

Multimodal Channel-Mixing: Channel and Spatial Masked AutoEncoder on Facial Action Unit Detection

Xiang Zhang¹ Huiyuan Yang² Taoyue Wang¹ Xiaotian Li¹ Lijun Yin¹

¹State University of New York at Binghamton ²Rice University

{zxiang4, twang61, xli210, lyin}@binghamton.edu hy48@rice.edu

Abstract

Recent studies have focused on utilizing multi-modal data to develop robust models for facial Action Unit (AU) detection. However, the heterogeneity of multi-modal data poses challenges in learning effective representations. One such challenge is extracting relevant features from multiple modalities using a single feature extractor. Moreover, previous studies have not fully explored the potential of multi-modal fusion strategies. In contrast to the extensive work on late fusion, there are limited investigations on early fusion for channel information exploration. This paper presents a novel multi-modal reconstruction network, named **Multimodal Channel-Mixing (MCM)**, as a pre-trained model to learn robust representation for facilitating multi-modal fusion. The approach follows an early fusion setup, integrating a Channel-Mixing module, where two out of five channels are randomly dropped. The dropped channels then are reconstructed from the remaining channels using masked autoencoder. This module not only reduces channel redundancy, but also facilitates multi-modal learning and reconstruction capabilities, resulting in robust feature learning. The encoder is fine-tuned on a downstream task of automatic facial action unit detection. Pre-training experiments were conducted on BP4D+, followed by fine-tuning on BP4D and DISFA to assess the effectiveness and robustness of the proposed framework. The results demonstrate that our method meets and surpasses the performance of state-of-the-art baseline methods.

1. Introduction

Facial Action Coding System (FACS) [10] is designed to detect subtle facial changes by describing facial muscle movement, while the combination of basic AUs (Action Units) characterizes human expressions. Automatic AUs detection has been studied for decades in human emotion analysis and human facial behavior analysis. AUs may occur individually or in combinations, hence AUs detection




Figure 1. The Channel-Mixing module in our approach involves concatenating the channels of visual RGB, Depth, and Thermal modalities (*RGBDT*). Subsequently, two channels are randomly dropped, and the three remaining channels in each patch are also randomly ordered. Hence, our pre-training task is to reconstruct the missing channels from the remaining ones.

has been mainly treated as a multi-label classification problem, which is a challenge in computer vision.

In recent years, deep learning networks have shown popularity and advancement in many fields, e.g., computer vision, natural language processing, and autonomous driving. Their successful application on large-scale datasets has also promoted a thriving research community on facial expression recognition and AU detection. By using 2D images, several methods achieve state-of-the-art performance in AU detection [7, 15, 17, 20, 29, 33, 42]. However, these supervised approaches are very scarce for labeled data, and AU annotation is a highly labor-intensive and time-consuming process, thus researchers attempt to utilize unlabeled data through self-supervised learning [23, 27]. Moreover, as some subtle facial muscle movements are not easily observed in the images, multimodal approaches are necessary for efficient AU detection. For instance, *AU6* (cheek raiser) involves the deformation of *Orbicularis oculi* and *pars orbitalis* muscles, which shows subtle differences when observed in images but appears distinct geometric variations on the 3D mesh. Similarly, human emotional variations can lead to changes in blood pressure, skin surface temperature, and heart rate, which are captured by physiological machines and thermal cameras.

Thanks to recently developed multi-modal databases [40, 44, 45], progress has been made by exploiting additional information in multi-modal learning on human emotion tasks. For example, Li et al. [18, 19] combined the 2D and 3D features for facial expression recognition. Irani et al. [14] utilized the visual, depth, and thermal modalities for pain study. In AU detection, Zhang et al. [43] and Yang et al. [41] present some fusion network with visual RGB, depth, or thermal images. They're all based on late or middle fusion, where each modality feature was extracted by a separate encoder, and fused into the latent space. Late fusion methods have the advantage of dealing with missing modes, but they are often accompanied by an increase in network parameters while ignoring low-level modal interactions. In contrast, early fusion, as the simplest example of fusion by concatenating individual modality features, hasn't been fully explored in recent multimodal learning studies and is typically used as a baseline method for comparison. Recent works [34, 47] on early fusion, have explored the utilization of visual RGB and depth modalities. However, in these approaches, the two modalities are simply concatenated on the channel dimension, which is a relatively under-explored. To address these problems, we propose a multi-modal method named Multi-modal Channel-mixing Network (MCM). This model not only learns early fusion features through an encoder, but also handle the downstream tasks even when some modalities are missing. Firstly, the 5-channel input from the early fusion of RGB, Depth, and Thermal is processed through a patch-based Channel-Mixing (CM) module, resulting in a 3-channel output. Figure 1 shows the overview structure of the channel mixing module. Therefore, the channel mixing module ensures that the resulting 3-channel output is compatible with the backbone pre-trained on ImageNet [16]. Additionally, this design allows for adaptation to self-supervised learning tasks such as image reconstruction. Secondly, the mixed image is inputted into a ViT-based masked encoder, which extracts the fused latent feature. Subsequently, the latent features, along with mask tokens, are fed into three parallel ViT-based decoders to reconstruct the visual RGB, depth, and thermal images. Similar to heavy spatial redundancy in an image, e.g., a missing patch can be recovered from neighboring patches [12], channel redundancy also exists across in visual-depth-thermal data. Therefore, the Channel-Mixing (CM) module not only reduces channel redundancy but also enhances multi-modal learning and reconstruction capabilities, leading to robust feature learning.

This architecture draws inspiration from the concept of masked autoencoder (MAE) [12], expanding its scope from single modality to multiple modalities, from spatial masks to channel masks, and from object recognition to facial action unit detection. Compared with related work, our contributions are as follows:

1. A multimodal channel-mixing (*MCM*) scheme has been proposed for building a multi-modal pre-training framework. Such a novel scheme not only helps to learn a robust representation but also facilitates multi-modal fusion.
2. We have also proposed a unified framework for visual RGB, depth, and thermal modalities. It is more effective and efficient than the conventional multi-modal fusion strategies. Additionally, our proposed framework demonstrates good performance when some modalities are missing.
3. Extensive experiments have been conducted on three public datasets to evaluate the effectiveness of the proposed method. The results show that our model reaches and surpasses the state-of-the-art.

2. Related Work

Automatic facial action recognition can be roughly categorized into single-modality-based methods and multi-modality-based approaches. The most commonly used modality is visual RGB image and a broad range of previous works [5, 20, 21, 33, 36, 37, 46] has studied to use of a region of interest (ROI), image patches, facial landmarks, heat-map, attention mechanism, and other methods to localize detailed facial parts and learn discriminative feature representations. Considering the structural information and dependencies among different AUs on human faces, some recent works [15, 42, 43] utilized the self-attention mechanism of transformer [39] to learn the semantic dependency among different AUs.

Unstructured real-world data can inherently take many forms, also known as modalities, often including diverse representations of content. Multi-modal learning refers to an embodied learning situation that engages multiple sensory systems and action systems of the model. In the past decade, multi-modal learning has been studied on both facial expression recognition and action unit detection tasks. Li et al. [19] proposed a deep fusion network to learn the optimal combination weights of 2D and 3D facial representations for multi-modal 2D+3D FER. Irani et al. [14] proposed a method to utilize RGB-Thermal-Depth images for pain estimation. Li et al. [22] explained that EEG signals show a strong correlation with facial actions and eye blinking of both posed and spontaneous expressions. They utilize the early fusion of EEG features and RGB features to boost both posed and authentic facial action detection. Liu et al. [24] designed a crossmodal translation network encoding the RGB images to reconstruct thermal images, which is an end-to-end model for AU detection. Yang et al. [41] proposed a model called AMF, which included a feature scoring module to select the most relevant feature representations from different modalities. Recently, the transformer-

based [39] attention mechanism has been widely applied in multi-modal features fusion. For example, MulT [38] first attempted to fuse multi-modalities by transformer, which was applied to image, audio, and text. TransFuser [30] incorporates the global context of the 3D scene into the feature extraction layers of different modalities. MFT [43] utilizes the RGB-Depth features fusion by multi-head fusion attention to fuse AU features of two modalities in a transformer.

All of these are late fusion or middle fusion works, where each modality has its encoder to encode the latent features, which are then fused at a deeper level. Hence the parameter size is increased than the single encoder and the low-level interaction between the modalities is ignored. Recently, the line between multimodal representation and fusion has been blurred for models such as deep neural networks where representation learning interacts with classification or regression objectives [2]. Hence, we refer to early fusion as the immediate integration of the data source or feature from the extractor, often simply by concatenating them into a single vector as the input to the learning model. Recently, compared to the number of late fusion works, only a few works utilize the early fusion strategy, e.g., [22, 34, 47]. One disadvantage of early fusion is the trained model usually can not handle the missing modality. Although the conversion technique such as projection can be applied, the input size of the model is fixed.

Inspired by the success of the masked language model [4, 8, 31, 32] in NLP, masked image encoding methods of self-supervised learning have been improved in computer vision, often focusing on different pretext tasks for pre-training. BEiT [3] design an autoencoder network similar to BERT [8] with the ViT encoder and discrete image token. MAE [12] proposed another ViT backbone network that only encodes the visible patches and inserts masked tokens in decoding for reconstruction, which speeds up the training process. Inspired by MAE, we propose a multimodal channel mixing (MCM) module to explore the effects of multimodal channel-based masking mechanism. **Differ** from the previous late fusion works, our method only needs one encoder; **Differ** from the early fusion works, our downstream model can be fed with only one modality.

3. Methodology

In this section, we describe our method in detail, including Channel-Mixing Module, Masked Autoencoder, and AU detection network. The overall pipeline of our framework is illustrated in Figure 2. Multiple modalities are fed into the channel mixing module to get the 3-channels output. Following a spatial mask, the patches are encoded by a ViT encoder. The set of encoded patches and three sets of mask tokens are processed by three decoders that reconstruct the original RGB, Depth, and Thermal respectively.

3.1. Channel-Mixing

We have visual RGB image I_v and its corresponding depth image I_d , and thermal image I_t , so the fused image is $I_f = I_v \oplus I_d \oplus I_t$, where \oplus represents concatenation by channels. Then the fused image $i_f \in \mathbb{R}^{H \times W \times C}$ is split to patches $i_p \in \mathbb{R}^{L \times (p^2 \cdot C)}$, where (H, W) is its resolution, C is the number of channels ($C = 5$ in our task), (P, P) is the resolution of each image patch, and $L = HW/P^2$ is the resulting number of patches [9]. Then each patch is randomly dropped with two channels and then shuffled in the channel, so the channel-mixing image set is $I_{mix} = \{i_{mix}\}$, where $i_{mix} \in \mathbb{R}^{L \times (p^2 \cdot (C-2))}$. Detail see in Figure 1.

The proposed channel mixing module has the following advantages:

1. It downsamples the fused modality to 3 channels, which can take advantage of the large pre-trained parameters on Imagenet, allowing the AU detection model to potentially accept all modalities as input.
2. It increases the difficulty of the reconstruction task so that the network focuses on both spatial and channel information.
3. The encoder is guided by the reconstruction task to learn the most relevant latent feature for all three modalities, resulting in an effective and robust multimodal representation.

3.2. Spatial Mask

Patches are randomly removed based on a uniform distribution. This approach has the potential to mask more patches near the center of the image, where the face is typically located. With high masking rates, redundancy is largely eliminated, resulting in a task that is not easily solved by visible neighboring patches inference. Since the following encoder only proceeds on unmasked patches, it speeds up the training process. In addition, we assess the performance of our model using various masking rates. The details and results of these evaluations are presented in the supplementary materials.

3.3. ViT Encoder

Our encoder is a ViT [9], and only proceeds with the visible, unmasked patches. Unlike using learnable 1D position embedding in the standard ViT, we use the 2-D sine-cosine variant [6] for position embedding. The embedding is then encoded by a stack of Transformer blocks following the design in [9]. Let w_l be the encoded features at the l -th layer in transformer blocks, W_0 being the input layer. The features at the l -th layer are obtained by applying a transformer block defined as:

$$w'_l = F_{LN}^l(w_{l-1} + F_{MSA}^l(w_{l-1})) \quad (1)$$

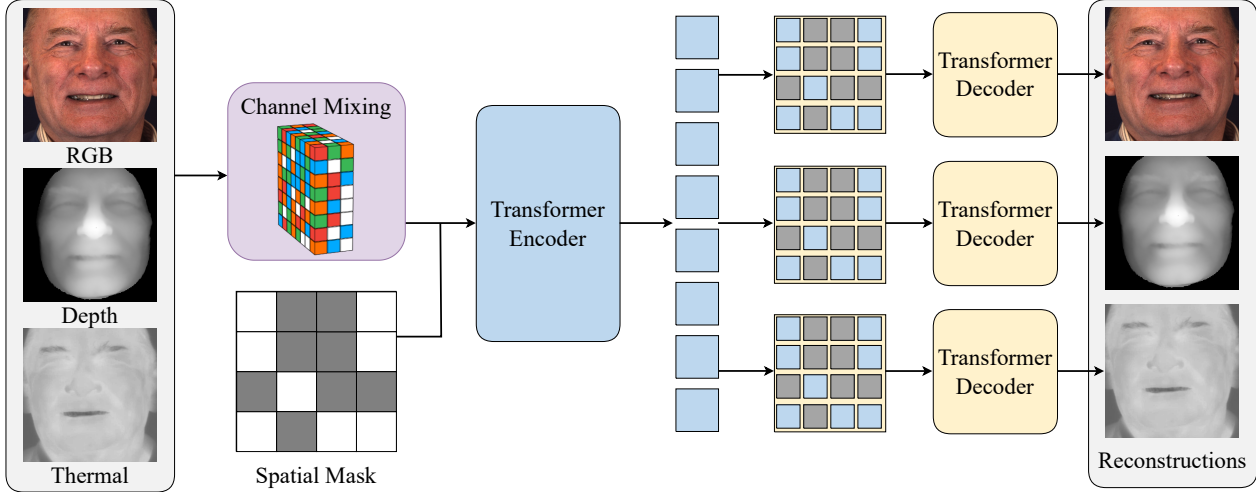


Figure 2. An overview of the proposed network, which is a multi-modal reconstruction network as a pre-train model for AU detection. In this model, 3 modalities are fused with 5 channels and randomly dropped 2 channels, resulting in 3 channels. Then a large subset of random patches is masked out, while the encoder is applied to the remaining unmasked patches. Three subsets of masked tokens combined with the shared visible encoder patches are proceeded by three small decoders to reconstruct three modalities respectively. Encoder and decoders are based on ViT and designed asymmetrical, i.e., 12, 8, 4, and 4 layers in the encoder, RGB decoder, depth decoder, and thermal decoder respectively.

$$w_l = F_{LN}^l(w_l' + F_{FF}^l(w_l')) \quad (2)$$

where $F_{MSA}(\cdot)$ is the multi-head self-attention, F_{LN} is LayerNorm, and F_{FF} is a fully connected feed-forward network. We apply the ViT-base with 12 blocks and 512-dim as the encoder in our model. Different from previous multi-modal AU detection works [41, 43], there's only one encoder for three modalities in the proposed network.

3.4. ViT Decoder

Since we have three modalities of reconstruction task, three decoders are involved in our network. The inputs of three decoders $De = \{de_v, de_d, de_t\}$ are consisting of the shared encoded visible patches e , and three modalities mask tokens $\{m_v, m_d, m_t\}$, so that we have $de_v = \{e, m_v\}$, $de_d = \{e, m_d\}$, and $de_t = \{e, m_t\}$, where v, d, t represent RGB, Depth, and Thermal. Each mask token is a shared embedding vector that represents a missing patch to be predicted in each decoder. Following the mask token design in [12], where removing the mask token from the encoder can not only improve the accuracy but also speed up the training process. The decoder position embedding is also added to the input, otherwise, it would have no information about their location in the image. The decoder has stacks of Transformer blocks, which are also defined as Equation.1 and Equation.2. The last layer of a decoder is a linear projection whose number of output channels equals the number of pixel values in a patch. The decoders are only used during pre-training to conduct the reconstruction tasks and

only the encoder is used later for multi-modal representation learning on AU detection.

In general, the last several layers in an autoencoder are more specialized for reconstruction but are less relevant for recognition [12]. The last layer of the encoder is adapted to the recognition task in fine-tuning, which means the decoder has less impact on improvement. Inspired by the asymmetric design of MAE, we also design different decoders for image, depth, and thermal due to the different complexity of reconstruction tasks, i.e., 12 layers of blocks in the encoder, 8, 4, and 4 layers of blocks in the decoder for visual RGB, depth, and thermal respectively.

3.5. Multi-modal Reconstruction

Our network reconstructs the shared latent feature by predicting the pixels for each masked patch on three modalities. Each element of the decoder output is a vector of pixel values representing a patch. For each modality, the loss function computes the mean squared error (MSE) between the masked patches of the reconstructed and original image. Define reconstruction losses \mathcal{L}_R^{rgb} for RGB, \mathcal{L}_R^{depth} for depth, and $\mathcal{L}_R^{thermal}$ for Thermal. Then we combine these losses from three modalities by,

$$\mathcal{L}_R = \mathcal{L}_R^{rgb} + \mathcal{L}_R^{depth} + \mathcal{L}_R^{thermal} \quad (3)$$

3.6. AU Detection

We do supervised training to evaluate the representations from the encoder with end-to-end fine-tuning on AU detection. Our model can accept mixed images as well as single

modal images, where input $x \in \{I_v, I_d, I_t, I_{mix}\}$. Similar to the MAE [12] we use the average pooling instead of the class token. As a multi-label task, AU detection faces the common problem of data unbalances. Hence, we apply the weighted binary cross-entropy (BCE) loss functions. The AU_k occurrence probability function $P(AU_k)$ and the correlate positive weight p_k of weighted BCE are given as:

$$P(AU_k) = \frac{\sum_{i=1}^N y_k^i}{\sum_{i=1}^N \sum_{k=1}^C y_k^i} \quad (4)$$

$$p_k = \frac{P(AU_k)}{\min\{P(AU_1), \dots, P(AU_C)\}} \quad (5)$$

where N is the total number of training set, C is the number of AUs. Then the weighted BCE loss is defined:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^C (p_k \times y_k^i \times \log(\hat{y}_k^i) + (1 - y_k^i) \times \log(1 - \hat{y}_k^i)) \quad (6)$$

where y_k^i and \hat{y}_k^i represent the ground truth label and prediction for AU_k respectively.

4. Experiment

Our pre-trained MCM is trained on the BP4D+ [45], where 300,000 random sample frames from all the frames are extracted (with/without AU notation). The pretraining needs three modalities for input, that is RGB, Depth, and Thermal. After the pretraining, we employ supervised training to finetune the encoder of AU detection on BP4D and DISFA [28]. Through our experiments, we aim to address the research questions with respect to the following four aspects: **RQ1:** *Can our proposed method outperform state-of-the-art baseline methods?* (section 4.4.1.); **RQ2:** *Will multiple modalities fusion be more effective?* (section 4.4.2); **RQ3:** *What are the image reconstruction results across the database?;* (section 4.5.1); **RQ4:** *How effective is MCM, and how efficient is it compared to traditional multimodal fusion methods?* (Section 4.5.2).

4.1. Data

BP4D [44] contains 41 subjects (23 females and 18 males) captured in laboratory environments. There are 8 tasks designed to elicit different spontaneous emotions. Expert coders select the most expressive 20s of each video clip for AU coding, resulting in 140,000 labeled frames. These labeled frames are used in AU detection fine-tuning. In the AU detection experiments, following the work [20], we split them into subject-exclusive 3-fold with 12 AUs. Besides the RGB images, the database provides the corresponding 3D face meshes as well, generating the Depth images to

use. **BP4D+** [45] consists of 140 subjects with a total of 1.5 M frames in the same laboratory environments. In **pre-training**, we equally sample 300,000 frames from all 1.5 M frames by 10 activity categories. Not only the RGB images, their corresponding Depth, and Thermal images are also used. **DISFA** [28] contains videos from the left view and right view of 27 subjects (12 females, 15 males). 12 AUs are labeled with AU intensity from 0 to 5, resulting in around 130,000 AU-coded images. We choose 8 of 12 AUs from the left camera with AU intensities higher or equal to 2 as positive samples. The recognition model trained on BP4D is then fine-tuned to the DISFA dataset, which is following the setting in [17,20]. F1-score is reported based on subject-exclusive 3-fold cross-validation.




Figure 3. An example of Homography Transform on Thermal images based on their bounding boxes.

4.2. Preprocessing

In our experiments, we use three modalities in BP4D+, 2D RGB, Depth, and Thermal images. The face regions of RGB images are detected and cropped by the publicly available library OpenCV. A raw 3D face contains 30k - 50k vertices in BP4D databases, including the face, hair, and neck areas of a subject. We first project the 3D mesh into depth images matching the visual image and then cropped the face area with the corresponding visual bounding box. Thermal images are normalized based on the original thermal temperature data. Since the RGB camera and Thermal camera are positioned differently, resulting in varying images, we apply *homography transform* to align the thermal images with their corresponding RGB images, which is shown in Figure. 3. For each subject, only one bounding box is necessary for RGB and one for thermal images. Prior to training, all the face images are resized to 244x244.

4.3. Implementation details

Pre-training. Following the standard ViT architecture, we use ViT-base as the encoder and decoders, where 12 layers are in the encoder, and 8, 4, and 4 layers in decoders for RGB, Depth, and Thermal, respectively. The imagenet pre-training weights from [12] are used to our encoder. Optimizer Adamw [26] with learning rate $2.e-4$, weight decay 0.05, and momentum (0.9, 0.95). We train 100 epochs with batch size 64 where the warmup epoch [11] is set with 40 and then perform cosine decay [25] with minimal lr by $1.e-$

Table 1. F1 scores in terms of 12 AUs are reported for the proposed method and the state-of-the-art methods based on single modality on the BP4D. Bold numbers indicate the best performance in each modality.

Methods	Modal	AU1	AU2	AU4	AU6	AU7	AU10	AU12	AU14	AU15	AU17	AU23	AU24	Avg
EAC	RGB	39.0	35.2	48.6	76.1	72.9	81.9	86.2	58.8	37.5	59.1	35.9	35.8	55.9
DSIN	RGB	51.7	40.4	56.0	76.1	73.5	79.9	85.4	62.7	37.3	62.9	38.8	41.6	58.9
JAA-Net	RGB	47.2	44.0	54.9	77.5	74.6	84.0	86.9	61.9	43.6	60.3	42.7	41.9	60.0
SRERL	RGB	46.9	45.3	55.6	77.1	78.4	83.5	87.6	63.9	52.2	63.9	47.1	53.3	62.9
HMP-PS	RGB	53.1	46.1	56.0	76.5	76.9	82.1	86.4	64.8	51.5	63.0	49.9	54.5	63.4
SEV-Net	RGB	58.2	50.4	58.3	81.9	73.9	87.8	87.5	61.6	52.6	62.2	44.6	47.6	63.9
FAUT	RGB	51.7	49.3	61.0	77.8	79.5	82.9	86.3	67.6	51.9	63.0	43.7	56.3	64.2
MFT	RGB	49.6	48.1	59.9	78.4	78.0	83.7	87.9	62.0	55.3	61.8	50.9	54.9	64.2
KSRL	RGB	53.3	47.4	56.2	79.4	80.7	85.1	89.0	67.4	55.9	61.9	48.5	49.0	64.5
MCM	RGB	54.4	48.5	60.6	79.1	77.0	84.0	89.1	61.7	59.3	64.7	53.0	60.5	66.0
ResNet-18	Depth	40.9	40.7	56.9	77.1	76.7	79.8	88.1	58.8	48.2	60.8	49.4	46.0	60.3
ResNet-50	Depth	44.5	41.3	57.2	76.5	72.6	80.9	86.8	54.3	50.2	62.7	48.1	44.5	60.0
MFT	Depth	38.6	37.3	44.2	84.2	89.0	89.7	89.2	79.8	44.7	46.0	53.2	37.0	61.1
MCM	Depth	44.3	40.8	51.3	77.8	70.9	82.9	87.3	65.5	50.8	63.1	48.1	54.5	61.4

Table 2. F1 scores in terms of 8 AUs are reported for the proposed method and the state-of-the-art methods on the DISFA dataset. Bold numbers indicate the best performance

Methods	Modal	AU1	AU2	AU4	AU6	AU9	AU12	AU25	AU26	Avg
EAC	RGB	41.5	26.4	66.4	50.7	80.5	89.3	88.9	15.6	48.5
DSIN	RGB	42.4	39.0	68.4	28.6	46.8	70.8	90.4	42.2	53.6
JAA-Net	RGB	43.7	46.2	56.0	41.4	44.7	69.6	88.3	58.4	56.0
SRERL	RGB	45.7	47.8	59.6	47.1	45.6	73.5	84.3	43.6	55.9
SEV-Net	RGB	55.3	53.1	61.5	53.6	38.2	71.6	95.7	41.5	58.8
HMP-PS	RGB	38.0	45.9	65.2	50.9	50.8	76.0	93.3	67.6	61.0
FAUT	RGB	46.1	48.6	72.8	56.7	50.0	72.1	90.8	55.4	61.5
KSRL	RGB	60.4	59.2	67.5	52.7	51.5	76.1	91.3	57.7	64.5
MCM	RGB	49.6	44.1	67.2	65.5	49.0	81.5	85.9	71.8	64.3

6. **Fine-tuning** The drop path [13] for ViT to 0.1 and layer-wise lr decay [3] is 0.75. We set the batch size to 64, train 5 epochs with no warmup and cosine decay in fine-tuning, and the reduction factor of lr is 10 at the 3rd epoch.

We implement our method with Pytorch-1.11.0 and cuda-11.3 on the GPU of NVIDIA RTX 3090. To evaluate the performance of AU detection accuracy, we use the F1-score for a comparative study with the state-of-the-art.

4.4. Comparison with related methods

4.4.1 Single-modal results (RQ1)

RGB. Our approach allows the use of a single modality in AU detection. Hence, we first compare our method to the state-of-the-art algorithms on an RGB modality in terms of the F1-score. Methods included on BP4D are EAC [20], DSIN [7], JAA-Net [33], SRERL [17], HMP-PS [35], SEV-Net [42], FAUT [15], MFT [43], and KSRL [5]. The top section in Table 1 shows the results of visual RGB images on BP4D, where our method achieves the best performance. Specifically, our method outperforms the SOTA methods in recognizing AU12, AU15, AU17, AU23, AU24.

Depth. To evaluate our method on Depth modality, we

compare with [43] and Resnet-18, Resnet-50, as the baseline functions, where our method outperforms the state of the art. Since depth has only one channel, it’s duplicated into 3 channels as the input when using Depth only. The result is reported in the bottom section of Table.1.

Table 2 shows the results on the DISFA dataset. The performance is evaluated on 8 action units, in line with prior works [5, 7, 15, 17, 20, 33, 35, 42]. For DISFA, following the protocol in [17, 20], we do not train the model from scratch, but fine-tune the model from training on BP4D. Our method achieves the state-of-the-art where only slightly less than KSRL [5].

4.4.2 Multi-modal results (RQ2)

Since our method also accepts a fused channel-mixing feature as the input, we compare our method with the state-of-the-art multi-modal methods, i.e., MTUT [1], TEMT-Net [24], AMF [41], and MFT [43]. The F1-score of MTUT and TEMT-Net are reported in the work [41], where the author implemented them with the fusion of RGB and Depth. The results on BP4D are shown in Table 3. Our method

Table 3. F1 scores in terms of 12 AUs are reported for the proposed method and the state-of-the-art methods based on RGB+Depth modalities on BP4D. Bold numbers indicate the best performance

Methods	Modal	AU1	AU2	AU4	AU6	AU7	AU10	AU12	AU14	AU15	AU17	AU23	AU24	Avg
MTUT	RGBD	51.3	50.2	62.2	77.2	71.7	83.8	88.2	61.4	54.3	57.9	45.8	42.2	62.2
TEMT-Net	RGBD	53.7	47.1	60.5	77.6	75.6	84.8	87.4	67.0	57.2	61.3	44.7	41.6	63.2
AMF	RGBD	55.1	58.3	62.0	82.5	75.6	87.2	89.6	60.9	59.1	62.4	45.0	52.0	65.8
MFT	RGBD	51.6	49.2	57.6	78.8	77.5	84.4	87.9	65.0	56.5	64.3	49.8	55.1	64.8
MCM	RGBD	56.5	47.4	62.2	79.9	78.2	84.3	87.9	67.8	55.8	64.8	52.3	57.2	66.2




Figure 4. Visualization of the reconstructed images on BP4D with different mask ratios is performed using a reconstruction model trained on BP4D+. Images in (b, c, and d) represent the masked channel-mixing image, reconstructed RGB, reconstructed Depth, and reconstructed Thermal, respectively.

outperforms all the related methods and achieves the highest F1-score at 66.2%. In terms of 12 AUs, our method performs best in 7 AUs: *AU1, AU4, AU7, AU14, AU17, AU23* and *AU24*. The fusion feature surpasses the performance of all single-modality methods, which validates the effectiveness of utilizing multiple modalities.




Figure 5. Visualization of the reconstructions on DISFA. The examples are reconstructed from the DISFA dataset using a pre-trained model on the BP4D+ and a mask ratio of 50%.

4.5. Ablation Studies

4.5.1 Image reconstructions (RQ3)

Figure 4 shows reconstructions result on BP4D by different mask ratios, i.e., 50%, 75% and 90%. Inputs with RGB and Depth through the channel-mixing resulting the final reconstruction result in RGB, Depth, and Thermal. The reconstructions are generated by the pre-trained model with 50% masks. As BP4D+ was using similar recording cameras to BP4D, our proposed method works well for the multi-modalities reconstruction task. Although all the identities in BP4D are unseen to the model trained in BP4D+, it successfully reconstructs the facial images of individuals using only a few visible patches. As the mask ratio increases, the quality of the reconstructed images tends to decrease, particularly with regards to preserving identity and expression. Higher mask ratios introduce more information loss, leading to more significant changes in the reconstructed images.

In the case of DISFA, where only RGB modality is available, our pre-training process utilizes only RGB images as the input. This also demonstrates the flexibility and adaptability of our method in handling different modalities across different datasets. The result is shown in Figure 5. The RGB image reconstruction on DISFA is not as good as on BP4D, which is caused by the following reason. The image color in DISFA is biased, i.e., the distribution of the im-

Table 4. Ablation study of the effectiveness of proposed Channel-Mixing module on BP4D database. IM: input modality in detection; CM: Channel-Mixing modalities; V: Visual RGB; D: Depth; T: Thermal. F1 scores of 12 AUs are reported for different variants on the BP4D.

Stage	IM	CM	AU1	AU2	AU4	AU6	AU7	AU10	AU12	AU14	AU15	AU17	AU23	AU24	Avg
Pretrain	V	✗	53.1	45.4	56.3	80.2	74.8	84.1	88.5	67.1	52.6	67.2	53.4	57.8	65.0
	V	VD	53.4	48.5	57.8	79.8	76.7	83.1	88.5	64.2	54.9	64.5	53.5	61.2	65.5
	V	VDT	54.4	48.5	60.6	79.1	77.0	84.0	89.1	61.7	59.3	64.7	53.0	60.5	66.0
	D	✗	44.5	41.4	51.0	77.5	69.8	82.1	87.2	62.1	48.3	65.7	48.7	53.1	61.0
	D	VD	44.3	40.8	51.3	77.8	70.9	82.9	87.3	65.5	50.8	63.1	48.1	54.5	61.4
	D	VDT	44.0	39.2	50.4	77.0	71.1	82.9	87.6	64.4	52.4	62.6	48.6	55.6	61.3
	VD	✗	54.6	49.1	58.4	79.7	76.0	83.7	88.6	61.9	55.6	64.6	52.6	60.3	65.4
	VD	VD	56.5	47.4	62.2	79.9	78.2	84.3	87.9	67.8	55.8	64.8	52.3	57.2	66.2
	VD	VDT	56.9	47.9	59.4	79.6	76.1	83.4	89.2	62.0	59.6	63.8	52.8	59.2	65.8
Finetune	I	✗	43.9	42.9	51.1	78.9	77.1	83.2	87.6	69.6	51.9	60.7	46.2	58.8	62.7
	D	✗	40.6	36.7	50.8	78.8	72.1	81.8	87.8	68.1	50.0	61.7	47.0	55.3	60.9
	VD	✗	44.0	42.7	48.5	79.0	76.7	83.2	87.8	69.0	52.3	63.1	48.8	54.7	62.5
	VD	✓	47.9	45.8	48.9	78.9	77.6	84.2	88.0	68.4	52.0	61.6	46.8	55.9	63.0

age ranges from warm to cool colors, which is much more deviated than BP4D+. See the enormous difference in the ground truth color between rows 1 and 3 in Figure 5. In such a case, the reconstruction model learned from BP4D+ is quite difficult to match all the colors on DISFA. The reconstructed images are quite warm, so the color from some examples is unmatched. However, the semantic information of the face, e.g., eyes, nose, and mouth, is well reconstructed. Meanwhile, due to the missing Depth modality, its Depth reconstruction result is also worse than BP4D. Considering there’s no same identity person across BP4D, BP4D+, and DISFA, our framework has a good quality reconstruction across the databases.

4.5.2 Effectiveness and Efficiency of CM (RQ4)

Effectiveness. We conducted comparison experiments on the BP4D dataset to carefully analyze the contribution of our proposed Channel-Mixing (CM) module, which was utilized in both the pre-training and fine-tuning stages. Table 4 shows the performance of various comparisons in terms of the F1-score of 12 AUs. We first examine the CM module in the pre-training stage, where we report the detection result of all modalities fine-tuning on different pre-trained multi-modal models. MAE without the CM module is used as the baseline function. Applying the proposed channel-mixing achieves a noticeable increase at 1.5%, 0.7%, and 1.2% in RGB, Depth, and RGBD respectively, leading to an improvement over the SOTA. To verify the efficiency of the CM module itself only in the fine-tuning stage, we then use a pre-trained ViT to evaluate. See the results in the Finetune section of Table 4. A 1x1 convolutional layer is used to reduce the 4-channel RGBD to 3 channels for embedding. Performance in RGB achieves the best by this baseline model, which is better than multimodal early fusion. However, by directly utilizing the CM mod-

ule, the RGBD early fusion outperforms the RGB modality. Comparing the two stages, MAE in pre-training improves more clearly. In the pre-training, channel mixed features are reconstructed through the encoder and decoder, effectively utilizing existing channel information to generate the missing channels. However, the mixing feature is fed simply to a ViT for AU pattern learning in fine-tuning.

Efficiency. The total parameter count of the pre-trained model is 125M, with 85M parameters located in the encoder. In comparison to the traditional three-encoder late fusion method, our proposed approach efficiently reduces approximately 57% (170M) of trainable parameters.

5. Conclusion

This paper has proposed a novel multi-modal learning approach that utilizes RGB, depth, and thermal modalities to train a reconstruction network on pre-training. A channel-mixing module empowered our method to process either a single modality or multi-modalities using a single-stream encoder. The proposed method for AU detection has surpassed or reached the state-of-the-art. Additionally, the experiments demonstrate the effectiveness of multi-modal fusion, which outperforms single-modality approaches. Ablation experiments further validate the efficiency and robustness of the proposed channel mixing module. **Limitation.** This approach has limitations in multi-modal, requiring that the modality be consistent in space and time. Feature-level Channel-Mixing will be investigated in the future.

6. Acknowledgment

This work is supported in part by NSF (CNS-1629898) and the Center of Imaging, Acoustics, and Perception Science (CIAPS) of the Research Foundation of Binghamton University.

References

- [1] Mahdi Abavisani, Hamid Reza Vaezi Joze, and Vishal M Patel. Improving the performance of unimodal dynamic hand-gesture recognition with multimodal training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1165–1174, 2019. 6
- [2] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41:423–443, 2018. 3
- [3] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEit: BERT pre-training of image transformers. In *International Conference on Learning Representations*, 2022. 3, 6
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 3
- [5] Yanan Chang and Shangfei Wang. Knowledge-driven self-supervised representation learning for facial action unit recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20417–20426, 2022. 2, 6
- [6] Xinlei Chen*, Saining Xie*, and Kaiming He. An empirical study of training self-supervised vision transformers. *arXiv preprint arXiv:2104.02057*, 2021. 3
- [7] Ciprian Corneanu, Meysam Madadi, and Sergio Escalera. Deep structure inference network for facial action unit recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 309–324. Springer International Publishing, 2018. 1, 6
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, 2019. 3
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 3
- [10] Paul Ekman and Erika L Rosenberg, editors. *What the face reveals: Basic and applied studies of spontaneous expression using the Facial Action Coding System (FACS)*. Oxford University Press, 1997. 1
- [11] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour, 2017. 5
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners, 2021. 2, 3, 4, 5
- [13] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. Deep networks with stochastic depth. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 646–661. Springer International Publishing, 2016. 6
- [14] Ramin Irani, Kamal Nasrollahi, Marc O Simon, Ciprian A Corneanu, Sergio Escalera, Chris Bahnsen, Dennis H Lundtoft, Thomas B Moeslund, Tanja L Pedersen, Maria-Louise Klitgaard, et al. Spatiotemporal analysis of rgb-dt facial images for multimodal pain level recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 88–95, 2015. 2
- [15] Geethu Miriam Jacob and Bjorn Stenger. Facial action unit detection with transformers. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2021. 1, 2, 6
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 2
- [17] Guanbin Li, Xin Zhu, Yirui Zeng, Qing Wang, and Liang Lin. Semantic relationships guided representation learning for facial action unit recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8594–8601, 2019. 1, 5, 6
- [18] Huibin Li, Huaxiong Ding, Di Huang, Yunhong Wang, Xi Zhao, Jean-Marie Morvan, and Liming Chen. An efficient multimodal 2d+ 3d feature-based approach to automatic facial expression recognition. *Computer Vision and Image Understanding*, 140:83–92, 2015. 2
- [19] Huibin Li, Jian Sun, Zongben Xu, and Liming Chen. Multimodal 2d+ 3d facial expression recognition with deep fusion convolutional neural network. *IEEE Transactions on Multimedia*, 19(12):2816–2831, 2017. 2
- [20] Wei Li, Farnaz Abtahi, Zhigang Zhu, and Lijun Yin. EACnet: A region-based deep enhancing and cropping approach for facial action unit detection. In *IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, 2017. 1, 2, 5, 6
- [21] Xiaotian Li, Zhihua Li, Huiyuan Yang, Geran Zhao, and Lijun Yin. Your “attention” deserves attention: A self-diversified multi-channel attention for facial action analysis. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, 2021. 2
- [22] Xiaotian Li, Xiang Zhang, Huiyuan Yang, Wenna Duan, Weiyang Dai, and Lijun Yin. An eeg-based multi-modal emotion database with both posed and authentic facial actions for emotion analysis. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, 2020. 2, 3
- [23] Yong Li, Jiabei Zeng, Shiguang Shan, and Xilin Chen. Self-supervised representation learning from videos for facial action unit detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2019. 1
- [24] Peng Liu, Zheng Zhang, Huiyuan Yang, and Lijun Yin. Multi-modality empowered network for facial action unit detection. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2175–2184. IEEE, 2019. 2, 6
- [25] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts, 2016. 5

- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5
- [27] Liupei Lu, Leili Tavabi, and Mohammad Soleymani. Self-supervised learning for facial action unit recognition through temporal consistency. In *BMVC*, 2020. 1
- [28] S. Mohammad Mavadati, Mohammad H. Mahoor, Kevin Bartlett, Philip Trinh, and Jeffrey F. Cohn. DISFA: A spontaneous facial action intensity database. *IEEE Transactions on Affective Computing*, 4(2):151–160, 2013. 5
- [29] Xuesong Niu, Hu Han, Songfan Yang, Yan Huang, and Shiguang Shan. Local relationship learning with person-specific shape regularization for facial action unit detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [30] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7077–7087, 2021. 3
- [31] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018. 3
- [32] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 3
- [33] Zhiwen Shao, Zhilei Liu, Jianfei Cai, and Lizhuang Ma. Deep adaptive attention for joint facial action unit detection and face alignment. In *Proceedings of the European conference on computer vision (ECCV)*, pages 705–720, 2018. 1, 2, 6
- [34] Hangke Song, Zhi Liu, Huan Du, Guangling Sun, Olivier Le Meur, and Tongwei Ren. Depth-aware salient object detection and segmentation via multiscale discriminative saliency fusion and bootstrap learning. *IEEE Transactions on Image Processing*, 26(9):4204–4216, 2017. 2, 3
- [35] Tengfei Song, Zijun Cui, Wenming Zheng, and Qiang Ji. Hybrid message passing with performance-driven structures for facial action unit detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6267–6276, June 2021. 6
- [36] Enrique Sánchez-Lozano, Georgios Tzimiropoulos, and Michel Valstar. Joint action unit localisation and intensity estimation through heatmap regression. In *BMVC*, 2018. 2
- [37] Yang Tang, Wangding Zeng, Dafei Zhao, and Honggang Zhang. Piap-df: Pixel-interested and anti person-specific facial action unit detection net with discrete feedback learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [38] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2019, page 6558. NIH Public Access, 2019. 3
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2, 3
- [40] Shangfei Wang, Zhilei Liu, Zhaoyu Wang, Guobing Wu, Peijia Shen, Shan He, and Xufa Wang. Analyses of a multi-modal spontaneous facial expression database. *IEEE Transactions on Affective Computing*, 4(1):34–46, 2012. 2
- [41] Huiyuan Yang, Taoyue Wang, and Lijun Yin. Adaptive multimodal fusion for facial action units recognition. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2982–2990, 2020. 2, 4, 6
- [42] Huiyuan Yang, Lijun Yin, Yi Zhou, and Jiuxiang Gu. Exploiting semantic embedding and visual feature for facial action unit detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10482–10491, 2021. 1, 2, 6
- [43] Xiang Zhang and Lijun Yin. Multi-modal learning for AU detection based on multi-head fused transformers. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*. IEEE, 2021. 2, 3, 4, 6
- [44] Xing Zhang, Lijun Yin, Jeffrey F Cohn, Shaun Canavan, Michael Reale, Andy Horowitz, Peng Liu, and Jeffrey M Girard. Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image and Vision Computing*, 32(10):692–706, 2014. 2, 5
- [45] Zheng Zhang, Jeff M Girard, Yue Wu, Xing Zhang, Lijun Yin, et al. Multimodal spontaneous emotion corpus for human behavior analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3438–3446, 2016. 2, 5
- [46] Kaili Zhao, Wen-Sheng Chu, Fernando De la Torre, Jeffrey F. Cohn, and Honggang Zhang. Joint patch and multi-label learning for facial action unit detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [47] Xiaoqi Zhao, Lihe Zhang, Youwei Pang, Huchuan Lu, and Lei Zhang. A single stream network for robust and real-time RGB-d salient object detection. In *Computer Vision – ECCV 2020*, pages 646–662. Springer International Publishing, 2020. 2, 3