

# Most Important Person-guided Dual-branch Cross-Patch Attention for Group Affect Recognition

Hongxia Xie<sup>1</sup>, Ming-Xian Lee<sup>2</sup>, Tzu-Jui Chen<sup>2</sup>, Hung-Jen Chen<sup>1</sup>, Hou-I Liu<sup>1</sup>, Hong-Han Shuai<sup>1</sup> and Wen-Huang Cheng<sup>3</sup>

<sup>1</sup>National Yang Ming Chiao Tung University, hongxiaxie.ee08@nycu.edu.tw

<sup>2</sup>University of Illinois Urbana-Champaign

<sup>3</sup>National Taiwan University, wenhuang@csie.ntu.edu.tw

## Abstract

Group affect refers to the subjective emotion that is evoked by an external stimulus in a group, which is an important factor that shapes group behavior and outcomes. Recognizing group affect involves identifying important individuals and salient objects among a crowd that can evoke emotions. However, most existing methods lack attention to affective meaning in group dynamics and fail to account for the contextual relevance of faces and objects in group-level images. In this work, we propose a solution by incorporating the psychological concept of the Most Important Person (MIP), which represents the most noteworthy face in a crowd and has affective semantic meaning. We present the Dual-branch Cross-Patch Attention Transformer (DCAT) which uses global image and MIP together as inputs. Specifically, we first learn the informative facial regions produced by the MIP and the global context separately. Then, the Cross-Patch Attention module is proposed to fuse the features of MIP and global context together to complement each other. Our proposed method outperforms state-of-the-art methods on GAF 3.0, GroupEmoW, and HECO datasets. Moreover, we demonstrate the potential for broader applications by showing that our proposed model can be transferred to another group affect task, group cohesion, and achieve comparable results.

## 1. Introduction

Humans are active and social creatures, using multi-modal interactions to convey their intentions, attitudes, and feelings. Such physical and emotional interactions between individuals within a group can generate group-level affect

Global +	All	3Random	3MIPs	1Random	MIP
Acc (%)	85.01	85.47	87.46	82.30	<b>88.86</b>

Table 1. Different face selection strategies. The “3Random” refers to selecting 3 faces randomly from the detected faces in the image, while the “3MIPs” strategy selects the top three detected faces based on their importance ranking. The dataset evaluated here is the GroupEmoW [22].

or group-level emotion [2]<sup>1</sup>.

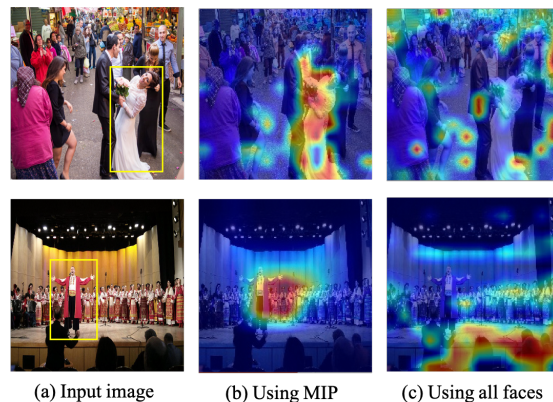


Figure 1. Feature map comparison between using MIP and all faces. The yellow box in (a) represents the MIP region.

By providing the group-level information, group affect prediction has various real-world applications, *e.g.*, work team outcome prediction [21], social relationship recognition [33], human-machine interaction systems [24]. As a result, the research topics surrounding group affect are di-

<sup>1</sup>In this work, groups range in size from two or three individuals up to hundreds.

verse, including categorical analysis, *e.g.*, group valence (positive, negative, neutral) prediction [16, 22], continuous intensity estimation, *e.g.*, group cohesion prediction [15].

Group affect is influenced by a combination of a group’s *affective context* (*e.g.*, salient objects in the funeral, party) and *affective composition* (*i.e.*, the combination of a group members’ state and trait affect) [2]. Previous methods for group affect recognition mainly adopted a bottom-up approach [9, 44, 16, 22], where features were extracted from individuals separately and then fused to output group-level emotion. This was done by utilizing multi-cues from separate pre-trained detectors, such as facial expressions of individuals, object proposals, and scene features. While these cues can provide valuable information, they may not fully capture the complex interplay of emotions within a group. As indicated in [12], current group-level recognition architectures tend to focus on detecting salient regions, but they may lack an understanding of the affective meaning in the image. Therefore, there is a need for a more holistic approach that considers the most important region in the crowd and how they interact with the scene.

Hence, in this work, we introduce a crucial psychological concept, namely, **Most Important Person (MIP)** [2], which goes beyond saliency and considers the importance of individuals for group affect recognition. The MIP of an image is often the group leaders who can influence the emotion of a group [11, 37, 34]. MIP has been explored in many related works in psychology and affective computing, such as the *Affective Transfer Process* phenomenon in psychology [2], which indicates that the MIP of an image can influence the group’s emotion. We conducted a preliminary experiment based on the CrossVit [5] model to demonstrate the effectiveness of the MIP in group affect. Our architecture uses both the global image and faces as two inputs. As depicted in Table 1, the results demonstrate that using all faces is considerably less effective than using the MIP, and even less effective than randomly selecting three facial images. Additionally, as observed in the feature maps illustrated in Figure 1, the model that employs all faces fails to predict accurately due to its focus on irrelevant background regions or faces. In contrast, the network that utilizes MIP can efficiently focus on relevant facial regions, thereby enhancing prediction accuracy. This is mainly due to the fact that in crowded scenes with many blurry faces, the model trained on all faces exhibits dispersed attention, while the model trained on the MIP region concentrates specifically on the relevant area.

Based on the above observations, we present a dual-pathway vision transformer model, the Dual-branch Cross-Patch Attention Transformer (DCAT), including a global branch and a MIP branch. With the dual-pathway, spatial-wise discriminative clues can be discovered both globally and locally in conjunction with self-attention learning. For

the dual-pathway interaction, since there are many interferences in the group, we propose the Cross-Patch Attention (CPA) module. CPA first utilizes the Token Ranking Module to select the important tokens in each path, and then the cross-attention is calculated between the selected query vectors and the entire key-value vectors from the other path. In this way, global and MIP contexts are able to complement and compensate for each other.

This paper makes the following main contributions:

- In this work, we validate that MIP plays a crucial role in group affect recognition. To the best of our knowledge, this is the first work introducing MIP into the group affect task.
- The MIP and global affective context information are integrated into the proposed dual-pathway vision transformer architecture for recognizing group emotions in a context-aware manner. We also propose a novel CPA module that focuses on the interactions among diverse scene contexts and facial informative regions.
- Experimental results show that the proposed DCAT outperforms both state-of-the-art group affect works and vision transformer models in terms of accuracy and model parameters. This is also the first work to explore group affect recognition that goes beyond the CNN-based approach. Moreover, our proposed model can be utilized in other group-level affect tasks, *i.e.*, group cohesion analysis.

## 2. Related works

### 2.1. Group affect recognition methods

Current group affect recognition methods primarily analyze the individual members of a group and then assess their contribution to the overall mood of the group by taking into account a variety of cues, such as faces, salient objects, and scenes. Fujii *et al.* [9] proposed a two-stage classification method, in which the first stage classifies facial expressions, and the second stage adds scene features into consideration to fuse faces and scenes through spectral clustering for further group affect recognition. Khan *et al.* [22] considered multi-cues, which include all faces, object proposals, and scene information. To achieve such an architecture, additional face detection and object proposal detection detectors are needed. Similarly, Guo *et al.* [16] also investigated multi-cues using graph neural networks, where the topology of the graph varies as the number of entities varies.

However, in large groups, it is not feasible to consider the expressions of all faces and objects individually before aggregating, resulting in lower efficiency and accuracy. As indicated in [9], considering the performance of the main subject estimation is needed, while our work fills the gap.

## 2.2. Most important person (MIP) detection

### 2.2.1 MIP detection

A MIP detection model understands a high-level pattern in a social event image to identify the most important person. Ramanathan *et al.* [35] used an RNN to learn time-varying attention weights that can be used for event detection and classification. Without any explicit annotations, the person with the highest attendance can be identified as the MIP. In order to overcome the inherent variability of human poses, PersonRank [27] produced a multiple Hyper-Interaction Graph, which treats each individual person as a node. The most active node can be identified using four types of edge message functions. POINT [26] proposed two types of interaction modules, *i.e.*, the person-person interaction module and the event-person interaction module. It then aggregated relation features and person features to form the important feature. To reduce the annotation effort, Hong *et al.* [19] formulated a POINT-based iterative learning method for semi-supervised important people detection. Thus, we incorporate people awareness into global-level affect analysis in this work.

### 2.3. Dual attention learning models

Visual attention with dual-pathway has been widely used to find important contextual regions. Based on the outstanding performance of vision transformers (ViTs) in various visual tasks [8], there are many self-attention-based dual-pathway networks in recent years. Zhu *et al.* [49] proposed a global-local cross-attention to reinforce the spatial-wise discriminative clues for fine-grained recognition. A pairwise cross-attention was designed to establish the interactions between image pairs. Dual-ViT [46] incorporated a compressed semantic pathway that served as prior information in learning finer pixel-level details. The semantic pathway and pixel pathway are integrated together and trained to spread enhanced self-attention information in parallel.

The most similar architecture to our work is the **Cross-ViT** [5], which was proposed to learn multi-scale features effectively. However, we differ from it in two main ways. Firstly, our network divides at different scales using global images and MIP images, whereas CrossViT divides at different scales using the same input image. Another difference is the cross-attention mechanism. The CrossViT exchanges class tokens, while we designed a Cross-Patch Attention module based on attentive queries from both paths.

## 3. Proposed Dual-branch Cross-Patch Attention Transformer (DCAT)

### 3.1. Overview

The architecture of our proposed Dual-branch Cross-Patch Attention Transformer (DCAT) for group affect

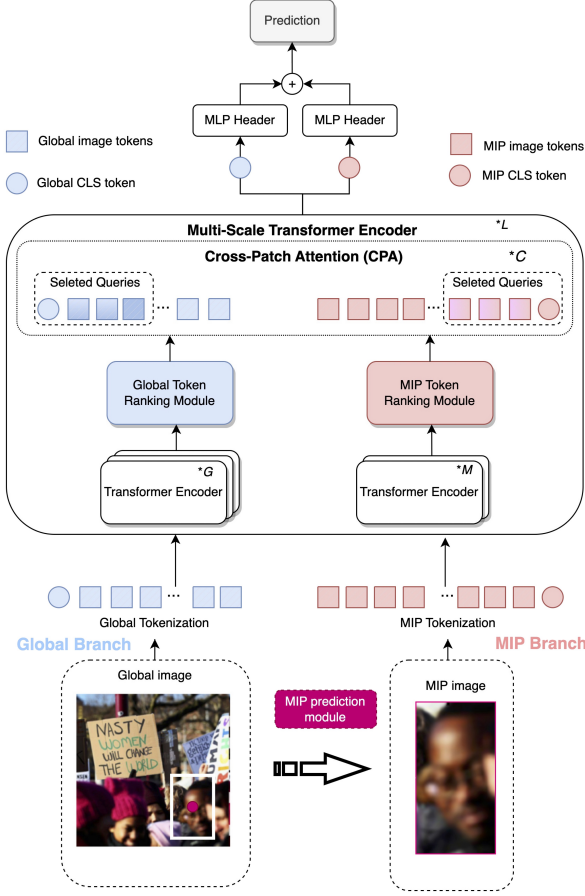
recognition is shown in Figure 2. The model contains two inputs, *i.e.*, the global image and the corresponding MIP image. The two images are tokenized into patches first and sent into a stack of Multi-Scale Transformer Encoder, which consists of dual-path Transformer Encoder, Token Ranking Module, and Cross-Patch Attention (CPA). Specifically, in a dual-path Transformer Encoder, the global image is sent to the coarse-level path to capture long-range information and further refine the encoded tokens for fine-grained details. On the other hand, the MIP image is fed to the fine-level path to obtain high-level semantic tokens. Next, since the global and MIP features are conceptual interdependent, we incorporate dual-path interaction in our model. In particular, we sort tokens according to their importance in each path based on the Token Ranking Module, then calculate the CPA using the key and value vectors from the other path. Finally, the class tokens from two branches are combined and sent into a linear layer for prediction.

### 3.2. Dual-pathway learning

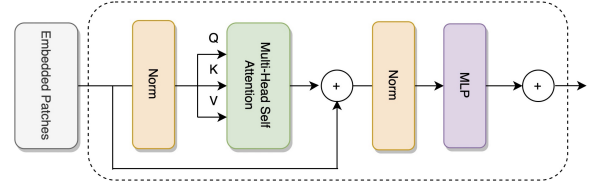
To leverage the global information of a group captured by the global image and the local-based semantic prior of the MIP, we propose a dual-pathway feature learning approach for group affect prediction. By integrating self-attention learning into the dual-pathway, our approach enables the discovery of spatial-wise discriminative clues both globally and locally. To validate our assumptions, we compute the CKA similarity [43] to measure the similarity of the intermediate token features in each layer and the final CLS token, following [43]. As illustrated in Figure 3, both the MIP and global tokens exhibit a positive correlation with the CLS token as the model depth increases. This suggests that both features are valuable for the final classification task and highlights the rationality and necessity of our dual-pathway input design.

Firstly, both the global and MIP image are tokenized into patches, then the class token and a learnable position embedding are added to both branches before the Multi-Scale Transformer Encoder. In order to balance computational costs, our Transformer Encoders in the two pathways include different numbers (*i.e.*,  $G=6$  and  $M=1$ ) inspired by CrossViT [5]. The global branch is the large (primary) branch with a coarse patch size (*i.e.*, 12), with a larger embedding size, and more transformer encoders. The MIP image is the input of a small (complementary) branch with fine-grained patch size (*i.e.*, 16), fewer encoders, and a smaller embedding size.

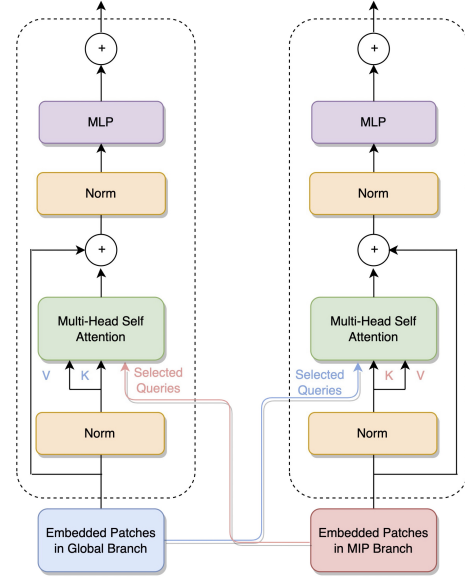
As the core of the vision transformer, the Multi-Head Self-Attention (MHSA) maps a query vector into a set of key and value vectors [8]. In the  $l$ -th Transformer Encoder block in both pathways (a basic Transformer Encoder is shown in Figure 2b), the output feature map is  $X_l \in \mathbb{R}^{(N+1) \times d}$ . The attention matrix  $A_l \in \mathbb{R}^{S \times (N+1) \times (N+1)}$



(a) The overall architecture of our proposed DCAT



(b) The Transformer Encoder



(c) The Cross-Patch Attention (CPA)

Figure 2. The overall architecture of our proposed Dual-branch Cross-Patch Attention Transformer (DCAT).

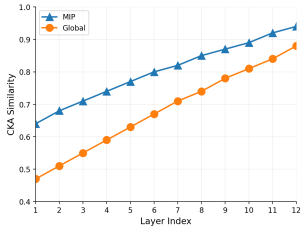


Figure 3. CKA similarity in each layer.

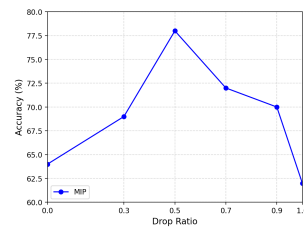


Figure 4. Different drop ratio of the MIP branch.

of MHSA module in the block is computed as:

$$A_l = \text{Softmax}\left(\frac{Q_l * K_l^T}{\sqrt{d}}\right), \quad (1)$$

where  $Q_l \in \mathbb{R}^{(N+1) \times d}$  and  $K_l \in \mathbb{R}^{(N+1) \times d}$  denote the queries and keys projected by  $X_l$  of self-attention operation in  $(l-1)$ -th transformer block, respectively.  $S$  represents the number of heads,  $T$  is a transpose operator,  $d$  is the di-

mension of the embedding.  $N$  is the number of tokenized patches. The input global and MIP images are divided into different patches numbers with different  $N$ . The calculation includes 1 class token in addition to the  $N$  tokens.

### 3.3. The interaction of global and MIP pathways

Due to the large amount of information in group-level images, extracting features by either global or MIP branch alone is not sufficient, and cross-attention of discriminative tokens from both can provide complementary effects.

Having obtained global and MIP tokens by the Transformer Encoder, we should consider the dual-path interactions. To achieve this, we propose the Cross-Patch Attention (CPA) mechanism to establish spatial interactions between the global and MIP images. Firstly, we sort the tokens based on their importance in each path using the Token Ranking Module and create a new query matrix by selecting the top  $\eta$  tokens. As depicted in Figure 4, the model's per-

formance, taking the MIP branch as an example<sup>2</sup>, is related to the drop ratio, where dropping a certain percentage of tokens can improve performance. Secondly, we compute the Cross-Patch Attention (CPA) using the new query matrix and key-value from the other branch. It is noteworthy that the gradients propagate through both pathways, which allows it to compensate for information loss on global feature compression through global-to-local interaction. As validated in [49], such cross-attention can also be viewed as a novel regularization method to regularize attention learning, especially for our small datasets.

### 3.3.1 Token Ranking Module

For Cross-Patch Attention, we need to remove unimportant contents from the complex group images and keep only the tokens that are relevant.

Our first step is to rank the tokens according to their importance. While the class token attention map  $A_{cls} \in \mathbb{R}^{1 \times (N+1)}$  reflects the importance of features [43, 29], we denote the token importance by the similarity scores between the global class token and each patch token:

$$A_{cls} = \text{Softmax}\left(\frac{q_{cls} * K^T}{\sqrt{d}}\right), \quad (2)$$

where  $q_{cls} \in \mathbb{R}^{1 \times d}$  is the query vector of the class token,  $K \in \mathbb{R}^{(N+1) \times d}$  is the key vector, and  $d$  is the dimension of the embedding.

According to the equation 2, the weights of each patch correlated with the class token are computed. In such a way, the  $A_{cls}$  reveals how much each patch contributes to the final classification since it captures the global interactions of the class token to all patches.

### 3.3.2 Cross-Patch Attention (CPA)

Based on the importance score of each token, we then construct a newly selected query matrix  $Q_{sel} \in \mathbb{R}^{(\alpha+1) \times d}$ , representing the attentive local embedding, by selecting the top  $\alpha$  query vectors that correspond to the top  $\alpha$  highest responses in the class attention map.

In each pathway, the Cross-Patch Attention (CPA) vectors is computed as below,

$$f_{CPA}(Q, K, V) = \text{Softmax}\left(\frac{Q_{sel} * K^T}{\sqrt{d}}\right)V, \quad (3)$$

where  $K \in \mathbb{R}^{(N+1) \times d}$  is the key vector,  $V \in \mathbb{R}^{(N+1) \times d}$  is the value vector and  $d$  is the dimension of the embedding.

As shown in Figure 2c, the CPA operation is bidirectional, with each of the two branches providing the selected tokens as the query vector. Both branches are fused together

<sup>2</sup>The global branch exhibits similar behavior.

$C$  times, which means the selected queries are updated in each CPA.

Finally, after repeating  $L$  Multi-Scale Transformer Encoder, the class tokens from two branches are combined and sent into a linear layer for prediction. The cross-entropy loss is used here for final classification [5, 8].

## 4. Experiments

### 4.1. Implemental details

Our model is partly pre-trained with the weight of CrossViT [5] on ImageNet. We train all our models for 300 epochs (30 warm-up epochs) on 2 GPUs (RTX 3090 Ti) with a batch size of 64. The numbers of Transformer Encoder in the dual path, CPA, and Multi-Scale Transformer Encoder are set as  $M=1$ ,  $G=6$ ,  $C=3$ , and  $L=1$ . For the token selection ratios  $\alpha$  of (global, MIP), we use (0.5, 0.5) for GAF 3.0 and (1, 0.5) for GroupEmoW throughout the paper. More details can be found in our supplementary materials.

### 4.2. Datasets

In this paper, we utilize several datasets for our experiments, including GroupEmoW [16], GAF 3.0 [13], GAF Cohesion [13, 15], and HECO [44], as presented in Table 2. The GroupEmoW dataset contains images obtained from the web via search engines, using keywords related to various social events such as funerals, birthdays, protests, conferences, and meetings. GAF 3.0 [13] and GAF Cohesion [13, 15] datasets were created through web crawling using a diverse set of keywords related to social events, including world cup winners, weddings, family gatherings, laughing clubs, birthday celebrations, siblings, riots, protests, and violence. The cohesion scores in GAF Cohesion range from *strongly agree*, *agree*, *disagree* to *strongly disagree*, using a 0-3 scale.

In our experiments, we utilized the MIP ground truth provided by HECO [44]. However, the ground truth for GAF 3.0 and GroupEmoW was not available. However, since the ground truth for GAF 3.0 and GroupEmoW was not available, we employed the pre-trained POINT model [26] to generate the MIP. The POINT model was trained on the Multi-scene Important People Image Dataset (MS Dataset).<sup>3</sup>

### 4.3. Comparison with state-of-the-art group affect methods and ViTs

In this work, we have built our proposed modules based on CrossViT [5] with vanilla Transformer Encoders [8], which denotes by **DCAT-ViT**. In addition, since there is a series of ViTs without the class token, *e.g.*, Swin [31] and

<sup>3</sup>To evaluate the effectiveness of the pre-trained POINT model, we labeled the MIP on the GAF 3.0 and GroupEmoW datasets, and further details can be found in the supplementary material.

Dataset	Label	Train/Val/Test	Scenario	Evaluation metric
GroupEmoW [13]	Neutral, Positive, Negative	11,127/3,178/1,589	Specific social events	Acc
GAF 3.0 [13, 15]	Neutral, Positive, Negative	9,815/4,349/3011	Specific social events	Acc
HECO [44]	Surprise, Excitement, Happiness, Peace, Disgust, Anger, Fear, and Sadness	6,570/1,877/938	In-the-wild images that do not limit scenes	mAP
GAF Cohesion [13, 15]	0-3	9,815/4,349/3011	Specific social events	MSE

Table 2. Datasets details used in this paper (Acc = Accuracy, mAP = mean average precision, MSE = mean square error).

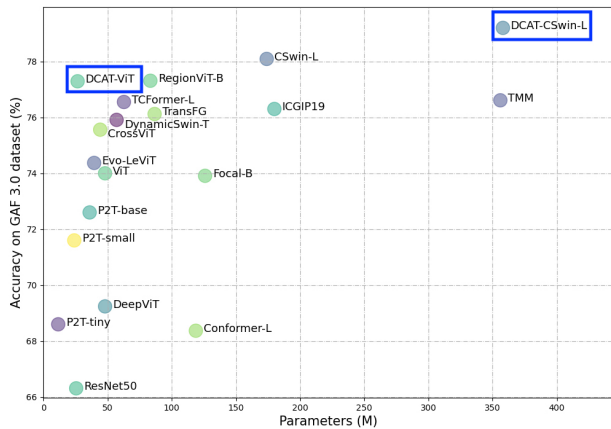


Figure 5. Model parameter and accuracy trade-off comparison on the GAF 3.0 dataset. Our proposed models are highlighted by the blue rectangular box.

CSwin [7]. In order to prove the generality of our approach, we have added the proposed dual-pathway and CPA module to the CSwin transformer [7], which is denoted by **DCAT-CSwin-L**.

A comparison of the proposed model with the state-of-the-art group affect approaches on GAF 3.0 and GroupEmoW datasets can be found in Figure 5 and Table 3, respectively. Since our proposed DCAT model is based on a vision transformer, we also compare it with recent vision transformer models, such as multi-scale-based (CrossViT [5], Focal-B [45], RegionViT-B [4], TCFormer-L [47]), CNN-transformer hybrid-based (Conformer-L [32]), and lightweight-based models (Evo-LeViT [43], DynamicSwin-T [36], DVT [40]).

**Comparing with state-of-the-art group affect methods.** Our DCAT-ViT and DCAT-CSwin-L achieve the best results compared to other group affect recognition methods on the GAF 3.0 dataset. Specifically, our DCAT-ViT achieves 77.29% accuracy with only 26.70M parameters, which is only about 7.6% of TMM [9], while our DCAT-CSwin-L with a larger number of parameters (358M) achieves 79.20% accuracy. On the GroupEmoW dataset, our DCAT-CSwin-L also achieves the best performance compared to other methods. It is worth noting that our proposed method only needs the global image and corre-

Methods	Acc (%)	Source
<i>Group affect methods</i>		
ResNet34 [18]	68.13	Global, Faces
SE-ResNet-50 [20]	69.79	Global
Efficientnet-b2 [38]	72.33	Global, Faces
CAER-Net [25]	80.61	Global, Faces
GNN [16]	82.38	Global
GNN [16]	84.62	Global, Faces, Objects
ConGNN [41]	85.59	Global, Faces, Objects
WACV 21 [22]	89.36	Global, Faces
GNN [16]	89.93	Global, Faces
WACV 21 [22]	90.18	Global, Faces, Objects
<i>Vision transformers</i>		
ViT [8]	87.50	Global
P2T-base [42]	86.50	Global
DeepViT [48]	83.90	Global
Conformer-L [32]	83.35	Global
Focal-B [45]	87.98	Global
CrossViT [5]	88.48	Global
TransFG [17]	89.47	Global
RegionViT-B [4]	89.49	Global
CSwin-L [7]	89.90	Global
DVT [40]	87.90	Global
Evo-LeViT [43]	87.83	Global
TCFormer-L [47]	89.24	Global
DCAT-ViT	89.55	Global, MIP
<b>DCAT-CSwin-L</b>	<b>90.47</b>	Global, MIP

Table 3. Overall Accuracy Comparison on the GroupEmoW dataset. "Source" refers to the types of features used in each paper.

sponding MIP image, while current group affect methods need multiple pre-trained detectors, which may be influenced by imperfect faces and objects. Furthermore, our method learns discriminative clues end-to-end in coordination with self-attention learning.

**Comparing with single focal attention methods.** A comparison is made between our dual-pathway attention model and methods specifically designed for local attention (Focal Transformers [45] and TransFG [17]). This finding highlights the fact that saliency does not always equate to

Method	mAP (%)
GCNs [23]	48.84
Depth [28]	49.02
OpenFace [1]+OpenPose [3]	50.28
R-FCN [6]+HOI-Net [30]	50.34
HECO [44]	50.24
<b>DCAT-ViT</b>	<b>53.49</b>
<b>DCAT-CSwin-L</b>	<b>58.06</b>

Table 4. Comparison on HECO dataset.

importance, and that our method’s attention mechanism can effectively identify and leverage the most critical affective regions.

Overall, our proposed DCAT-ViT and DCAT-CSwin-L perform better than other group affect recognition models and ViTs in terms of both accuracy and model complexity.

#### 4.4. The scalability of our method

Our proposed architecture can be adapted to accommodate the concept of a *primary agent*, similar to the idea of MIP, which was introduced by HECO [44]. Specifically, we can modify the input to the MIP branch from the MIP image to the primary agent while keeping the global branch unchanged.

Due to the wide variety of human-object interactions (HOIs) in natural scenes and the fact that interactions of agents’ actions with objects can induce emotional arousal of the primary agent in the scene, a useful approach for emotion recognition in images is to detect HOIs. In previous approaches [44, 25], multiple contexts, such as regions of interest (ROIs) and HOIs of all objects, were considered. Our proposed method outperforms previous approaches, as shown in Table 4. Our proposed method achieved superior performance on the dataset, which may be attributed to the wide variety of HOI categories in various challenging scenarios. Previous methods relied on pre-trained HOI models that were limited by predefined interaction categories. In contrast, our Cross-Patch Attention (CPA) module leverages the ability to learn HOI interactions directly from the data, allowing for greater flexibility and adaptability in recognizing new and diverse HOI categories. Moreover, our dual-pathway architecture significantly simplifies the computational process, eliminating the need for a heuristic selection of affective regions.

#### 4.5. Transfer learning for group cohesion estimation

Group cohesion is an affective phenomenon and a measure of agreement between members. Since cohesion and emotion (or valence) are intertwined in social science and affective computing, our DCAT for group valence prediction can be extended to group cohesion estimation. In this work, we finetune our model pre-trained on GroupEmoW

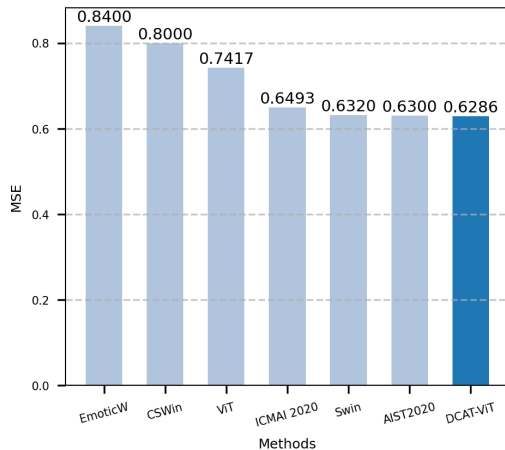


Figure 6. Group cohesion score comparison in the GAF Cohesion dataset (MSE is the evaluation metric, the lower the better).

Module	GAF 3.0	GroupEmoW
CrossVit	75.56	88.48
CrossVit+A	75.79	88.86
CrossVit+A+B	76.21	88.96
CrossVit+A+B+C	<b>77.29</b>	<b>89.55</b>
CSwin-L	78.09	89.90
CSwin-L+A	78.93	90.12
CSwin-L+A+B	<b>79.20</b>	<b>90.47</b>

Table 5. The accuracy (%) of each proposed module. *A*, *B*, and *C* represent dual-pathway input, CPA module, and Token Ranking Module, respectively. The downsampling operation in each stage of the CSwin-L can be regarded as a token selection mechanism, so the Token Ranking Module is not included in it.

for the GAF Cohesion dataset. The results in Figure 6 assure that our models still have good generalization ability rather than only fit to group affect datasets compared to other group affect recognition methods, *e.g.*, EmoticW [13], ICMAT2020 [50], AIST2020 [10], as well as the ViTs, *i.e.*, ViT [8], CSwin [7], Swin transformer [31].

#### 4.6. Ablation study

**The effectiveness of each module.** The effectiveness of each module proposed in our work, *i.e.*, dual-pathway input, CPA module, Token Ranking Module, is shown in Table 5. For both the DCAT-ViT and DCAT-CSwin-L, each of the modules contributes to our overall results.

**The effectiveness of dual-input.** Hence, we should also dynamically consider the global affective context while learning group affect, as validated in Table 6, where MIP and global inputs together perform better than either alone.

**Cross-Cls-Attention (CCA) vs Our proposed Cross-**

Input	GAF 3.0	GroupEmoW
Only MIP	68.02	80.08
Only global	75.56	88.48
<b>MIP + global</b>	<b>77.29</b>	<b>89.55</b>

Table 6. Effectiveness of dual-input on GAF 3.0 [14] and GroupEmoW [22] dataset (evaluation metric is ACC (%)).

	GAF 3.0	GroupEmoW
CCA	75.41%	87.86%
<b>CPA</b>	<b>77.29%</b>	<b>89.55%</b>

Table 7. Comparison between the CCA and our proposed CPA (using the DCAT-ViT) on GAF 3.0 and GroupEmoW dataset.

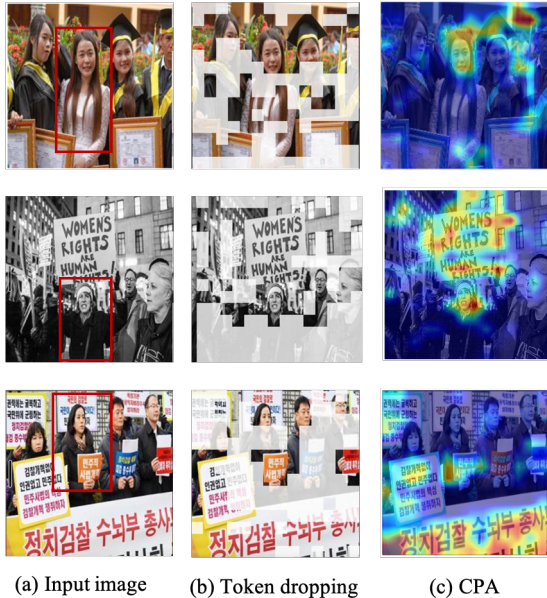


Figure 7. The visualization of the interaction of global and MIP.

**Patch-Attention (CPA).** It is worth comparing between Cross-Cls-Attention (CCA) proposed in CrossViT [5] and our proposed CPA (as provided in Table 7). The CCA only exchanges the class token in the cross-attention mechanism, which limits its ability to capture global interactions between class tokens and patches in each branch. In complex scenarios, such as those encountered in group affect recognition, our proposed CPA approach is more effective for modeling the interdependencies between global and MIP features.

#### 4.7. Visualization of the interaction of global and MIP pathways

The interaction of the global and MIP pathways, including token dropping and CPA, is demonstrated in Fig-

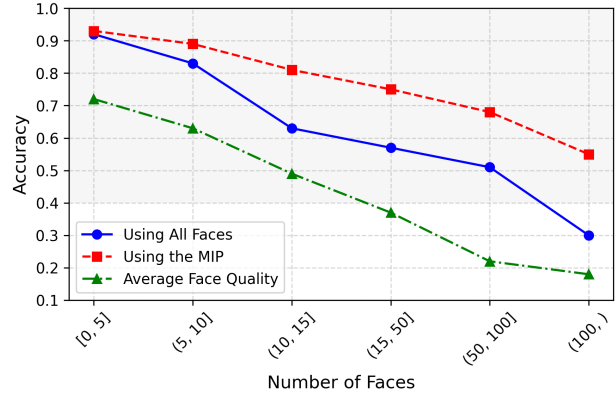


Figure 8. The performance of average face quality, using all faces and using the MIP under different face count intervals in the image. The GroupEmoW dataset is used here.

ure 7. As depicted in column (b), irrelevant background and other facial details are eliminated, while the MIP region and significant background regions (such as banners in protest scenes) are preserved for cross-patch attention. This demonstrates the effectiveness of our MIP-global interaction and highlights its efficacy in capturing discriminative information.

#### 4.8. Case study: the relationship between group size and performance

Previous studies have suggested that distinguishing between the *negative* and *neutral* categories in a group can be challenging [9]. In this paper, we present new insights into this issue. We first analyze the number of individuals in each image and find that as the number of individuals increases, the average quality of their facial features decreases<sup>4</sup>, as illustrated in Figure 8. This indicates that facial details become increasingly blurred and difficult to extract, posing a significant challenge, particularly when trying to distinguish between negative and neutral categories, which are inherently difficult to distinguish. In contrast to previous works that typically average features across all detected faces [9], our approach selectively focuses on the most informative facial region, i.e., the MIP, and outperforms models that utilize all faces in images with larger populations. This approach avoids interference from low-quality faces and effectively interacts between the features of the MIP and those of the entire global image through our proposed CPA module, enhancing the accuracy of our results. Therefore, our approach yields higher accuracy for distinguishing between the negative and neutral categories compared to previous methods. The exact numerical results can be found in the supplementary material.

<sup>4</sup>We quantify face quality based on [39].



## 5. Conclusions

In this work, by incorporating the MIP concept, we propose a dual-pathway vision transformer model, the Dual-branch Cross-Patch Attention Transformer (DCAT) for group affect recognition. Moreover, we propose the CPA module so that the important tokens from both paths can be interacted with. In terms of accuracy and parameters, the proposed DCAT outperforms both state-of-the-art group affect models and vision transformer models. It is also the first study to explore group affect recognition beyond CNNs. Moreover, our proposed model can be utilized in other group-level affect tasks, *i.e.*, group cohesion analysis.

## Acknowledgement

This work was supported in part by the National Science and Technology Council of Taiwan under Grants NSTC-109-2223-E-002-005-MY3, NSTC-112-2628-E-002-033-MY4 and NSTC-112-2634-F-002-002-MBK.

## References

- [1] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. Openface: an open source facial behavior analysis toolkit. In *2016 IEEE winter conference on applications of computer vision (WACV)*, pages 1–10. IEEE, 2016. 7
- [2] Sigal G Barsade and Donald E Gibson. Group affect: Its influence on individual and group outcomes. *Current Directions in Psychological Science*, 21(2):119–123, 2012. 1, 2
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 7
- [4] Chun-Fu Chen, Rameswar Panda, and Quanfu Fan. Regionvit: Regional-to-local attention for vision transformers. *arXiv preprint arXiv:2106.02689*, 2021. 6
- [5] Chun-Fu Richard Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 357–366, 2021. 2, 3, 5, 6, 8
- [6] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. *Advances in neural information processing systems*, 29, 2016. 7
- [7] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12124–12134, 2022. 6, 7
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 3, 5, 6, 7
- [9] Katsuya Fujii, Daisuke Sugimura, and Takayuki Hamamoto. Hierarchical group-level emotion recognition. *IEEE Transactions on Multimedia*, 23:3892–3906, 2020. 2, 6, 8
- [10] Ilya Gavrikov and Andrey V Savchenko. Efficient group-based cohesion prediction in images using facial descriptors. In *International Conference on Analysis of Images, Social Networks and Texts*, pages 140–148. Springer, 2020. 7
- [11] Jennifer M George. Leader positive mood and group performance: The case of customer service. *Journal of Applied Social Psychology*, 25(9):778–794, 1995. 2
- [12] Shreya Ghosh and Abhinav Dhall. Role of group level affect to find the most influential person in images. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 2
- [13] Shreya Ghosh, Abhinav Dhall, Nicu Sebe, and Tom Gedeon. Predicting group cohesiveness in images. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019. 5, 6, 7
- [14] S. Ghosh, A. Dhall, N. Sebe, and T. Gedeon. Predicting group cohesiveness in images. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2019. 8
- [15] Shreya Ghosh, Abhinav Dhall, Nicu Sebe, and Tom Gedeon. Automatic prediction of group cohesiveness in images. *IEEE Transactions on Affective Computing*, 2020. 2, 5, 6
- [16] Xin Guo, Luisa Polania, Bin Zhu, Charles Boncelet, and Kenneth Barner. Graph neural networks for image understanding based on multiple cues: Group emotion recognition and event recognition as use cases. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2921–2930, 2020. 2, 5, 6
- [17] Ju He, Jie-Neng Chen, Shuai Liu, Adam Kortylewski, Cheng Yang, Yutong Bai, and Changhu Wang. Transfg: A transformer architecture for fine-grained recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 852–860, 2022. 6
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [19] Fa-Ting Hong, Wei-Hong Li, and Wei-Shi Zheng. Learning to detect important people in unlabelled images for semi-supervised important people detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4146–4154, 2020. 3
- [20] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 6
- [21] Janice R Kelly and Sigal G Barsade. Mood and emotions in small groups and work teams. *Organizational behavior and human decision processes*, 86(1):99–130, 2001. 1
- [22] Ahmed Shehab Khan, Zhiyuan Li, Jie Cai, and Yan Tong. Regional attention networks with context-aware fusion for group emotion recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1150–1159, 2021. 1, 2, 6, 8
- [23] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 7
- [24] R Benjamin Knapp, Jonghwa Kim, and Elisabeth André. Physiological signals and their use in augmenting emotion recognition for human-machine interaction. In *Emotion-oriented systems*, pages 133–159. Springer, 2011. 1
- [25] Jiyoung Lee, Seungryong Kim, Sunok Kim, Jungin Park, and Kwanghoon Sohn. Context-aware emotion recognition networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10143–10152, 2019. 6, 7
- [26] Wei-Hong Li, Fa-Ting Hong, and Wei-Shi Zheng. Learning to learn relation for important people detection in still images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5003–5011, 2019. 3, 5
- [27] Wei-Hong Li, Benchao Li, and Wei-Shi Zheng. Personrank: Detecting important people in images. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 234–241. IEEE, 2018. 3

- [28] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2041–2050, 2018. 7
- [29] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. In *The International Conference on Learning Representations (ICLR)*, 2022. 5
- [30] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jiashi Feng. Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 482–490, 2020. 7
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 5, 7
- [32] Zhiliang Peng, Wei Huang, Shanzhi Gu, Lingxi Xie, Yaowei Wang, Jianbin Jiao, and Qixiang Ye. Conformer: Local features coupling global representations for visual recognition. in 2021 IEEE. In *CVF International Conference on Computer Vision, ICCV*, pages 357–366, 2021. 6
- [33] Claire Polo, Kristine Lund, Christian Plantin, and Gerald P Niccolai. Group emotions: The social and cognitive functions of emotions in argumentation. *International Journal of Computer-Supported Collaborative Learning*, 11(2):123–156, 2016. 1
- [34] Fan Qi, Xiaoshan Yang, and Changsheng Xu. Zero-shot video emotion recognition via multimodal protagonist-aware transformer network. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1074–1083, 2021. 2
- [35] Vignesh Ramanathan, Jonathan Huang, Sami Abu-El-Haija, Alexander Gorban, Kevin Murphy, and Li Fei-Fei. Detecting events and key actors in multi-person videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3043–3053, 2016. 3
- [36] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. *Advances in neural information processing systems*, 34:13937–13949, 2021. 6
- [37] Thomas Sy, Stéphane Côté, and Richard Saavedra. The contagious leader: impact of the leader’s mood on the mood of group members, group affective tone, and group processes. *Journal of applied psychology*, 90(2):295, 2005. 2
- [38] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 6
- [39] Philipp Terhörst, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. SER-FIQ: unsupervised estimation of face image quality based on stochastic embedding robustness. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 5650–5659. IEEE, 2020. 8
- [40] Yulin Wang, Rui Huang, Shiji Song, Zeyi Huang, and Gao Huang. Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. *Advances in Neural Information Processing Systems*, 34:11960–11973, 2021. 6
- [41] Yu Wang, Shunping Zhou, Yuanyuan Liu, Kunpeng Wang, Fang Fang, and Haoyue Qian. Congnn: Context-consistent cross-graph neural network for group emotion recognition in the wild. *Information Sciences*, 610:707–724, 2022. 6
- [42] Yu-Huan Wu, Yun Liu, Xin Zhan, and Ming-Ming Cheng. P2t: Pyramid pooling transformer for scene understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 6
- [43] Yifan Xu, Zhijie Zhang, Mengdan Zhang, Kekai Sheng, Ke Li, Weiming Dong, Liqing Zhang, Changsheng Xu, and Xing Sun. Evo-vit: Slow-fast token evolution for dynamic vision transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2964–2972, 2022. 3, 5, 6
- [44] Dingkan Yang, Shuai Huang, Shunli Wang, Yang Liu, Peng Zhai, Liuzhen Su, Mingcheng Li, and Lihua Zhang. Emotion recognition for multiple context awareness. *ECCV*, 2022. 2, 5, 6, 7
- [45] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers, 2021. 6
- [46] Ting Yao, Yehao Li, Yingwei Pan, Yu Wang, Xiao-Ping Zhang, and Tao Mei. Dual vision transformer. *arXiv preprint arXiv:2207.04976*, 2022. 3
- [47] Wang Zeng, Sheng Jin, Wentao Liu, Chen Qian, Ping Luo, Wanli Ouyang, and Xiaogang Wang. Not all tokens are equal: Human-centric visual analysis via token clustering transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11101–11111, 2022. 6
- [48] Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*, 2021. 6
- [49] Haowei Zhu, Wenjing Ke, Dong Li, Ji Liu, Lu Tian, and Yi Shan. Dual cross-attention learning for fine-grained visual categorization and object re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4692–4702, 2022. 3, 5
- [50] Bochao Zou, Zhifeng Lin, Haoyi Wang, Yingxue Wang, Xianguwen Lyu, and Haiyong Xie. Joint prediction of group-level emotion and cohesiveness with multi-task loss. In *Proceedings of the 2020 5th International Conference on Mathematics and Artificial Intelligence*, pages 24–28, 2020. 7