Multi-modal Framework for Analyzing the Affect of a Group of People

Xiaohua Huang, Abhinav Dhall, Roland Goecke, *Member, IEEE*, Matti Pietikäinen, *Fellow, IEEE*, and Guoying Zhao*, *Senior Member, IEEE*

Abstract—With the advances in multimedia and the world wide web, users upload millions of images and videos everyone on social networking platforms on the Internet. From the perspective of automatic human behavior understanding, it is of interest to analyze and model the affects that are exhibited by groups of people who are participating in social events in these images. However, the analysis of the affect that is expressed by multiple people is challenging due to the varied indoor and outdoor settings. Recently, a few interesting works have investigated face-based Group-level Emotion Recognition (GER). In this paper, we propose a multi-modal framework for enhancing the affective analysis ability of GER in challenging environments. Specifically, for encoding a person's information in a group-level image, we first propose an information aggregation method for generating feature descriptions of face, upper body and scene. Later, we revisit localized multiple kernel learning for fusing face, upper body and scene information for GER against challenging environments. Intensive experiments are performed on two challenging group-level emotion databases (HAPPEI and GAFF) to investigate the roles of the face, upper body, scene information and the multi-modal framework. Experimental results demonstrate that the multi-modal framework achieves promising performance for GER.

Index Terms—Facial expression recognition, Group-level emotion recognition, Feature descriptor, Information aggregation, Multi-modality

1 INTRODUCTION

The presence of a large pool of data on multimedia enables us to explore images that contain more than one person. To advance the research in multimedia, it is of interest to understand the affect¹ that is exhibited by a group of people in an image, which is called group²-level emotion recognition (**GER**). The analysis of the emotion³ of multiple people in an image has various applications in multimedia such as image management and retrieval, photo

album creation [1] and event detection [2]. However, few research studies have examined automatic GER. This paper addresses the scenario in which a group of people in an image are posing during a social event. It considers social events such as a convocation, where many people gather together. To infer the perceived mood of the group of people in each image, we propose a multi-modal approach that based on face-level, body-level and scene-level cues.

As suggested by social psychology studies, group emotion can be conceptualized in different ways. Approaches for representing group emotion can be divided into two categories: bottom-up and top-down approaches [3], [4]. In the bottom-up methods, group emotion information is inferred from the subjects' attributes, while the top-down techniques consider external factors to the group members such as the effect of the scene. However, for group affective analysis, a single approach, such as a bottom-up or top-down approach, may lose useful information. One of the earliest works in GER involves fusing the top-down and bottom-up components of a group of people [5]. In [5], the top-down component contains the effect of the scene and group structure on the group, where the scene is a representation of the event and the surroundings of a group. An example of the effect of the scene is the scenario in which the same joke being told to a group of people who are sitting in a cafe and to a group of people in an official meeting room may result into different responses (affects). In contrast, the bottom-up component is the group members, together with their attributes, such as spontaneous expressions, clothes, age and gender. Furthermore, Dhall et al. [5] explored a topic-model-based approach, which takes into consideration the contributions of the members of the group and the effect of the neighbors. They demonstrated that combining the bottom-up and top-down components is very beneficial for GER.

Recently, several group expression models (**GEMs**) were proposed that combine bottom-up and top-down components for

- X. Huang and M. Pietikäinen are with the Center for Machine Vision and Signal Analysis, University of Oulu, Finland.
E-mail: xiaohua.huang@oulu.fi, matti.pietikainen@oulu.fi
- A. Dhall is with Department of Computer Science and Engineering, Indian Institute of Technology Ropar, India.
E-mail: abhinav@iitrpr.ac.in
- R. Goecke is with Human-Centred Technology Research Centre, University of Canberra, Australia.
E-mail: roland.goecke@ieee.org
- G. Zhao is with the Center for Machine Vision and Signal Analysis, University of Oulu, Finland and with the School of Information and Technology, Northwest University, 710069, China.
E-mail: guoying.zhao@oulu.fi

The asterisk indicates corresponding author.

This work was supported by Jorma Ollila Grant of Nokia Foundation, Central Fund of Finnish Cultural Foundation, AI Grant of Kaute Foundation, Academy of Finland, Tekes Fidipro Program (Grant No. 1849/31/2015) and Tekes project (Grant No. 3116/31/2017), Infotech, National Natural Science Foundation of China (No. 61772419). We gratefully acknowledge Dr. Xiuyan Shao from University of Oulu for performing the statistical analysis. We also acknowledge the support of NVIDIA Corporation with the donation of Titan Xp and Quadro P6000 GPUs used for this research. The authors wish to acknowledge CSC IT Center for Science, Finland, for generous computational resources.

1. In affective computing, affect is a concept that describes the experience of feeling or emotion. However, in this paper, affect is specified as the collective mood of a group.

2. In this paper, a group is defined as two or more persons interacting during a social event. Interaction is interpreted that the acts actions or practices of two or more people mutually oriented towards each other's selves.

3. In this paper, emotion represents the basic facial expressions or happiness intensity. We use emotion and mood interchangeably.

face-level GER [6], [7]. The GEMs are inspired by works in social psychology in which group emotions are conceptualized as pairs of bottom-up and top-down approaches [3], [4]. In [6], [7], global attributes such as the effect of neighboring group members were considered as the top-down component and local attributes such as an individual's features were considered as the bottom-up component. However, existing GEMs on GER may suffer from failure of face processing, when the face processing pipeline is noisy due to challenging conditions such as illumination variation and head pose changes.

As we learn from the user survey in [6], the face may not be the only attribute that is useful for understanding the group-level affect. Additionally, psychology studies [3], [4] suggested that combining bottom-up and top-down components can well represent the group emotion in social events. Moreover, research on facial expression recognition demonstrates that multiple sources increase the robustness of the emotion recognition system [8], [9], [10]. It is necessary to exploit a multi-modal approach to make the GER robust 'in the wild', which refers to attributes such as different backgrounds, head motion, illumination conditions, occlusion and the presence of multiple people.

Overall, this work makes five main contributions: (1) Information aggregation is proposed, in which an improved Fisher Vector is used to encode feature vectors from multiple persons in a group of images. (2) Three modalities, namely, face, upper body and scene, are explored and discussed for the task of GER; (3) A super-pixel based approach is explored for analyzing the scene in GER; (4) A robust multi-modal approach is exploited to infer the emotional state of a group of people in an image; (5) Comprehensive experiments on two 'in the wild' databases demonstrate the superiority of information aggregation and the multi-modal framework over previous methods.

To explain the concepts that are used in our approach, the paper is organized as follows: In Section 2, we present the recent work on multi-modal approaches for GER, feature extraction from face, body and scene, and feature fusion. In Section 3, we discuss face, upper body and scene feature descriptor extraction. In Section 4, we encode face, upper body and scene information for representing group-level images. In Section 5, we present our multi-modal framework. In Section 6, we present the results for face, upper body and scene and the multi-modal approach. Finally, we present our conclusions in Section 7.

2 RELATED WORK

Currently, there are two multi-modal approaches that combine multiple sources for GER 'in the wild' [11], [12]. In [11], the facial action unit and facial features are regarded as the bottom-up component, while the scene feature is viewed as the top-down component. Specifically, they used a bag-of-words method to encode the facial action unit and facial features in a group-level image. Then, they employed two scene descriptors to extract scene features. The combination of face and scene achieved promising performance.

Researchers have reported that body expressions are of significant importance when displaying and understanding expressions [13], [14], [15]. For example, Joshi et al. [13] used the key points of the full body for depression analysis. Mou et al. [12] implemented body information for GER. Another interesting multi-modal work [12] combined face and body information to predict the valence and arousal of a group of people. Both works [11],

[12] demonstrate that a multi-modal framework can perform better than a uni-modal approach. In addition, multiple-source fusions are demonstrated to increase the robustness of the system [8], [9], [10]. Therefore, in this paper, we use multiple sources for GER. Both works [11], [12] have limitations. For example, the body information is neglected in the model of [11]; Mou et al. [12] ignored scene information and experimented on specific groups based on the fixed number of faces and bodies. According to the user survey in [11], 50% of the participants mentioned that the scene affects their labeling of the emotions of a group of people. Recent studies [16], [17], [18] have shown that the perception of emotion of a subject is influenced by the body expression of the subject and by its surrounding scene. Motivated by these studies, we propose a new multi-modal approach that considers the face, upper body and scene for strengthening the recognition ability of GER. The multimodal framework that is proposed in this paper is depicted in Figure 1. Different from [12], the upper body region is defined as the shoulder information, excluding the face image. We define the scene as the background, excluding the upper body and face regions, which is different from [11]. The main advantage of the new definition is that it is convenient to observe the contribution of each factor. From a computer vision perspective, it is intuitive to take the scene-level descriptor and body expression features into consideration. Similar to [6], [7], [11], [12], feature extraction and data fusion are two important issues that need to be resolved for our multi-modal approach.

2.1 Feature extraction

Feature extraction from an image that contains multiple people is an open problem in GER. In other words, we aim at using a single feature vector to represent multiple persons' information in a group-level image. In the GEMs for GER [6], [7], the authors first extracted features from the detected faces of group-level images. Then, they constructed GEM for group-level images according to these facial features. However, they did not extract the features from a group-level image. Bag-of-Words is used by Dhall et al. [11] to accumulate a histogram from multiple faces for a group-level image. However, it is observed that the obtained feature is very sparse and not stable due to mis-alignment problems that are caused by head pose changes. Instead, the Fisher Vector (**FV**) representation of images can be viewed as an extension of the Bag-of-Words method. It has been widely used in computer vision problems such as depression analysis [19] and emotion and membership recognition in group videos [20]. This representation has the advantage of providing better classification performance than a Bag-of-Words that is obtained with supervised visual vocabularies. Recently, an interesting improved Fisher Vector (**IFV**) method was proposed by [21] for improving FV by using L2 and power normalization strategies. It is used to describe the facial features by encoding the scale-invariant feature transform (**SIFT**) feature of pixels in an image [22]. An IFV is also applied for image retrieval by constructing an image representation from a set of a few hundred windows [23]. Motivated by [23], we propose information aggregation based on an improved FV for encoding regional features, which makes the features more compact for GER than the Bag-of-Words approach [11]. Moreover, our proposed information aggregation method can generate three feature vectors for the face, upper body and scene in a group-level image, thereby leading to the use of classical feature fusion approaches. In the following, we will discuss the available regional features for the face, body and scene.

Fig. 1: Pictorial illustration of the proposed multi-modal framework for group-level emotion recognition, where “RVLBP”, “PHOG”, “LBP” and “SIFT” represent Riesz-based volumed local binary pattern, pyramid histogram of oriented gradients, local binary pattern and scale-invariant feature transform, respectively. In this framework, we exploit face, upper body and scene channels to strengthen the recognition ability of group-level emotion recognition. For each channel, we propose an “information aggregation” component for generating a compact feature for a group-level image. “Information aggregation” includes two stages: visual vocabulary generation and feature encoding. In the final layer of the multi-modal framework, we propose revisited localized multiple kernel learning for fusing three features. In this picture, we also illustrate the procedure of the “feature extraction on superpixel (SIFT)” component for the scene-level channel.

Face: Facial expression descriptors are broadly categorized as geometric- and appearance-based. Comparing with geometric-based features, appearance-based features have the advantage that they are more stable to global changes that are caused by inaccurate alignment and illumination changes. Recently, Zhang et al. combined the Gabor and local binary pattern (LBP) descriptors to form local Gabor binary patterns (LGBP), to improve the face recognition performance [24], because Gabor wavelets are used to capture the local structure that corresponds to the spatial frequency, spatial localization and orientation selective [25] and LBP [26] is also robust to global changes. However, LGBP suffers from two critical problems: (1) the patterns are not optimal if the broad spectral information with maximal spatial localization needs to be obtained and (2) their maximum bandwidth is restricted to approximately one octave. Recent studies have demonstrated the local image information can be well characterized in a unified theoretical framework, namely, the Riesz transform. Felsberg and Sommer [27] proposed using the first-order Riesz transform for image processing. The Riesz transform has attracted substantial interest from researchers in the fields of texture classification [28] and face analysis [29]. However, these works designed the texture descriptors based on the first-order Riesz transform and lost important complex structures, such as corners. According to the intrinsic dimension theorem [30], the first-order Riesz transform is designed for an intrinsic 1D signal and the higher-order Riesz transforms reused to characterize the intrinsic 2D local structures. To describe the intrinsic 2D local structures, high-order Riesz

transforms have been developed for biometric recognition [29] and texture recognition [31]. Recently, we proposed a texture descriptor that is based on higher-order Riesz transforms and LBP for characterizing facial expression [7]. It has been demonstrated to well characterize facial expressions and achieve more promising performance than the state-of-the-art texture descriptors in the tasks of facial expression recognition, smile detection and happiness intensity estimation in the wild.

Body: As a dense version of the dominating scale-invariant feature transform (SIFT) feature, pyramid histogram of oriented gradients (PHOG) [32] has shown great success in human upper body estimation [33] and human detection [34]. It has been widely accepted as one of the best features for capturing edge or local shape information. In addition, the LBP operator [26] is an exceptional texture descriptor. The LBP operator is highly discriminative and its key advantages, namely, its invariance to monotonic gray-level changes and computational efficiency, make it suitable for image analysis tasks such as pose estimation. The combination of PHOG and LBP descriptors has been demonstrated to robustly describe body information against a challenging background [35]. PHOG performs poorly when the background is clustered with noisy edges. LBP is complementary in this aspect. It can filter out noises using the concept of a uniform pattern. The appearance of the human upper body can be better described if edge/local shape information and texture information are combined.

Scene: Two widely used scene analysis descriptors, namely, census transform histogram (CENTRIST) [36] and GIST [37],

have been employed by [11] to analyze the emotion of a group of people. They model the scene at the holistic level. As the scene includes many objects, multiple people and a complicated background in a group-level image, the holistic level will destroy semantic information of the scene. Instead, we aim at extracting the features of the scene at the local level, which will be helpful for GER. Furthermore, due to the semantic information of scene, it is not reasonable to divide the scene into several fixed blocks. Alternatively, a superpixel [38] method can be used to resolve the above-mentioned problems, and enables us to explore the scene feature in a semantically meaningful subregion.

2.2 Data fusion

Given face, upper body and scene features for a group-level image, we aim at combining them in a way that increases their discriminatory power. Generally, concatenating all features together is a simple way to perform feature fusion. However, it is not guaranteed that the complementary information will be captured. It is possible that increasing the feature dimension will reduce the efficiency of computation. Alternatively, there are several ways to efficiently combine multiple features in different fields using multi-view learning methods [39]. According to [39], multi-view learning methods can be categorized into three groups: co-training, subspace learning and multiple kernel learning. Researchers developed the variant subspace learning method for image classification, with the aim of obtaining a latent subspace that is shared by multiple views. For example, multi-set statistical uncorrelated projection analysis and multi-set discriminating uncorrelated projection analysis were proposed by Wu et al. [40] for exploring discriminant features from three color components for face recognition. Luo et al. [41] proposed a large-margin multi-modal multi-task feature extraction framework for image classification, in which they simultaneously utilize the information that is shared between tasks and the complementarity of different modalities to extract strong predictive features. Wu et al. [42] developed a multi-view low-rank dictionary learning approach by introducing a multi-view dictionary low-rank regularization term and designing a structural incoherence constraint for multi-view dictionary learning when dealing with face recognition, object classification and digit classification tasks. It is found that these methods extract discriminative information based on the label information. However, the task of GER is to estimate the happiness intensity of a group of people. Without label information, subspace learning may not obtain the satisfactory performance for GER. Recently, multiple kernel learning (MKL) has been used for image classification [43] and audio-video emotion recognition [44]. It aims at exploiting kernels that naturally correspond to different views and combining kernels to improve learning performance. Gönen et al. [45], [46] developed the localized multiple kernel learning (LMKL) framework by fusing different kernel functions on the same input data for classification and regression problems. Interestingly, the LMKL algorithm utilized a gating model to select the appropriate kernel function locally, while also coupling the localizing gating model and the kernel-based classifier in a joint manner. Han et al. [43] developed a novel sample-wise alternating optimization algorithm for training LMKL, namely, S-LMKL, which achieved consistent accuracy improvement for classification. For GER, since there are two tasks in this paper, namely, estimating the happiness intensity and classifying emotion, we aim at using an objective function that can be shared between the two

tasks. As LMKL shared a similar objective function for regression and classification tasks, we only focus on the original LMKL [45], [46] for achieving two different tasks in GER. It is observed that the gating model in LMKL is defined up to a set of parameters that are learned from the same input data. However, for feature fusion, the gating model needs to consider the locality of different modalities. Therefore, we will revisit LMKL for multi-modal GER in this paper.

3 INFORMATION EXTRACTION

Analysis of the affect that is expressed by multiple people is challenging due to difficult situations such as head and body pose variations. A group-level image may contain face, upper body and scene information, which contribute to emotion perception of the affect that is expressed by multiple people. Recently, multi-modal emotion recognition has gained ground [8]. Thus, we explore face, upper body and scene features for GER. To facilitate understanding, we define the detected object, such as a face, as a sub-image. Considering local information, we segment the sub-image into numerous local regions. For each local region, we use a feature descriptor to extract its information. In this section, we will introduce information extraction for local regions from the face, upper body and scene.

3.1 Facial features

For extracting facial features, we employ Riesz-based Volume Local Binary Pattern (RVLBP) [7]. Specifically, a facial image is divided into $P \times Q$ overlapping blocks. Then, RVLBP is used for each block. In this paper, we choose the commonly utilized log-Gabor filter [47] with 5 scales and 8 orientations. The first-order and second-order Riesz transform functions in the case of a 2D image are expressed as:

$$h_x(\mathbf{x}) = g * \frac{x}{2\pi|\mathbf{x}|^3}, \quad h_y(\mathbf{x}) = g * \frac{y}{2\pi|\mathbf{x}|^3}, \quad (1)$$

$$h_{xx}(\mathbf{x}) = h_x(\mathbf{x}) * h_x(\mathbf{x}), \quad (2)$$

$$h_{yy}(\mathbf{x}) = h_y(\mathbf{x}) * h_y(\mathbf{x}), \quad (3)$$

$$h_{xy}(\mathbf{x}) = h_x(\mathbf{x}) * h_y(\mathbf{x}), \quad (4)$$

where g is the log-Gabor filter with the u -th scale and v -th orientation. We convolve an image $\mathbf{I}$ with the first-order and second-order Riesz transforms and obtain the Riesz faces as $\mathbf{R}_x = \mathbf{I} * h_x(\mathbf{x})$, $\mathbf{R}_y = \mathbf{I} * h_y(\mathbf{x})$, $\mathbf{R}_{xx} = \mathbf{I} * h_{xx}(\mathbf{x})$, $\mathbf{R}_{yy} = \mathbf{I} * h_{yy}(\mathbf{x})$ and $\mathbf{R}_{xy} = \mathbf{I} * h_{xy}(\mathbf{x})$. For convenience, superscripts x , y , xx , yy and xy of R are omitted when we discuss the subsequent feature extraction.

We view volume-based Riesz faces $\mathbf{R}$ as a video sequence. In [48], the authors considered the video sequence as three orthogonal planes and concatenated feature in these three direction. They efficiently obtained the LBP features from three orthogonal planes that represented the appearance and motion information. Additionally, their methods are computationally simple. Motivated by them, we employ the LBP on three orthogonal planes: the XY, XZ and YZ planes of $\mathbf{R}$. Finally, we combine the histograms of these planes to represents faces.

3.2 Upper body features

Due to the low resolution, size and variations in pose and lighting, in naturalistic settings, a face may provide less reliable evidence for analyzing affect. In this scenario, we extract body features

as an additional source of information for group-level emotion recognition. For each person in the image, we use the upper body detection [49] to obtain an upper-body region and omit the face region from the upper body region since the face may provide redundant information to the upper body. One example is illustrated in Figure 1. The upper body contains shoulder information, which is similar to [12]. The presence of varied backgrounds, illumination change and partial occlusion in challenging conditions may make the GER even more difficult. It is necessary to perform feature augmentation by computing low-level features on the upper body region. In our method, we divide the detected upper body image into $P \times Q$ overlapped blocks with an overlap ratio of 0.7. For each block, PHOG is computed by applying edge detection on a sub-region, followed by histogram computation in a pyramid fashion. Local binary patterns are computed over each pixel, which is neighboring to eight pixels. The combination of these two features is robust to scale and illumination changes.

3.3 Scene-level information

Scene-level information has been investigated in scene image classification [37], [50]. However, there is scant research on using scene-level information for GER. The background may provide complementary information for GER. According to the user survey in [11], attributes such as a pleasant scene (background/situation) and clothing affect the perception of human beings regarding the affect of a group of people in an image. For convenience, we define the scene information in the group-level image as the region that excludes the face and upper body. Therefore, we exploit the usefulness of scene analysis features as the global information for the multi-modal framework.

According to our empirically experimental result, we segment each group-level image into N_s superpixels by using the Linear Spectral Clustering Superpixels segmentation algorithm in [38], in which one superpixel represents the segmented region. A superpixel is roughly uniform in color and naturally preserves the boundaries of objects. Two exemplars are shown in Figure 2. The main advantage is that the superpixel algorithm can explore many semantically meaningful subregions, such as hand posture and clothing. However, it is important to choose a suitable number of superpixels for scene information, since using many superpixels may destroy some objective structures. For example, in Figure 2(a), the hand posture can be better segmented when we use 200 superpixels. To encode appearance information into segmented regions, we describe each segmented region by using a SIFT descriptor. SIFT has been widely used as a local descriptor to characterize local gradient information [51] in object recognition [52]. SIFT has been accepted as one of the best features in terms of resistance to common image deformations. In [51], a SIFT descriptor is a sparse feature representation that consists of both feature extraction and detection. In this paper, however, we only use the feature extraction component. Given a segmented region, we divide the neighborhood (16×16) of every pixel into a 4×4 cell array, quantize the orientation into 8 bins in each cell, and obtain a 128-dimensional vector as the SIFT representation for the pixel. The procedure is shown in the green rectangle of Figure 1. The feature of a segmented region can be computed as the average of these per-pixel SIFT descriptors as $\frac{1}{M} \sum_{i=1}^M \mathbf{f}_i$, where M is the number of pixels in the segmented region and $\mathbf{f}_i$ represents the SIFT feature of the i -th pixel.

4 INFORMATION AGGREGATION

As previously presented in Section 3, for one region/block, we obtain features $\mathbf{h}_{\text{face}}$, $\mathbf{h}_{\text{body}}$ and $\mathbf{h}_{\text{scene}}$ for the face, upper body and scene, respectively. Generally, concatenating all features is a simple way of representing a group-level image. However, there are different numbers of persons in different group-level images; for example, in Figure 3, there are three and twelve subjects in the two group-level images. The varied concatenated feature dimension renders it difficult to measure two group-level images. Moreover, it is not conducive to classification. Therefore, we propose a new method for resolving the above-mentioned problem in group-level expression recognition. For convenience, we first discuss the problem formulation for the face and omit the superscript of $\mathbf{h}$.

The proposed framework is illustrated in the ‘‘Information aggregation’’ component of Figure 1. After feature extraction, there are multiple feature vectors for a single group-level image. For two group-level images, we obtain two different numbers of feature vectors, resulting in difficulty measuring the similarity of the two feature sets. For GER, aggregating $\mathbf{h}_1, \dots, \mathbf{h}_N$ into a single feature vector $\mathbf{x}$ is one problem that needs to be resolved, where N is the number of subjects in a group-level image. The problem is defined as follows:

$$\mathbf{x} = f(\mathbf{h}_1, \dots, \mathbf{h}_N), \quad (5)$$

where $f(\cdot)$ denotes the aggregation function.

As is well known, a conventional way for the aggregation function $f(\cdot)$ to encode all persons’ features is by using Bag-of-Words (**BoW**) method in [11] and the Fisher Vector strategy in [19], [20]. However, the obtained feature is very sparse. Instead, we propose an INformation Aggregation (**INFA**) method that is based on the improved Fisher Vector [53] for $f(\cdot)$ for encoding region-based features of multiple persons into a compact feature for a group-level image.

As mentioned in Section 3, we divide each face into $P \times Q$ blocks, where P and Q are the numbers of rows and columns of blocks, respectively. This leads to $P \times Q$ regional features for the i -th subject $\mathbf{h}_i$. Based on $P \times Q$ blocks, we obtain $P \times Q \times N$ regional features for one group-level image, where N is the number of subjects in the group-level image. We denote them as $\{\mathbf{g}_j\}_{j=1}^M$ where $M = P \times Q \times N$. Considering local information, Equation 5 is modified as:

$$\mathbf{x} = f(\mathbf{g}_1, \dots, \mathbf{g}_M). \quad (6)$$

To encode $\{\mathbf{g}_1, \dots, \mathbf{g}_M\}$, we exploit the implementation of visual vocabulary. Given all regional features, we train a visual vocabulary by using a GMM with diagonal covariances under word size K , where the word size is defined as the number of Gaussians, and consider derivatives with respect to the Gaussian mean and variance. We obtain the visual vocabulary as $\Delta_k = \{\omega_k, \mu_k, \sigma_k\}$, where $k = 1, \dots, K$, and ω_k, μ_k and σ_k are the mixture weight, mean, and diagonal covariance of the GMM, respectively. The procedure is shown in stage 1 of the ‘‘Information aggregation’’ component of Figure 1. Before obtaining GMM, we apply Principal Component Analysis (**PCA**) on the local regional features to decorrelate the features over all regions. Based on Δ_k , this leads to a representation that captures the average first- and second-order differences between the regional features and each of the GMM centers:

$$\Phi_k^{(1)} = \frac{1}{M\sqrt{\omega_k}} \sum_{j=1}^M \gamma_j^k \left(\frac{\mathbf{g}_j - \mu_k}{\sigma_k} \right), \quad (7)$$

Fig. 2: Superpixel segmentation results on two group-level images, where we implemented 100, 200 and 500 superpixels, respectively.

Fig. 3: Proposed framework for resolving the subject imbalance problem in group-level emotion recognition, where the number under each face image represents the order index of the image that is obtained by face detection.

$$\Phi_k^{(2)} = \frac{1}{M\sqrt{2w_k}} \sum_{j=1}^M \gamma_j^k \left(\frac{(g_j - \mu_k)^2}{\sigma_k^2} - 1 \right), \quad (8)$$

where γ_j^k is the soft assignment weight of g_j to the k -th word size. For a group-level image, its feature $\mathbf{x}$ is obtained by stacking the differences: $\mathbf{x} = [\Phi_1^{(1)}, \Phi_1^{(2)}, \dots, \Phi_K^{(1)}, \Phi_K^{(2)}]$. Therefore, it is $2KD$ -dimensional, where D and K are the reduced dimensionality of PCA and the word size, respectively. Furthermore, we use power normalization for $\mathbf{x}$, which is defined as $\mathbf{x} = \text{sign}(\mathbf{x})|\mathbf{x}|$. Finally, $\mathbf{x}$ is normalized by the L2 norm.

Last, information aggregation is used to encode all blocks of multiple faces into one feature vector. The upper body has been handled in the same way that we divided it into fixed blocks. The problem for the upper body is formulated in the same way. For the scene, we used a superpixel segmentation algorithm for extracting N_s semantic regions for a group-level image. Therefore, different from the face/upper body, we define M in Equation 6 as N_s , where N_s is the number of superpixels.

5 REVISITED LOCALIZED MULTIPLE KERNEL LEARNING FOR THE MULTIMODAL FRAMEWORK

Given a group-level image, the face, upper body and scene features that are obtained by using information aggregation are denoted as $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$, respectively. Its corresponding label is represented

as y . We aim at combining the face, upper body and scene features to (1) estimate the happiness intensity of the group-level image on the HAPPEI database, which is viewed as a regression problem, and (2) classify the group-level image into one of three emotion categories (positive, neutral or negative) on the GAFF database.

Fig. 4: Structures of (a) LMKR and (b) RLMKR.

Recently, Gönen et al. [45], [46] developed the localized multiple kernel learning (**LMKL**) framework for classification and regression problems. Interestingly, the LMKL algorithm [45], [46] utilized a gating model to select the appropriate kernel function locally and coupled the localizing gating model and the kernel-based classifier in a joint manner. LMKL was used to [45], [46] investigate the gating model for fusing multiple kernels based on the same input data, as shown in Figure 4(a). The gating model in LMKL is defined up to a set of parameters which are learned from the same input data. In contrast, the gating model in feature fusion (in Figure 4(b)) should consider the locality of different modalities. Therefore, we will revisit LMKL (**RLMKR**) for the multi-modal framework. Motivated by LMKL, for fusing the face, upper body and scene for group-level affective state analysis, our problem is formulated by considering the gating function as:

$$f(\mathbf{x}) = \sum_{i=1}^3 \eta(\mathbf{x}_i) \langle \beta_i, g(\mathbf{x}_i) \rangle + b, \quad (9)$$

where β_i is the weight for $\mathbf{x}_i$, and $\eta(\mathbf{x}_i)$ and $g(\mathbf{x}_i)$ are the gating and mapping functions for $\mathbf{x}_i$, respectively.

Given n group-level images, their feature sets that are obtained by using information aggregation are denoted as $\mathbf{X}_1$, $\mathbf{X}_2$ and $\mathbf{X}_3$ for the face, upper body and scene, respectively. For efficient computation, PCA whitening is applied to each modality. Then the objective function in Equation 9 on the whitened preprocessed data is reformulated as follows:

$$f(\mathbf{x}) = \sum_{i=1}^3 \eta(\mathbf{x}_i) \sum_{j=1}^n \alpha_j y_j \mathbf{K}(\mathbf{U}_i^T \mathbf{x}_i, \mathbf{U}_i^T \mathbf{x}_i^j) + b, \quad (10)$$

where η is a gating function, $\mathbf{x}_i^j$ is the j -th data sample of $\mathbf{X}_i$, $\mathbf{U}_i$ is the whitened matrix of $\mathbf{X}_i$ by using PCA, and $\mathbf{K}$ is the kernel function. The same kernel function is employed for the three modalities. For the kernel function, we investigate the linear kernel, the Gaussian kernel and the Histogram Intersection (**HI**) kernel.

By modifying the original SVM formulation with Equation 10, we obtain the following optimization problem for (1) classification:

$$\begin{aligned} \min & \frac{1}{2} \sum_{i=1}^3 \|\mathbf{w}_i\|^2 + C \sum_{j=1}^n \xi_j \\ \text{w.r.t. } & \mathbf{w}_i, b, \xi, \eta \\ \text{s.t. } & y_j \left(\sum_{i=1}^3 \eta(\mathbf{x}_i^j) \langle \mathbf{w}_i, g(\mathbf{x}_i^j) \rangle + b \right) \geq 1 - \xi_j, \\ & \xi_j \geq 0, \forall j, \end{aligned} \quad (11)$$

and (2) regression:

$$\begin{aligned} \min & \frac{1}{2} \sum_{i=1}^3 \|\mathbf{w}_i\|^2 + C \sum_{j=1}^n (\xi_j^+ + \xi_j^-) \\ \text{w.r.t. } & \mathbf{w}_i, b, \xi_j^+, \xi_j^-, \eta \\ \text{s.t. } & \epsilon + \xi_j^+ \geq y_j, \epsilon + \xi_j^- \leq y_j, \\ & \xi_j^+ \geq 0, \xi_j^- \geq 0, \forall j, \end{aligned} \quad (12)$$

where C is the regularization parameter, $\{\xi, \xi^+, \xi^-\}$ are slack variables, and ϵ is the tube width. Following [45], the optimization problem of classification (Equation 11) and regression

(Equation 12) is resolved by a two-step alternating optimization algorithm.

(1) For the classification problem, the first step is to resolve Equation 11 with respect to $\mathbf{w}_i, b$, and ξ while fixing η_i , while the second step is to update the parameters of η_i using a gradient-descent step calculated from the objective function (Equation 11). Based on the two-step alternating optimization algorithm, the dual formulation is obtained as:

$$\begin{aligned} \max J &= \sum_{j=1}^n \alpha_j - \frac{1}{2} \sum_{j=1}^n \sum_{k=1}^n \alpha_j \alpha_k y_j y_k \mathbf{K}_\eta, \\ \text{w.r.t. } & \alpha \\ \text{s.t. } & \sum_{j=1}^n \alpha_j y_j = 0 \text{ and } C \geq \alpha_j \geq 0, \forall j \end{aligned} \quad (13)$$

where $\mathbf{K}_\eta = \sum_{i=1}^3 \eta(\mathbf{x}_i^j) \mathbf{K}(\mathbf{U}_i^T \mathbf{x}_i^j, \mathbf{U}_i^T \mathbf{x}_i^k) \eta(\mathbf{x}_i^k)$, which is the locally combined kernel matrix.

(2) For the regression problem, following the step of the classification problem, the dual formulation is easily obtained as:

$$\begin{aligned} \max J &= \sum_{j=1}^n y_j (\alpha_j^+ - \alpha_j^-) + \epsilon \sum_{j=1}^n y_j (\alpha_j^+ + \alpha_j^-) \\ & - \frac{1}{2} \sum_{j=1}^n \sum_{k=1}^n (\alpha_j^+ - \alpha_j^-) (\alpha_k^+ - \alpha_k^-) \mathbf{K}_\eta, \\ \text{w.r.t. } & \alpha_j^+, \alpha_j^- \\ \text{s.t. } & \sum_{j=1}^L (\alpha_j^+ - \alpha_j^-) = 0 \\ & C \geq \alpha_j^+ \geq 0, C \geq \alpha_j^- \geq 0, \forall j \end{aligned} \quad (14)$$

For the gating model η , we implement the Softmax function, which can be expressed as:

$$\eta(\mathbf{x}_i^j) = \frac{\exp(\langle v_i, \mathbf{x}_i^j \rangle + v_{i0})}{\sum_{k=1}^3 \exp(\langle v_k, \mathbf{x}_i^j \rangle + v_{k0})}, \quad (15)$$

where v_i and v_{i0} are the parameters of this gating model and the Softmax guarantees non-negativity. We can simply use the objective function of Equation 13 or Equation 14 as function $J(\eta)$ to calculate the gradients of the primal objective with respect to the parameters of η_i . To train the gating model, we take derivatives of $J(\eta)$ with respect to v_i and v_{i0} and use gradient-descent:

$$\begin{aligned} \frac{\partial J(\eta)}{\partial v_{i0}} &= -\frac{1}{2} \sum_{j=1}^n \sum_{k=1}^n \sum_{m=1}^3 \alpha_j \alpha_k y_j y_k \eta(\mathbf{x}_m^j) \mathbf{K}(\mathbf{U}_m^T \mathbf{x}_m^j, \mathbf{U}_m^T \mathbf{x}_m^k) \\ & \quad \eta(\mathbf{x}_m^k) (\delta_i^m - \eta(\mathbf{x}_i^j) + \delta_i^m - \eta(\mathbf{x}_i^k)), \end{aligned} \quad (16)$$

and

$$\begin{aligned} \frac{\partial J(\eta)}{\partial v_i} &= -\frac{1}{2} \sum_{j=1}^n \sum_{k=1}^n \sum_{m=1}^3 \alpha_j \alpha_k y_j y_k \eta(\mathbf{x}_m^j) \mathbf{K}(\mathbf{U}_m^T \mathbf{x}_m^j, \mathbf{U}_m^T \mathbf{x}_m^k) \\ & \quad \eta(\mathbf{x}_m^k) (\mathbf{x}_j [\delta_i^m - \eta(\mathbf{x}_i^j)] + \mathbf{x}_k [\delta_i^m - \eta(\mathbf{x}_i^k)]), \end{aligned} \quad (17)$$

where δ_i^m is 1 if $i = m$ and 0 otherwise. After updating the parameters of η , we can solve a single kernel SVM with $\mathbf{K}_\eta$ at each step.

6 EXPERIMENTS

In the previous section, we discussed two novel ideas for GER: INformation Aggregation (INFA) for encoding information of multiple persons and a multi-modal framework. To the best of our knowledge, there are only a few group-level emotion recognition databases. Two of them are available publicly. Therefore, we focus on two ‘in the wild’ databases: the HAPpy PEople Images (HAPPEI) and Group AFFect (GAFF) databases. In this section, we will thoroughly evaluate all parameters of INFA and the multi-modal framework on the HAPPEI database [6]. Following [7], the four-fold cross-validation method is used, where 1,500 images are used for training and 500 for testing, and the process is repeated 4 times. The main task is to predict the intensity levels of group-level images. Thus, mean absolute error (MAE) is used as metric. Finally, we evaluate the generalization ability of our proposed method on the GAFF database [11] using the well-designed parameters of the HAPPEI database.

6.1 Database descriptions

HAPPEI database: This database was collected by Dhall et al. [6] and contains 2,886 images. Six exemplar images are shown in Figure 5. All images were annotated with a group-level mood intensity. Moreover, 2,886 images were manually annotated for happiness intensity by four human labelers, who annotated different images. The mood was represented by the happiness intensity, which corresponds to six stages of happiness (0-5): Neutral, Small smile, Large smile, Small laugh, Large laugh and Thrilled. The aim of this database in [6] is to infer the perceived group mood as closely as possible to human observers. An interesting application of this database is the estimation of the happiness intensity of group-level images.

Fig. 5: Images that are labeled with six happiness intensities in the HAPPEI database, where the number at the bottom of an image indicates the intensity of happiness of the group-level image.

GAFF database: Dhall et al. [11] extended the HAPPEI database from positive affect only [6] to a wider variety of emotions (Positive, Neutral and Negative) of a group of people. They first developed a user study to understand attributes that

affect the perception of the affective state of a group. Then, they acquired the GAFF database [11] by searching group-level images from Flickr and Google according to keywords. All images are labeled with three emotion categories (Positive, Neutral and Negative). Examples are shown in Figure 6. In the GAFF database, the main task is to classify each group-level image into an emotion category.

The databases and experimental protocols that are used in the experiments are summarized in Table 1.

Fig. 6: Six examples with three emotion labels on the GAFF database: (a) positive, (b) neutral and (c) negative.

TABLE 1: Description of two emotion recognition databases. In this table, we provide the emotion categories, the number of images that are used in the experiments and the protocol.

Database	Emotion Category	Sample Size	Protocol
HAPPEI	Six-level happiness intensity	2000 images	4-fold cross-validation
GAFF	Positive, Neutral, Negative	417 images	213 for training 204 for testing

6.2 Evaluation of INFA to the face, upper body and scene

As was previously mentioned in Section 4, the number of blocks of a face or upper body area, the PCA dimension and word size are three important parameters in INFA. For convenience, we denote them as $P \times Q$, D and K , respectively, in the following analysis. Additionally, different regional features may provide varying levels of performance to INFA. In this experiment, we will focus on the influence of these three parameters, various regional features and different encoding methods on face-level information.

TABLE 2: Mean absolute error (MAE) of INFA based on face information for various PCA dimensions and number of blocks on the HAPPEI database, where the bold number is the smallest MAE along the row and the underlined bold number is the lowest MAE in the table.

Word size	Number of blocks	PCA Dimension					
		16	32	64	128	256	512
50	4×4	0.6381	0.5873	0.5565	0.5488	0.5659	0.5545
50	8×8	0.6397	0.5951	0.5896	0.5761	0.5556	0.5540
50	16×16	0.6094	0.5975	0.569	0.5472	0.5432	0.5477
100	4×4	0.6257	0.5790	0.5507	0.5687	0.5584	0.5556
100	8×8	0.6131	0.5905	0.586	0.5508	0.5364	0.547
100	16×16	0.6052	0.5716	0.5659	0.5452	0.5293	0.5318
150	4×4	0.6186	0.5696	0.5681	0.5644	0.5682	0.5606
150	8×8	0.6094	0.5962	0.5804	0.5484	0.5363	0.5388
150	16×16	0.6014	0.5637	0.5537	0.5381	0.5255	0.5286

We will further investigate the effects of these three parameters on body-level and scene-level information.

6.2.1 Face

First, we evaluate the influence of the PCA dimension and the number of blocks. The word size of INFA is set as 50, 100 and 150. Three numbers of blocks $\{4 \times 4, 8 \times 8, 16 \times 16\}$ and six PCA dimensions $\{16, 32, 64, 128, 256, 512\}$ are considered. Table 2 shows the results of INFA using different numbers of blocks and PCA dimensions on the HAPPEI database.

Number of blocks: According to Table 2, the MAE is favorably decreased when more blocks are used. This means that GMM has enough training features to learn the feature distribution.

PCA dimension: According to Table 2, increasing the PCA dimension improves the performance. The best results are obtained at $D = 256$ for various word sizes and numbers of blocks on the HAPPEI database.

Fig. 7: Performance using INFA based on face information under various word sizes on the HAPPEI database.

Word size: Based on the above parameter setup, we further discuss the influence of word size K on the performance of INFA. Figure 7 shows the effects of different values of K on INFA. According to Figure 7, the MAE is considerably decreased as the word size is increased for all D . The INFA method obtains a comparative MAE of 0.5187 at $K = 180$. Large word size and suitable PCA dimension can provide promising performance of INFA. Based on the optimal word size, we further examine whether the descriptor differs significantly on various happiness intensity levels. Due to space limitations, we only show examples of the significance analysis on two descriptors under levels 0

(a)

(b)

Fig. 8: Two INFA descriptors of (a) a sample with happiness intensity level 0 and (b) a sample with happiness intensity level 5.

and 5. We randomly choose two descriptors after information aggregation, as shown in Figure 8. We conduct a t-test for the comparison between the two descriptors. We obtain $p = 1.8171e - 07$ ($p < .05$), which indicates that the two descriptors are found to be significantly different. For other cross-levels, we have also obtained the significance results for two descriptors. This shows that the feature that is obtained by INFA has a compact representation for different happiness intensity levels.

Regional features: We compare RVLBP with LBP and Local Phase Quantization (LPQ) [54] as region descriptors for INFA. The comparison is shown in Table 3. INFA using RVLBP outperforms LBP and LPQ. RVLBP provides more useful and discriminative information than LBP and LPQ because RVLBP

incorporates the spatial information and the co-occurrence statistics on the frequency and orientation domains of the higher-order Riesz transform.

TABLE 3: Performances of different feature descriptors as regional features for INFA on the HAPPEI database. The results are reported in terms of Mean Absolute Error (MAE). The best result is shown in bold.

Algorithm	Parameters	MAE
LBP	16×16 , $D = 32$, $W = 150$	0.6092
LPQ	16×16 , $D = 128$, $W = 160$	0.6047
RVLBP	16×16 , $D = 256$, $W = 180$	0.5187

To determine the impact of descriptor and word size, we conduct statistical significance analysis for the face on three feature descriptors: LBP, LPQ and RVLBP. For simplicity, we set D as 16 and choose 16×16 blocks. Figure 9 describes the MAEs of the three feature descriptors across word size [10 200]. We conduct t-tests for the pairwise comparisons between descriptors. We obtain $p = 6.35e - 07$ ($p < .05$), $p = 7.2441e - 13$ ($p < .05$), and $p = 3.8208e - 13$ ($p < .05$) for LBP-LPQ, LBP-RVLBP, and LPQ-RVLBP, respectively. Furthermore, we conduct t-tests for the pairwise comparison between word sizes. For most pairwise size comparisons, p is larger than 0.05. On average, $p = 0.2469$. The analysis results indicate that the descriptor significantly influences the MAE, while the word size has less impact on MAE. Therefore, using an appropriate descriptor is more important for group emotion representation than word size.

Fig. 9: Comparative performance of INFA based on three descriptors on HAPPEI database.

TABLE 4: Comparison of BoW, VLAD, FV and INFA for encoding face regional features on the HAPPEI database. The best performance is highlighted in bold.

Measurement	Algorithm			
	BoW	VLAD	FV	INFA
MAE	0.7015	0.6384	0.5383	0.5187

Finally, we compare INFA with Bag-of-Words (**BoW**), Vector of Locally Aggregated Descriptor (**VLAD**) and Fisher Vector (**FV**), where the k-means method is used for all regional features for BoW and a visual dictionary is obtained by vector quantization for VLAD. The values of K and D for FV are the same as for INFA. Support vector regression (**SVR**) based on a linear kernel

is used. Table 4 shows the MAEs of BoW, VLAD and INFA. According to Table 4, BoW and VLAD fail to yield promising results, while FV yields similar results to INFA. Moreover, INFA outperforms FV, as INFA uses the non-linear additive kernel and power normalization.

6.2.2 Upper body

The upper body images are resized to 256×256 . A three-level PHOG [32] is computed first on the upper body sub-regions. The number of bins for the histogram representation is set to 10 and orientation range is $[0, 360]$. For the LBP descriptor, the number of neighbors is 8 and the radius is 3 for one sub-region. After concatenation, a 1106-dimensional feature vector is obtained for each sub-region. In this section, we aim at evaluating the contribution of body information to GER. Parameter setup and classification are the same as in Section 6.2.1.

PCA dimension and number of blocks: Tables 5 shows the results of INFA based on body information with different PCA dimensions and numbers of blocks. According to the table, INFA with 8×8 blocks achieves the lowest MAE of 0.7339. This demonstrates that increasing the number of divided blocks can promisingly improve the performance of INFA on body information. The results in Tables 2 and 5 suggest that using more blocks can consistently make GMM more compact in terms of the distribution of samples.

Word size: We perform the experiment to evaluate how varying the word size affects INFA based on body information. The effect of K is presented in Figure 10. For the parameter setup, the number of blocks is set as 8×8 and the PCA dimension as 16. It is observed that the increasing word size does not always significantly degrade the MAE on the upper body information. INFA yields an MAE of 0.7164 at a word size of 40 for the HAPPEI database.

Finally, we observe that PHOG and LBP achieve MAEs of 0.7647 and 0.7515 with the above-mentioned parameters, respectively. PHOG and LBP have complementary information to each other.

6.2.3 Scene

In this section, we aim at determining the benefit of using scene information for GER. As was mentioned in Section 3.3, the parameter for the number of blocks is eliminated by using the superpixel method. According to our experimental results, we choose $N_s = 200$, as it achieves the best performance under a PCA dimension of 16. In this section, we discuss how word size and PCA dimension leverage the performance of INFA when fixing 200 superpixels for the group-level image. The results are shown in Figure 11. According to Figure 11, the scene achieves the lowest MAE of 0.7151 using $K = 30$ and $D = 16$. The scene obtains promising results using a small word size comparing with the face-level and upper body level. This may be explained by the following: (1) we used the superpixel method to extract more semantic information of the scene than the block division method and (2) SIFT is used to well characterize the local gradient information of each superpixel region. These two approaches may provide discriminative information to INFA. On the other hand, according to the group-level images that are shown in Figures 5 and 6, the scene may provide interesting information regarding the face and upper body, such as the clothing color and gender. As surveyed in [11], most of the participants mentioned information such as age, gender and attractiveness as attributes that affect their

TABLE 5: MAE of INFA based on upper body information for various PCA dimensions and numbers of blocks on the HAPPEI database, where the bold number is the lowest MAE along the row and the underlined bold number is the lowest MAE on all results.

Word size	Number of Blocks	PCA Dimension						
		16	32	64	128	256	512	1024
60	4×4	0.7558	0.7629	0.7593	0.7553	0.7552	0.7464	0.7420
60	8×8	0.7339	0.7460	0.7494	0.7428	0.7388	0.7344	0.7338
60	16×16	0.7457	0.7577	0.7410	0.7468	0.7459	0.7340	0.7383
140	4×4	0.7462	0.7635	0.7612	0.7541	0.7522	0.7439	0.7482
140	8×8	0.7530	0.7519	0.7520	0.7492	0.7507	0.7438	0.7397
140	16×16	0.7550	0.7499	0.7358	0.7426	0.7374	0.7379	0.7393
200	4×4	0.7501	0.7671	0.7607	0.7639	0.7582	0.7491	0.7519
200	8×8	0.7545	0.7497	0.7517	0.7519	0.7463	0.7493	0.7403
200	16×16	0.7501	0.7433	0.7358	0.7407	0.7407	0.7397	0.7375

Fig. 10: Performance of INFA based on upper body information under word sizes that range in [10 200] on the HAPPEI database.

Fig. 11: Performance of INFA using scene information under word sizes that range in [10 200] and three dimensionalities on the HAPPEI database.

perception of the affect of a group of people in an image. Thus, for the scene-level, more information is stored per word size than for the face-level and upper body-level. Therefore, a small word size is sufficient for the scene-level.

Following [11], we employ CENTRIST, GIST and the transferred deep learning features that are based on available VGG

TABLE 6: Comparison of CENTRIST, GIST and INFA_{Scene} on the HAPPEI database. The best result is shown in bold.

Database	Algorithm		
	CENTRIST [36]	GIST [37]	INFA _{Scene}
HAPPEI	0.7433	0.7517	0.7151

TABLE 7: Comparison of transfer features based on deep VGG models and INFA_{Scene} on the HAPPEI database. Bold indicates the best result across the FC layer.

Model	Layer		
	FC6	FC7	FC8
VGG-deep-16 [55]	1.1674	0.7941	0.749
VGG-deep-19 [55]	1.1637	0.7643	0.7523
CNN F [56]	1.1707	0.7887	0.7527
CNN S [56]	1.0903	0.7797	0.7591
CNN M [56]	1.0827	0.7732	0.7715
CNN M 128 [56]	0.9434	0.7412	0.7349
CNN M 1024 [56]	0.9898	0.7713	0.7473
CNN M 2048 [56]	0.9908	0.7493	0.7396

models⁴ [55], [56] for comparison to our INFA features. The publicly available GIST implementation⁵ is used with its default parameters: Orientations per scale is [8, 8, 8, 8]; number of blocks is 4. Similarly, for the CENTRIST descriptor, a publicly available implementation⁶ is used. For the deep learning features, we use the high-level features that are transferred from a trained deep model for the scene. For selecting high-level features, we extract the deep features from the ‘FC6’, ‘FC7’ and ‘FC8’ layers, respectively. The comparison results are shown in Tables 6 and 7. According to Table 6, CENTRIST performs better than INFA but their results are similar. INFA outperforms GIST. Among the features that are derived from VGG models, CNN M 128 achieves the lowest MAE of 0.7412. INFA_{Scene} outperforms CNN M 128 by decreasing MAE to 0.7151. This demonstrates that INFA can still achieve a satisfactory performance for GER.

6.3 Performance of the multi-modal framework

According to above-mentioned experimental setup, we obtain high-performing parameters for the face, upper body and scene

4. <http://www.vlfeat.org/matconvnet/pretrained/>

5. <http://people.csail.mit.edu/torralba/code/spatialenvelope/>

6. <http://github.com/sometimesfood/spact-matlab>

on the HAPPEI database. Table 8 presents the designed parameters in this scenario, where INFA on the face, upperbody and scene obtains MAEs of 0.5187, 0.7164 and 0.7151, respectively. According to Table 8, the face performs the best, followed closely by the scene and more distantly by the upper body. As these results show, it is empirically demonstrated that the face contributes most to GER. In the following scenario, we will use revisited localized multiple kernel learning to automatically learn the contributions of the three modalities to GER and determine how they are weighted for GER.

TABLE 8: Designed parameters (region number, reduced dimensionality and word size) and MAEs of INFA on the HAPPEI database for the face, upper body and scene. The best performance among the face, upper body and scene is highlighted in bold.

Information	Region number	Reduced dimensionality	Word size	MAE
Face	16×16	256	180	0.5187
Upperbody	8×8	16	40	0.7164
Scene	200	16	30	0.7151

TABLE 9: Performance of the multi-modal framework, where F, B and S represents the face, body and scene, respectively. The best result is highlighted in bold.

Kernel	Modality			
	F+B	F+S	B+S	F+B+S
Linear	0.5148	0.5109	0.6154	0.4821
Gaussian	0.5242	0.5177	0.6349	0.5107
HI	0.5309	0.5148	0.7243	0.5002

We test the performances of revisited localized multiple kernel learning based on linear, Gaussian and HI kernels. Specifically, the same kernel function is employed for the three modalities. We choose the optimal values of C and ϵ using 10-fold-cross-validation on the training set. For the Gaussian kernel, we choose $s = 10$ for the standard deviation. For the testing procedure, face/upper body detection may fail to work. To address this situation, we set its learned MKL weight as 0 in our implementation. This way, it can ensure that the failed face/upper body cannot provide unnecessary information to the classification. Table 9 shows the results of combining different modalities. Revisited LMKL based on the linear kernel function performs better than using the two other kernel functions. The performance of the face is improved by using the scene feature as complementary information, as attributes such as pleasant scene and clothes may affect the perception of human beings of the affect of a group of people in an image [11]. It is observed that combining scene with face features yields better performance than combining upper body information. This demonstrates that the upper body information provides little information to the face. It is very interesting that the fusion of the face, upper body and scene information performs the best out of all tested configurations. Combining three modalities yields the very promising results compared to each individual modality. Since we conduct the experiments using the four-fold cross-validation protocol, we average the weights of the three modalities that are learned by RLMKL across four folds. The weights are 0.6912, 0.1014 and 0.2074 for the face, upper body and scene, respectively.

The contributions that are indicated by the learned weights yield the same conclusion as Table 8. Additionally, that based on the fusion of the upper body and the scene performs less well than that based on the fusion of face and the upper body/scene. Given the ‘in the wild’ nature of the images, face detection may fail. Therefore, we need to add complementary information, which is obtained from the scene-level and body-level descriptors.

Some exemplar results are shown in Figure 12. The multi-modal framework performs better than the state-of-the-art algorithms [6], [7] in Figures 12(b), 12(c) and 12(f). The estimated happiness intensity by using the multi-modal framework is similar to the perception of human beings. The clothing color (for example, in Figure 12(b)) and active body pose (for example, in Figure 12(c)) may provide useful multi-modal information to the multi-modal framework. Unfortunately, it is also observed that no approach can well predict the intensity level of 5 in Figures 12(c) and 12(d); however, the approaches perform better for intensity levels of 2 and 3. In the HAPPEI database, an intensity level of 5 occurs in only 35 of 2000 images, while other levels occur in many more images. For example, level 2 occurs in 600 images. Therefore, the poor performance for intensity level of 5 is due to small number of instances. According to Figure 12(d), the face-level information can yield better performance than the recent GEM model, but the multi-modal framework works as well as GEM based on CCRF. It may be because the complicated background and the proximity of the body to the camera cause noise in the face-level information. In Figures 12(a) and 12(e), the multi-modal framework achieves satisfactory performance in challenging situations, for example, in bad illumination.

6.4 Algorithm comparison

Dhall et al. [6] proposed the GEM based on the average, weight and Latent Dirichlet Allocation (**LDA**) for estimating the happiness intensity of a group of people. Huang et al. [7] proposed a new GEM that based on Continuous Conditional Random Field (**CCRF**) for happiness intensity estimation.

TABLE 10: Algorithm comparison in terms of mean absolute error (MAE) on the HAPPEI database, where the results of GEM are from [7]. The numbers inside brackets are the correlation coefficients. The best performance is highlighted in bold.

Algorithm	MAE	Algorithm	MAE
GEM _{avg}	0.5622 (0.6887)	INFA _{Face}	0.5187 (0.7291)
GEM _{weight}	0.5469 (0.6902)	INFA _{Upperbody}	0.7164 (0.1348)
GEM _{LDA}	0.5407 (0.6684)	INFA _{Scene}	0.7151 (0.1521)
GEM _{CCRF}	0.5292 (0.7209)	Multi-modal	0.4821 (0.7521)

Additionally, the deep neural networks are proposed for predicting the happiness that is displayed by a group of people in images [57], [58], [59]. In [57], they use Long Short Term Memory (LSTM) networks, including face happiness intensity and the spatial distribution of faces in a group. Sun et al. used an LSTM and GEM model for estimating the happiness of a group of

Fig. 12: Six exemplar results using the multi-modal framework on the HAPPEI database. The number in brackets represents the absolute value of the difference between the estimated intensity and the ground truth. A smaller number means that the predicted value is more similar to the ground truth. The best result and the second best result are highlighted in bold red and bold blue, respectively.

people [58]. In [59], Li et al. also proposed LSTM with feature-level augmentation by considering two sets of feature vectors that are trained with the same settings but different initializations of the facial feature extraction set and use the scene features as extra information in the LSTM model. In this scenario, we compare INFA with all GEM models [7] and three state-of-the-art works that use deep neural networks [57], [58], [59]. The comparative results are illustrated in Tables 10 and 11. According to Table 10, INFA using face information achieves competitive results comparing with all GEM models. It decreases mean absolute error by 0.0457 and increases correlation coefficient by 0.0082. However, for upper body and scene information, INFA performs worse than GEM models that are based on face information. Combining face, upper body and scene information acceptably improves the performance of the face in terms of mean absolute error and increases the correlation coefficients, since it uses additional information regarding the face from the upper body and the scene.

Statistical significance analysis: We also conduct t-tests for the comparisons between the multi-modal approach and the state-of-the-art methods in Tables 10 and 11. In Tables 10 and 11, we obtain $p = .00005$ ($p < 0.05$), which indicates that compared with the state-of-the-art methods, the improvements that are achieved by the multi-modal approach are significant. Furthermore, we perform t-tests for the comparison between multi-modal and sole modalities. We obtain $p = .0309$ ($p < 0.05$). This indicates that the multi-modal framework outperforms the sole modalities.

Based on intensive comparisons on the HAPPEI database, our

TABLE 11: Comparison of the proposed framework with three state-of-the-art algorithms [57], [58], [59]. Root-mean square error (RMSE) is used as performance measurement. Bold indicates the best result.

Methods	RMSE
Face-CEN [57]	0.55
LSTM+Scene+Feature Aug. [59]	0.516
CNN-LSTM [58]	0.494
$INFA_{Face}$	0.482
$INFA_{Upperbody}$	0.9626
$INFA_{Scene}$	0.9514
Multi-modal	0.4516

multi-modal method achieves competitive performance for group-level happiness intensity estimation. Additionally, we also show that upper body and scene information contribute to the analysis of the emotional state of a group of people.

6.5 Evaluation of the multi-modal system on the GAFF database

Based on the well-designed parameters, we evaluate the performance of our proposed method on the GAFF database [11]. In the experiment, 417 out of 504 images are chosen in our experiments, since face detection failed to work on 87 images. Following [11], 213 images are chosen for the training set and 204 for the test

set. Different from the HAPPEI database, the task in the GAFF database is to classify group-level images into ‘Positive’, ‘Neutral’ or ‘Negative’. The recognition rate is used as a measure. Using the designed parameters in Table 8, we carry out the evaluation of INFA and multi-modal methods on the GAFF database. We also compare our results with those of the baseline algorithm [11]. In [11], Dhall et al. proposed using the Bag-of-Words model to obtain the features for group-level facial expression images. Then, they combined Action-unit-based face representation (BoW_AU) and low-level features (BoW_LL) for face and scene information (Scene_GIST and Scene_CENTRIST) for affective information on the GAFF database. The results of the baseline algorithm [11] and our methods are presented in Table 12.

TABLE 12: Algorithm comparison on the GAFF database, where results of the baseline algorithms are from [11]. The best result is shown in bold.

Algorithm	Recognition rate
BoW_AU [11]	50.43%
BoW_LL [11]	50.98%
Scene_GIST [11]	42.16%
Scene_CENTRIST [11]	45.58%
BoW_LL+BoW_AU+Scene_CENTRIST [11]	67.64%
Deep Neural Network + Bayesian Network [60]	64.68 %
INFA _{Face}	58.33%
INFA _{Upperbody}	46.08%
INFA _{Scene}	48.28%
INFA _{Face} +INFA _{Upperbody}	62.75%
INFA _{Face} +INFA _{Scene}	60.78%
INFA _{Scene} +INFA _{Upperbody}	49.51%
Multi-modal	69.77%

According to Table 12, on the face-level, BoW_LL obtained an accuracy of 50.98%, while INFA obtained 58.33%. INFA outperforms their face-level approach (BoW_LL) on the GAFF database. According to this table, INFA obtains the recognition rates of 58.33%, 46.08% and 48.28% for face, upper body and scene information, respectively. We use revisited localized multiple kernel learning based on a linear kernel to combine face, upper body and scene information. By adding sole information such as upper body or scene information, INFA based on the face considerably boosts the recognition rate. Combining upper body and scene information results in only a small increase in performance of 1.1%. The face plays an important role in recognizing emotion recognition, followed by upper body and the scene. The multi-modal framework obtains a recognition rate of 66.67%. Comparing with the baseline algorithm, Dhall et al. [11] combined their features, including Action-unit-based face representation and scene features, and achieved an accuracy of 67.64%, which is lower than that of our multi-modal framework.

Statistical significance analysis: We conduct t-tests for the comparisons between the multi-modal approach and the state-of-the-art methods on the GAFF database. In addition, we perform t-tests to compare with sole modalities. We obtain $p = .0176$ ($p < 0.05$) and $p = .0377$ ($p < 0.05$), respectively. These results indicate that the multi-modal approach achieves substantial improvement compared with the state-of-the-art methods and sole modalities.

According to the algorithm comparison and statistically significant test results on the GAFF database, our multi-modal method achieves promising performance in predicting group-level emo-

tion. Moreover, the multi-modal framework achieves significant improvement over the sole modalities.

7 CONCLUSIONS

In this paper, a multi-modal method that combines face, upper body and scene information has been presented for analyzing the affect state of a group of people in an image. First, for robustness, we exploit three interesting types of information, namely, face, upper body and scene information, in an image. The face and upper body are viewed as the bottom-up components while the scene is viewed as the top-down component. To represent an image, information aggregation was proposed for encoding multiple people’s information for a group-level image. A robust multi-modal framework that fuses face, upper body and scene information is finally presented to determine the affective state of a group of people.

We have conducted experiments on two challenging group-level emotion recognition databases. We show that INFA considerably improves the performance for group-level emotion recognition. Additionally, we evaluate the multi-modal framework on the HAPPEI and GAFF databases. Intensive experiments demonstrate that our multi-modal framework predicts the perceived group mood more accurately. As mentioned in Section 3.3, we set the number of superpixels as 200 for images according to our experimental results. However, the image resolutions of group-level images may be quite different. As a result, this number may not well explore the semantic subregions. Additionally, according to Figure 2, we found that superpixel algorithm segments a region that excludes the face and upper body. It may provide the unnecessary and redundant information for information aggregation. In future work, we will attempt to find a flexible method for alleviating this problem that is caused by superpixel algorithm for group-level emotion recognition.

REFERENCES

- [1] A. Dhall, A. Asthana, and R. Goecke, “Facial expression based automatic album creation,” in *ICONIP*, 2010, pp. 485–492.
- [2] T. Vandal, D. McDuff, and R. Kaliouby, “Event detection: ultra large-scale clustering of facial expression,” in *AFGR*, 2015, pp. 1–8.
- [3] S. Barsäde and D. Gibson, “Group emotion: A view from top and bottom,” *Research on Managing in Group and Teams*, vol. 1, pp. 81–102, 1998.
- [4] J. Kelly and S. Barsäde, “Mood and emotions in small groups and work teams,” *Organizational behavior and human decision processes*, vol. 86, no. 1, pp. 99–130, 2001.
- [5] A. Dhall, J. Joshi, I. Radwan, and R. Goecke, “Finding happiest moments in a social context,” in *ACCV*, 2012, pp. 613–626.
- [6] A. Dhall, R. Goecke, and T. Gedeon, “Automatic group happiness intensity analysis,” *IEEE Trans. Affective Computing*, vol. 6, no. 1, pp. 13–26, 2015.
- [7] X. Huang, A. Dhall, G. Zhao, R. Goecke, and M. Pietikäinen, “Riesz-based volume local binary pattern and a novel group expression model for group happiness intensity analysis,” in *BMVC*, 2015, pp. 1–9.
- [8] Z. Zeng, J. Tu, M. Liu, T. Huang, B. Pianfetti, D. Roth, and S. Levinson, “Audio-visual affect recognition,” *IEEE Trans. Multimedia*, vol. 9, no. 2, pp. 424–428, 2007.
- [9] J. Lin, C. Wu, and W. Wei, “Error weighted semi-coupled hidden markov model for audio-visual emotion recognition,” *IEEE Trans. Multimedia*, vol. 14, no. 1, pp. 142–156, 2012.
- [10] C. Wu, J. Lin, and W. Wei, “Two-level hierarchical alignment for semi-coupled HMM-based audiovisual emotion recognition with temporal course,” *IEEE Trans. Multimedia*, vol. 15, no. 8, pp. 1880–1895, 2013.
- [11] A. Dhall, J. Joshi, K. Sikka, R. Goecke, and N. Sebe, “The more the merrier: analysing the affect of a group of people in images,” in *AFGR*, 2015, pp. 1–6.

- [12] W. Mou, O. Celiktutan, and H. Gunes, "Group-level arousal and valence recognition in static images: face, body and context," in *AFGR*, 2015, pp. 1–6.
- [13] J. Joshi, A. Dhall, R. Goecke, and J. Cohn, "Relative body parts movement for automatic depression analysis," in *ACII*, 2013, pp. 492–497.
- [14] M. Hoai and A. Zisserman, "Talking heads: Detecting humans and recognizing their interaction," in *CVPR*, 2014, pp. 875–882.
- [15] K. Schindler, L. Gool, and B. Gelder, "Recognizing emotions expressed by body pose: A biologically inspired neural model," *Neural Network*, vol. 21, no. 9, pp. 1238–1246, 2008.
- [16] M. Kret, K. Roelofs, J. Stekelenburg, and B. Gelder, "Emotional signals from faces, bodies and scenes influence observers' face expressions, fixations and pupil-size," *Front. Hum. Neurosci.*, vol. 7, no. 810, pp. 1–9, 2013.
- [17] R. Righart and B. Gelder, "Context influences early perceptual analysis of faces - an electrophysiological study," *Cereb. Cortex*, vol. 16, no. 1, pp. 1249–1257, 2006.
- [18] —, "Rapid influence of emotional scenes on encoding of facial expressions: an erp study," *Soc. Cogn. Affect. Neurosci.*, vol. 3, no. 3, pp. 270–278, 2008.
- [19] A. Dhall and R. Goecke, "A temporally piece-wise fisher vector approach for depression analysis," in *ACII*, 2015, pp. 255–259.
- [20] W. Mou, H. Gunes, and I. Patras, "Automatic recognition of emotions and membership in group videos," in *CVPRW*, 2016, pp. 1478–1486.
- [21] F. Perronnin, J. Sánchez, and T. Mensink, "Improving the fisher kernel for large-scale image classification," in *ECCV*, 2010, pp. 143–156.
- [22] K. Simonyan, O. Parkhi, A. Vedaldi, and A. Zisserman, "Fisher vector faces in the wild," in *BMVC*, 2013, pp. 1–6.
- [23] T. Uricchio, M. Bertini, L. Seidenari, and A. Bimbo, "Fisher encoded convolutional bag-of-windows for efficient image retrieval and social image tagging," in *ICCV Workshop*, 2015, pp. 1–7.
- [24] W. Zhang, S. Shan, W. Gao, and H. Zhang, "Local Gabor binary pattern histogram sequence (LGBPHS): a novel non-statistical model for face representation and recognition," in *ICCV*, 2005, pp. 786–791.
- [25] C. Liu and H. Wechsler, "Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition," *IEEE Trans. on Image Process.*, vol. 11, no. 4, pp. 467–476, 2002.
- [26] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary pattern," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 24, no. 7, pp. 971–987, 2002.
- [27] M. Felsberg and G. Sommer, "The monogenic signal," *IEEE Trans. Signal Process.*, vol. 49, no. 12, pp. 3136–3144, 2001.
- [28] L. Zhang, D. Zhang, Z. Guo, and D. Zhang, "Monogenic-LBP: a new approach for rotation invariant texture classification," in *ICIP*, 2010, pp. 2677–2680.
- [29] L. Zhang and H. Li, "Encoding local image patterns using Riesz transform: With applications to palmpring and finger-knuckle-print recognition," *Image Vis Comput.*, vol. 30, no. 12, pp. 1043–1051, 2012.
- [30] C. Zetsche and E. Barth, "Fundamental limits of linear filters in the visual processing of two dimensional signals," *Vision Research*, vol. 30, no. 7, pp. 1111–1117, 1990.
- [31] J. Pan and Y. Tang, "Texture classification based on BIMF monogenic signals," in *ACCV*, 2012, pp. 236–243.
- [32] A. Bosch, A. Zisserman, and X. Munoz, "Representing shape with a spatial pyramid kernel," in *CIVR*, 2007, pp. 401–408.
- [33] C. Weinrich, C. Vollmer, and H. Gross, "Estimation of human upper body orientation for mobile robotics using an svm decision tree on monocular images," in *IROS*, 2012, pp. 2147–2152.
- [34] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *CVPR*, 2005, pp. 886–893.
- [35] X. Wang, T. Han, and S. Yan, "An HOG-LBP human detector with partial occlusion handling," in *ICCV*, 2009, pp. 32–39.
- [36] J. Wu and J. Rehg, "Centrist: a visual descriptor for scene categorization," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 33, no. 8, pp. 1489–1501, 2011.
- [37] A. Oliva and A. Torralba, "Modeling the shape of the scene: a holistic representation of the spatial envelop," *IJCV*, vol. 42, no. 3, pp. 145–175, 2001.
- [38] Z. Li and J. Chen, "Superpixel segmentation using linear spectral clustering," in *CVPR*, 2015, pp. 1356–1363.
- [39] C. Xu, D. Tao, and C. Xu, "A survey of multi-view learning," *arXiv preprint arXiv:1304.5634*, pp. 1–59, 2016.
- [40] F. Wu, X. Jing, X. Dong, Q. Ge, S. Wu, Q. Liu, D. Yue, and J. Yang, "Uncorrelated multi-set feature learning for color face recognition," *Pattern Recognition*, vol. 60, pp. 630–646, 2016.
- [41] Y. Luo, Y. Wen, D. Tao, J. Gui, and C. Xu, "Large margin multi-modal multi-task feature extraction for image classification," *IEEE Trans. Image Process.*, vol. 25, no. 1, pp. 414–427, 2016.
- [42] F. Wu, X. Jing, X. You, D. Yue, R. Hu, and J. Yang, "Multi-view low-rank dictionary learning for image classification," *Pattern Recognition*, vol. 50, pp. 143–154, 2016.
- [43] Y. Han, K. Yang, Y. Ma, and G. Liu, "Localized multiple kernel learning via sample-wise alternating optimization," *IEEE Trans. Cybernetics*, vol. 44, no. 1, pp. 137–148, 2014.
- [44] K. Sikka, K. Dykstra, S. Sathyanarayana, G. Littlewort, and M. Bartlett, "Multiple kernel learning for emotion recognition in the wild," in *ICMI*, 2013, pp. 517–524.
- [45] M. Gönen and E. Alpaydin, "Localized multiple kernel learning," in *ICML*, 2008, pp. 1–8.
- [46] —, "Localized multiple kernel regression," in *ICPR*, 2010, pp. 1425–1428.
- [47] S. Fischer, F. Šroubek, L. Perrinet, R. Redondo, and G. Cristóbal, "Self-invertible 2D log-Gabor wavelets," *IJCV*, vol. 75, no. 2, pp. 231–246, 2007.
- [48] G. Zhao, M. Barnard, and M. Pietikäinen, "Lipreading with local spatiotemporal descriptors," *IEEE Trans. Multimedia*, vol. 11, no. 7, pp. 1254–1265, 2009.
- [49] M. Eichner and V. Ferrari, "Better appearance models for pictorial structures," in *BMVC*, 2009, pp. 1–11.
- [50] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bag of features: spatial pyramid matching for recognizing natural scene categories," in *CVPR*, 2006, pp. 2169–2178.
- [51] D. Lowe, "Object recognition from local scale-invariant features," in *ICCV*, 1999, pp. 1150–1157.
- [52] L. Chiu, T. Chang, J. Chen, and N. Chang, "Fast sift design for real-time visual feature extraction," *IEEE Trans. Image Process.*, vol. 22, no. 8, pp. 3158–3167, 2013.
- [53] F. Perronin, J. Sanchez, and T. Mensink, "Improving the fisher kernel for large-scale image classification," in *ECCV*, 2010, pp. 143–156.
- [54] V. Ojansivu and J. Heikkilä, "Blur insensitive texture classification using local phase quantization," in *ICISP*, 2008, pp. 236–243.
- [55] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015, pp. 1–14.
- [56] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional networks," in *BMVC*, 2014, pp. 1–12.
- [57] A. Cerekovic, "A deep look into group happiness prediction from images," in *ICMI*, 2016, pp. 437–444.
- [58] B. Sun, Q. Wei, L. Li, Q. Xu, J. He, and L. Yu, "LSTM for dynamic emotion and group emotion recognition in the wild," in *ICMI*, 2016, pp. 451–457.
- [59] J. Li, S. Roy, J. Feng, and T. Sim, "Happiness level prediction with sequential inputs via multiple regressions," in *ICMI*, 2016, pp. 487–493.
- [60] L. Surace, M. Patacchiola, E. Sönmez, W. Spataro, and A. Cangelosi, "Emotion recognition in the wild using deep neural networks and bayesian classifiers," in *ICMI*, 2017, pp. 593–597.