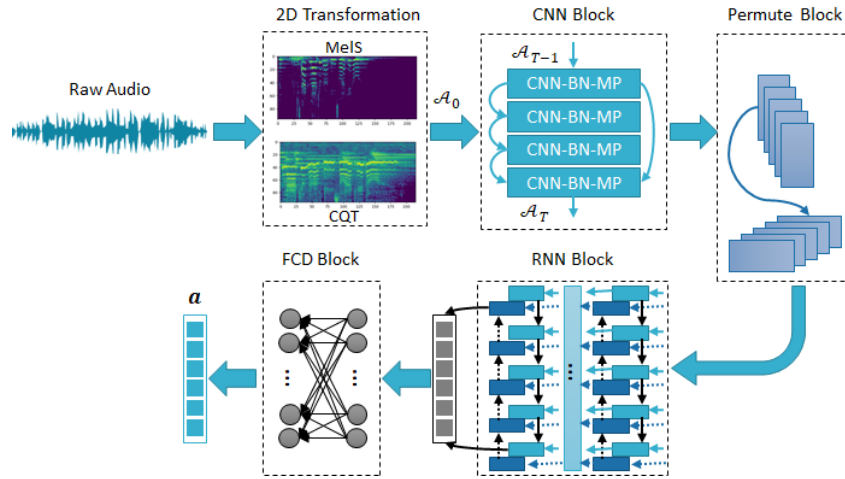


Figure 4: The DRCNN strategy for modeling acoustic content

where $\mathbf{h}_i (0 \leq i \leq N - 1)$ is the i^{th} column of \mathcal{H} , $\hat{\mathbf{h}}_i$ is a latent vector computed from \mathbf{h}_i using the tanh function, α_i is the attention weight scalar computed from $\hat{\mathbf{h}}_i$ and a random initialized context vector $\hat{\mathbf{h}}_c$, \mathbf{W}_u and \mathbf{b}_u are parameters learned from training. The final textual representation \mathbf{t} is a weighted sum of $\mathbf{h}_j (0 \leq j \leq N - 1)$ on $\alpha_j (0 \leq j \leq N - 1)$.

3.1.2 Audio modeling

In this paper, we focus on polyphonic audio, such as music clips. To capture useful acoustic features, we design a densely connected convolutional recurrent neural network called DCRNN, which takes raw audio as input, and generates a distributed representation \mathbf{a} to effectively represent the acoustic information. DCRNN consists of five phases, as illustrated in Fig. 4.

The first phase is named as 2D transformation, which transfers raw audio into two-dimensional vectors. Following the majorities of the deep learning approaches in music information retrieval (Choi et al., 2017b), we use two dimensional spectrograms instead of the discrete audio signals. We choose two kinds of spectrograms (i.e., mel-spectrogram (Boashash, 1996) and constant-Q transform (Brown and Puckette, 1992)). A mel-spectrogram (MelS) is a visual representation of the spectrum of frequencies of audio and is optimized for human auditory perception. A constant-Q transform (CQT) is computed with the logarithmic-scale of the central frequencies. The CQT is well suited with the frequency distribution of music pitch. MelS and CQT have been used frequently in various acoustic tasks (Zhu et al., 2017; Jansen et al., 2017), and are utilized as coarse-grained features in the current paper. We use \mathcal{A}_0 to indicate the 2D spectrograms.

The second phase is a convolutional neural network (CNN) based block (called CNN block). \mathcal{A}_0 is used as the input of the block. According to (Choi et al., 2017a), early layers of CNN have the ability to capture local information such as pitch and harmony, whereas deeper layers can capture global information such as melody. We introduce the densely connected CNN, which has been proved powerful at learning both local and global features of images (Huang et al., 2017). The CNN block is composed of several densely connected ‘‘CNN-BN-MP’’ sub-blocks, where ‘‘BN’’ represents batch normalization layers, and ‘‘MP’’ indicates max pooling layers. Suppose there are N densely connected ‘‘CNN-BN-MP’’ sub-blocks in the CNN block, each can be defined as:

$$\begin{aligned}
 \mathcal{A}_h^1 &= \text{CNN} - \text{BN} - \text{MP}(\mathcal{A}_{T-1}), \\
 \mathcal{A}_h^2 &= \text{CNN} - \text{BN} - \text{MP}(\mathcal{A}_h^1), \\
 \mathcal{A}_h^3 &= \text{CNN} - \text{BN} - \text{MP}(\mathcal{A}_h^2), \\
 \mathcal{A}_T &= \text{CNN} - \text{BN} - \text{MP}(\mathcal{A}_h^1 \oplus \mathcal{A}_h^3),
 \end{aligned} \tag{2}$$

where $T \in (1, N)$, \mathcal{A}_{T-1} represents the output of the $(T - 1)^{th}$ densely connected ‘‘CNN-BN-MP’’ sub-block, \mathcal{A}_h^* indicate the latent variables learned by the CNN block, \oplus denotes the concatenation operation, and \mathcal{A}_T is the output of the T^{th} sub-block.

The permute block is adopted from the technique introduced in (Panwar et al., 2017). The output of the CNN block is permuted to time major form, and fed into the following recurrent neural network.

The recurrent neural network (RNN) block is built based on bidirectional RNN. The stacked RNN layers are used to learn the long-short term temporal context information. We take the concatenation of the last hidden state (in vector format) in the forward direction and the first hidden state in the backward direction as the output of this phase.

Several fully connected dense layers constitutes the fully connected dense (i.e., FCD) block. We use the output of the last dense layer as the acoustic representation \mathbf{a} . Note that all the parameters presented in this section, e.g., number of layers, are optimally tuned and will be clarified in the experiment section.

3.1.3 Image modeling

Extracting rewarding features from images is a well explored problem; hence, following the tradition of other image-related tasks (Li et al., 2017; Chen et al., 2017), we use pre-trained approaches on external visual tasks.

VGG (Simonyan and Zisserman, 2015) is a deep convolutional network with nineteen weight layers, ResNet (He et al., 2016) is a residual learning framework with 101 weight layers, and DenseNet is a deep densely connected convolutional network with 121 weight layer. All these models are pre-trained on the ImageNet dataset ¹, and are fine-tuned on the experimental datasets by modifying the final densely-connected layers. The last dense layer is used as the visual representation \mathbf{v} .

3.2 Feature Aggregation

In this section, we propose a feature aggregation network, which consists of the cross-modal fusion (CMF) and the attentive pooling (AP) strategies to generate a unified representation of multi-modal information, as shown in the bottom part of Fig. 2. We call the feature aggregation approach as CMF-AP.

The first phase is cross-modal fusion. We introduce the concepts of the inner and outer information. The inner information refers to information remaining in single modality, like the visual information of an image or the acoustic information of a song clip. The outer mutual information cross modalities refers to the relevant information between two modalities. In addition to the inner information generated from the feature extraction module, we compute the matrix product (\times) pairwise on all modal vectors to calculate the cross-modal aware matrices. The cross-modal aware matrices contain outer mutual information cross modalities and the inner information.

However, handling the cross-modal aware matrices directly is not feasible, and thus, we reconstruct \mathbf{t} , \mathbf{a} , and \mathbf{v} from the cross-modal aware matrices. The attentive pooling strategy is adapted from Eq. 1 to conduct row pooling and column pooling.

After the above procedure, two reconstructed \mathbf{t} , two reconstructed \mathbf{a} , and two reconstructed \mathbf{v} are obtained. For instance, from the text-image matrix, we can obtain a text-enhanced visual vector and an image-enhanced textual vector. The reconstructed vectors carry both inner and outer information. Densely connected layers are deployed for dimensionality reduction purpose. The final unified representation \mathbf{u} is the concatenation of the reconstructed vectors.

The feature aggregation module is formulated as follows:

¹<http://www.image-net.org>.

$$\begin{aligned}
& \mathbf{C}_{mn} = \mathbf{m} \times \mathbf{n}^\top; \\
& \begin{cases} \mathbf{C}_m = \tanh(\mathbf{W}_m \mathbf{C}_{mn} + \mathbf{b}_m), \\ \boldsymbol{\alpha}_m = \text{softmax}(\mathbf{C}_m^\top \mathbf{u}_m), \\ \hat{\mathbf{m}} = \sigma(\hat{\mathbf{W}}_m \mathbf{C}_{mn} \boldsymbol{\alpha}_m + \hat{\mathbf{b}}_m); \end{cases} \\
& \begin{cases} \mathbf{C}_n = \tanh(\mathbf{W}_n \mathbf{C}_{mn}^\top + \mathbf{b}_n), \\ \boldsymbol{\alpha}_n = \text{softmax}(\mathbf{C}_n^\top \mathbf{u}_n), \\ \hat{\mathbf{n}} = \sigma(\hat{\mathbf{W}}_n \mathbf{C}_{mn}^\top \boldsymbol{\alpha}_n + \hat{\mathbf{b}}_n); \end{cases} \\
& \mathbf{u}_{mn} = \hat{\mathbf{m}} \oplus \hat{\mathbf{n}};
\end{aligned} \tag{3}$$

where $\mathbf{m}, \mathbf{n} \in \{t, a, v\}$ and $\mathbf{m} \neq \mathbf{n}$, σ represents an activation function, $\hat{\mathbf{m}}$ and $\hat{\mathbf{n}}$ refer to the reconstructed vectors of \mathbf{m} and \mathbf{n} , respectively, and the remaining variables stand for latent parameters learned from training. The final unified representation \mathbf{u} is defined as $\mathbf{u} = \mathbf{u}_{tv} \oplus \mathbf{u}_{ta} \oplus \mathbf{u}_{av}$.

3.3 Training and Optimization

The Adam algorithm (Kingma and Ba, 2015) is used as optimizer to minimize the binary cross entropy function. The binary cross entropy is defined as follows:

$$\mathcal{L} = -\frac{1}{N} \sum_{n=1}^N [y_n \log \hat{y}_n + (1 - y_n) \log(1 - \hat{y}_n)], \tag{4}$$

where N denotes the number of target values, y_n denotes the n^{th} ($1 \leq n \leq N$) true value, and \hat{y}_n denotes the corresponding predicted value. For the information retrieval task, given a query, we randomly generate a fake candidate along with the correct candidate from the candidate corpus for effective training.

4 Experimental Evaluation

Evaluation is conducted on two kinds of tasks, sentiment analysis and music information retrieval, as presented in Sections 4.1 and 4.2, respectively².

4.1 Sentiment Analysis

Task Description. Sentiment analysis is a task that deals with the computational detection and extraction of opinions, beliefs, and emotions in given content (Paltoglou and Thelwall, 2017). In this paper, we analyze posts from a social media platform ShutterSong. Given a post, which includes the visual (image), textual (song lyrics and user defined caption), and acoustic (song clip) parts, the goal is to tag a “positive” or “negative” label for the multi-modal content. We use “1” to indicate the positive label, and “0” to indicate the negative label. The sentiment analysis task can be regarded as a binary classification task.

Dataset and Parameter Settings. We sort a sub dataset from the ShutterSong dataset³ to conduct the sentiment analysis task. The user-defined mood tags are used as labels, and 3260 items have available moods. After removing duplications, we obtain 272 mood tags. We divide these mood tags into “positive” and “negative” on the basis of the related meaning manually. We obtain 2297 positive and 963 negative samples, with each including a user posted image, a song clip, and its corresponding lyrics, part of the samples have available captions. We call this dataset as MoodShutter, and separate it into train (80%), validation (10%), and test (10%) sets.

The song lyrics are truncated at 100 words, and captions are cleaned into five words according to a caption corpus. We use a 300 dimensional embedding matrix, which is pre-trained on the English

²The source code and more details on the experimental settings are available at <http://github.com/mengshor/JTAV>

³<https://drive.google.com/file/d/0B2N8XiDRrEgISXFJSXBEMWpUMDA/view>.

Wikipedia dataset⁴ through FastText, as W_e . The number of the hidden units of the BiGRU encoder is set as 50. Raw song clips are transformed into spectrograms with the *librosa* tool (McFee et al., 2017). We cut all the audio files into 10-second-segments. Audio sample rates are set as 22050Hz, hop lengths are 1024, and the numbers of frequency bins are 96. The parameter N of DCRNN is set as $N = 2$. The images are sized to $224 \times 224 \times 3$.

Baseline approaches. We compare our JTAV framework with:

- doc2vec (Le and Mikolov, 2014), a paragraph embedding approach;
- lyricsHAN (Alexandros, 2017), a hierarchical attention network which encodes songs into vectors;
- BiGRU, a bidirectional GRU network which utilizes the song lyrics as input;
- attBiGRU, our novel text modeling approach which adds the available captions as the supplying role based on BiGRU;
- MFCC (McFee et al., 2017), a type of cepstral representation of the audio clip generated by *librosa*;
- convnet (Choi et al., 2017a), an acoustic feature learning method based on transfer learning;
- DCRNN-MelS, our novel audio modeling approach which takes MelSs as input;
- DCRNN-CQT, our novel audio modeling approach which takes CQTs as input;
- F-VGG (Simonyan and Zisserman, 2015), a VGG network pre-trained on the ImageNet dataset and fine tuned on the MoodShutter dataset;
- F-ResNet (He et al., 2016), a ResNet pre-trained on the ImageNet dataset and fine tuned on the MoodShutter dataset;
- F-DenseNet (Huang et al., 2017), a DenseNet pre-trained on the ImageNet dataset and fine tuned on the MoodShutter dataset; and
- early fusion, a method which combines features learned by our methods before classification or regression tasks, and is very popularly used (Chen et al., 2017; Oramas et al., 2017).

Evaluation Metrics. The positive and negative samples are unbalanced and the ratio is about 7 : 3. The weighted average area under the curve (AUC) score, F1 score, and precision score are chosen as evaluation metrics⁵.

Results and Discussion. Table 1 presents the results between JTAV and the baseline approaches. For fair comparison, all feature vectors are obtained through optimized training; the classic logistic regression approach is utilized as the classifier; and the reported results are the average values of 10 runs. The values in boldface represent the best results among all the approaches. The proposed JTAV performed the best on all metrics, with 0.623 AUC score, 0.691 F1 score, 68.8% precision score.

Observations on single modality. For textual content, attBiGRU obtained the best results, and was superior to doc2vec with about 7 percentage points promotion of AUC score, 11 percentage points promotion of F1 score, and 6 percentage points promotion of precision score. These results demonstrate the proposed BiGRU approach extracts more powerful features than doc2vec and lyricsHAN. And taking the omni textual component into consideration, the attBiGRU approach can generate better results. For acoustic content, DCRNN outperforms MFCC and convnet because it can learn acoustic information both locally and globally. Long-term temporal context dependency is also retained. We also observe that CQT works better than MelS. DenseNet is the best choice for encoding images into vectors as compared with VGG and ResNet.

⁴https://en.wikipedia.org/wiki/Wikipedia:Database_download#English-language_Wikipedia.

⁵We utilize the implementation in sklearn, <http://scikit-learn.org/>.

Table 1: Results of the sentiment analysis task

modal	materials	approaches	AUC score	F1 score	precision score
text	lyrics	doc2vec (Le and Mikolov, 2014)	0.513	0.545	0.593
		lyricsHAN (Alexandros, 2017)	0.518	0.575	0.596
		BiGRU	0.572	0.602	0.640
	lyrics+caption	attBiGRU	0.581	0.652	0.650
audio	song clip	MFCC (McFee et al., 2017)	0.505	0.549	0.586
		convnet (Choi et al., 2017a)	0.518	0.614	0.621
		DCRNN-MelS	0.536	0.635	0.634
		DCRNN-CQT	0.559	0.651	0.643
image	image	F-VGG (Simonyan and Zisserman, 2015)	0.546	0.594	0.619
		F-ResNet (He et al., 2016)	0.578	0.618	0.645
		F-DenseNet (Huang et al., 2017)	0.588	0.627	0.653
text+audio	lyrics+caption+ song clip	early fusion	0.586	0.660	0.656
		CMF-AP	0.597	0.673	0.668
text+image	lyrics+caption+ image	early fusion	0.593	0.663	0.661
		CMF-AP	0.611	0.688	0.683
audio+image	image+song clip	early fusion	0.589	0.624	0.653
		CMF-AP	0.603	0.654	0.665
text+audio+image	lyrics+caption+ image+song clip	early fusion	0.602	0.671	0.669
		CMF-AP	0.623	0.691	0.688
JTAV=attBiGRU+(F-DenseNet)+(DCRNN-CQT)+(CMF-AP)					

Observations on multiple modalities. When more modalities emerge, more information is included, and better results are obtained. For example, the performance of combining textual and acoustic information is much better than that of using texts or audio solely. Our CMF-AP approach works steadily better than the baseline approach, because it not only extracts the inner features inside single modality, but also the the outer information cross modalities. Moreover, the best results are presented when we utilize all modalities and all available materials (i.e., JTAV). The AUC score is 0.623, representing an improvement of 2.1 percentage points to the early fusion approach. The F1 score is 0.691, performing an improvement of 2 percentage points over the best baseline approach. The precision score is 0.688, representing about 1.9 percentage points improvement to the best baseline approach. It seems like that the promotions of the proposed JTAV over the best baseline approach is not that huge. As a matter of fact, the early fusion approach utilizes the effective features generated by our proposed work (i.e., attBiGRU, DCRNN-CQT, and F-DenseNet). We believe that, if compared to the original version of early fusion method which only uses the previous features, the margin would be larger. For the sake of space limitation, in Table 1 the early fusion approach on previous features are omitted.

4.2 Music Information Retrieval

We conduct an extended experiment, which is can be explained as given an image query and named as image2song (Li et al., 2017), to verify the general effectiveness of JTAV. Image2song is a music information retrieval task and aims to find the most relevant song.

We perform the image2song task on two benchmark datasets, they are the Shattersong[†] dataset and the Shattersong[§] dataset (Li et al., 2017). Both datasets include no-repeat 620 songs with 3100 images, that is, each song is related to five images. Part of the images are attached with user-defined captions, and each song has an acoustic song clip and corresponding song lyrics. The difference lies in the partition strategy of the train and test dataset. In Shattersong[†], 100 songs and related images are selected randomly for testing, and the rest for training. While in Shattersong[§], the train and test set share the whole song




Figure 5: Results of image2song on Shattersong† Figure 6: Results of image2song on Shattersong§

set, while one of the five images of each song is chosen randomly for testing. The preprocessing settings of images, captions, song lyrics and song clips is the same as that of the sentiment analysis task.

We utilize the rank-based evaluation metrics to compare the results. For fair comparison, we adopt the evaluate metrics in (Li et al., 2017). Med r represents the medium rank of the ground truth retrieved song, and lower values indicate better performance. Recall@k (R@k for short), is the percentage of a ground truth song retrieved in the top-k ranked items, and higher values indicate better performance. In dataset†, $k = \{1, 5, 10\}$, and $1 \leq \text{Med r} \leq 100$; whereas in Shattersong§, $k = \{10, 50, 100\}$, and $1 \leq \text{Med r} \leq 620$.

We compare the proposed JTAV with the state-of-the-art approach (Li et al., 2017). For a more comprehensive comparison, we design several baseline approaches based on our proposed strategies: I+L uses images as queries and song lyrics as candidates with a suit of BiGRU, F-DenseNet and CMF-AP; I+C+L uses images and available captions as queries while song lyrics as candidates with a suit of attBiGRU, F-DenseNet and CMF-AP; I+S uses images as queries and song clips as candidates with a suit of DCRNN-CQT, F-DenseNet and CMF-AP; and I+C+S uses images and available captions as queries while song clips as candidates with a suit of DCRNN-CQT, F-DenseNet and CMF-AP. In JTAV, the images and available captions form the queries, the song clips and lyrics form the candidates.

Fig. 5 and Fig. 6 display the results on Shattersong† and Shattersong§, respectively. Both prove JTAV is the most powerful one among all testing approaches. JTAV turns the image2song task from finding the most similar item to finding the most relevant item, which is nearer the real situation in cross-modal retrieval tasks. Notably, JTAV exceeds the state-of-the-art approaches soundly. In Shattersong†, the Med r is no more than 5, and the R@1 score is more than 95%, which is more than 90% better than that obtained in (Li et al., 2017). In Shattersong§, the Med r is no more than 20, which is less a tenth of that obtained in (Li et al., 2017), the R@10 score is more than 83%, and the R@100 score is more than 97%.

5 Conclusion

In this paper, we aim to address the issue of learning social media content from the multi-modal view. We have designed effective approaches (i.e., attBiGRU) for textual content, F-DenseNet for visual content, and DCRNN for acoustic content, to extract fine-grained features. We have introduced CMF-AP to generate cross-modal aware matrices and reconstructed modal vectors, which are used to produce an unified representation of the multi-modal content. The proposed framework has been intensively evaluated, and the experimental results have demonstrated the general validity of JTAV by comparing with the state-of-the-art approaches.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant No.U1636116, 11431006, 61772288, the Research Fund for International Young Scientists under Grant No. 61650110510 and 61750110530, and the Ministry of education of Humanities and Social Science project under grant 16YJC790123.

References

- Eugene Agichtein, Carlos Castillo, Debora Donato, Aristides Gionis, and Gilad Mishne. 2008. Finding high-quality content in social media. In *Proceedings of the international conference on web search and data mining*, pages 183–194. ACM.
- Tsaptsinos Alexandros. 2017. Lyrics-based music genre classification using a hierarchical attention network. In *Proceedings of the 18th International Society of Music Information Retrieval Conference*, pages 694–701. ISMIR.
- Boualem Boashash. 1996. Time frequency signal analysis: Past, present and future trends. In *Control and Dynamic Systems*, volume 78, pages 1–69. Elsevier.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Judith C Brown and Miller S Puckette. 1992. An efficient algorithm for the calculation of a constant Q transform. *The Journal of the Acoustical Society of America*, 92(5):2698–2701.
- Xingyue Chen, Yunhong Wang, and Qingjie Liu. 2017. Visual and textual sentiment analysis using deep fusion convolutional neural networks. In *Proceedings of the 2017 IEEE International Conference on Image Processing*, pages 296–300. IEEE.
- Keunwoo Choi, George Fazekas, Mark Sandler, and Kyunghyun Cho. 2017a. Transfer learning for music classification and regression tasks. In *Proceedings of the International Society of Music Information Retrieval Conference*. ISMIR.
- Keunwoo Choi, György Fazekas, Kyunghyun Cho, and Mark B. Sandler. 2017b. A tutorial on deep learning for music information retrieval. *CoRR*, abs/1709.04396.
- Venkata Rama Kiran Garimella, Abdulrahman Alfayad, and Ingmar Weber. 2016. Social media image analysis for public health. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 5543–5547. ACM.
- Bo Han, Paul Cook, and Timothy Baldwin. 2013. Lexical normalization for social media text. *ACM Transactions on Intelligent Systems and Technology*, 4(1):5.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778. IEEE.
- Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4700–4708.
- Clayton J Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the 8th international AAAI conference on weblogs and social media*.
- Aren Jansen, Manoj Plakal, Ratheet Pandya, Dan Ellis, Shawn Hershey, Jiayang Liu, Channing Moore, and Rif A. Saurous. 2017. Towards learning semantic audio representations from unlabeled data. In *Proceedings of Workshop on Machine Learning for Audio Signal Processing at the 31st Conference on Neural Information Processing Systems*. NIPS.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3th International Conference on Learning Representations*. ICLR.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of International Conference on Machine Learning*, pages 1188–1196.
- Xuelong Li, Di Hu, and Xiaoqiang Lu. 2017. Image2song: Song retrieval via bridging image content and lyric words. In *Proceedings of IEEE International Conference on Computer Vision*, pages 5650–5659.
- Brian McFee, Matt McVicar, Oriol Nieto, Stefan Balke, Carl Thome, Dawen Liang, Eric Battenberg, Josh Moore, Rachel Bittner, Ryuichi Yamamoto, et al. 2017. librosa 0.5.0.
- Sergio Oramas, Oriol Nieto, Francesco Barbieri, and Xavier Serra. 2017. Multi-label music genre classification from audio, text, and images using deep features. In *Proceedings of the 18th International Society of Music Information Retrieval Conference*. ISMIR.

- Georgios Paltoglou and Mike Thelwall. 2017. Sensing social media: a range of approaches for sentiment analysis. In *Cyberemotions*, pages 97–117. Springer.
- Sharaj Panwar, Arun Das, Mehdi Roopaei, and Paul Rad. 2017. A deep learning approach for mapping music genres. In *Proceedings of the 12th System of Systems Engineering Conference*, pages 1–5. IEEE.
- Gwangbeen Park and Woobin Im. 2016. Image-text multi-modal representation learning by adversarial backpropagation. *CoRR*, abs/1612.08354.
- Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the 3th International Conference on Learning Representations*.
- Zhi Wang, Wenwu Zhu, Peng Cui, Lifeng Sun, and Shiqiang Yang. 2013. Social media recommendation. In *Social Media Retrieval*, pages 23–42. Springer.
- Yilin Wang, Suhang Wang, Jiliang Tang, Huan Liu, and Baoxin Li. 2015. Unsupervised sentiment analysis for social media images. In *Proceedings of the International Conference on Artificial Intelligence*, pages 2378–2379. AAAI Press.
- Chao Wu, Yaoxue Zhang, Jia Jia, and Wenwu Zhu. 2017. Mobile contextual recommender system for online social media. *IEEE Transactions on Mobile Computing*, 16(12):3403–3416.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J Smola, and Eduard H Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 14th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489. NAACL.
- Bilei Zhu, Fuzhang Wu, Ke Li, Yongjian Wu, Feiyue Huang, and Yunsheng Wu. 2017. Fusing transcription results from polyphonic and monophonic audio for singing melody transcription in polyphonic music. In *Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE.