

Prototypical Calibrating Ambiguous Samples for Micro-Action Recognition

Kun Li¹, Dan Guo^{*1,2}, Guoliang Chen^{*1}, Chunxiao Fan¹, Jingyuan Xu¹, Zhiliang Wu³,
Hehe Fan³, Meng Wang^{*1,2}

¹School of Computer Science and Information Engineering, Hefei University of Technology

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

³ReLER Lab, CCAI, Zhejiang University

kunli.hfut@gmail.com, guodan@hfut.edu.cn, guoliangchen.hfut@gmail.com, fanchunxiao@hfut.edu.cn,
xujingyuan@hfut.edu.cn, wu_zhiliang@zju.edu.cn, hehefan@zju.edu.cn, eric.mengwang@gmail.com

Abstract

Micro-Action Recognition (MAR) has gained increasing attention due to its crucial role as a form of non-verbal communication in social interactions, with promising potential for applications in human communication and emotion analysis. However, current approaches often overlook the inherent ambiguity in micro-actions, which arises from the wide category range and subtle visual differences between categories. This oversight hampers the accuracy of micro-action recognition. In this paper, we propose a novel Prototypical Calibrating Ambiguous Network (PCAN) to unleash and mitigate the ambiguity of MAR. **Firstly**, we employ a hierarchical action-tree to identify the ambiguous sample, categorizing them into distinct sets of ambiguous samples of false negatives and false positives, considering both body- and action-level categories. **Secondly**, we implement an ambiguous contrastive refinement module to calibrate these ambiguous samples by regulating the distance between ambiguous samples and their corresponding prototypes. This calibration process aims to pull false negative (FN) samples closer to their respective prototypes and push false positive (FP) samples apart from their affiliated prototypes. In addition, we propose a new prototypical diversity amplification loss to strengthen the model’s capacity by amplifying the differences between different prototypes. **Finally**, we propose a prototype-guided rectification to rectify prediction by incorporating the representability of prototypes. Extensive experiments conducted on the benchmark dataset demonstrate the superior performance of our method compared to existing approaches.

1 Introduction

Human body actions are an important form of non-verbal communication in social interactions (Aviezer, Trope, and Todorov 2012). In recent years, there has been increasing interest in human-centric action understanding (Balazia et al. 2022; Liu et al. 2021; Chen et al. 2023; Guo et al. 2024a; Li et al. 2023b, 2024b). For instance, body language in naturalistic multiview group conversations influences perceptions of leadership and relationships during social interactions (Balazia et al. 2022; Li et al. 2023a). Micro-Action (MA), as a subset, has emerged as a significant research area

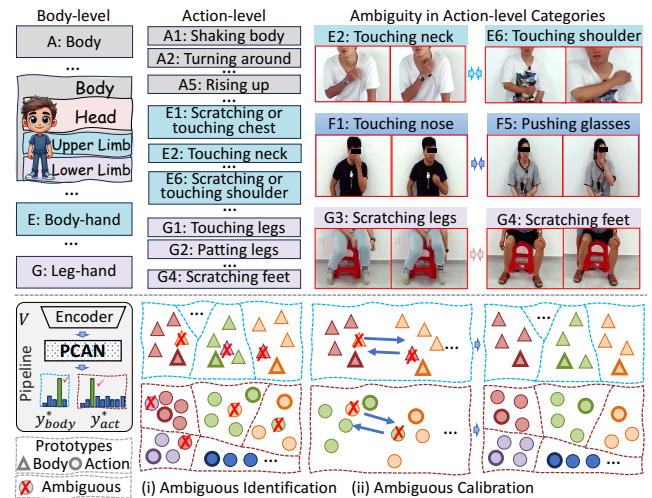


Figure 1: **TOP:** Micro-Action Recognition (MAR) aims to recognize the body-level and action-level micro-action categories, particularly when dealing with ambiguous samples. For example, “touching shoulder” and “touching neck” belong to the “body-hand” but exhibit subtle visual differences. **BOTTOM:** Our approach is motivated by the need to address ambiguities in MAR. We begin by identifying ambiguous samples (marked in \times) that are prone to misclassification. Subsequently, we construct prototypes for each category within the body and action levels, and then align the ambiguous samples with the corresponding prototypes within the feature space. Please Zoom in for details.

due to its potential applications in human-to-human communication and analyzing human emotional states (Liu et al. 2021; Chen et al. 2023; Guo et al. 2024a,b; Li et al. 2024a).

As body language is an essential part of understanding human emotions according to psychological studies (Shiffrar, Kaiser, and Chouhroulou 2011; Aviezer, Trope, and Todorov 2012), iMiGUE (Liu et al. 2021) and SMG (Chen et al. 2023) explored spontaneous micro-gestures in upper limbs driven by a person’s inner feelings, revealing the deep emotional states hidden behind the subtle actions. To better analyze and understand the spontaneous micro-actions, whole-body movement, and fine-grained action labels are

*Corresponding authors

crucial. Therefore, (Guo et al. 2024a) collected a large-scale spontaneous whole-body micro-actions dataset that includes the head, hands, and upper and lower limbs, using face-to-face psychological interviews way. This work extends micro-action from the upper limb to the whole body and builds coarse-to-grained action understanding, benchmarking the datasets, and bringing challenges for human-centric micro-action understanding and analysis. Compared to generic action recognition (Lin, Gan, and Han 2019; Feichtenhofer et al. 2019; Liu et al. 2025; Mou et al. 2024) and skeleton-based action recognition (Wang and Koniusz 2023; Zhu et al. 2022; Hao et al. 2021), micro-actions encompass a multitude of approximate intricate samples, extending across various body regions.

As shown in Fig. 1, MAR aims to recognize categories at the *body-level* and the *action-level*. *Body-level* category denotes the body part of micro-action occurring, such as body, body-hand, and leg-hand. *Action-level* category denotes the exact name of micro-actions. Intuitively, body-level and action-level categories enjoy the nature of hierarchy. The actions of “touching shoulder” and “touching neck” are very similar, both involving hand movement around the upper body. However, current methods typically overlook the inherent ambiguity in micro-actions. To reduce the ambiguity in MAR, one solution is to build prototypes that represent the features of each category. However, most of the existing works on few-shot (Zhu et al. 2021; Wang et al. 2024) or egocentric action recognition (Wang et al. 2021; Dai et al. 2023) and can not be directly applied to tackle the challenges associated with action ambiguity in MAR.

To address the challenges of ambiguity of micro-actions, we propose a Prototypical Calibrating Ambiguous Network (PCAN). As shown in Fig 2, PCAN contains the following designs: 1) **Ambiguous Samples Identification** (§3.2). Guided by preliminary predicted scores from the backbone model, we initially identify ambiguous samples at body-level and action-level. Concurrently, we mine confidence samples - those correctly classified based on the predicted scores - to capture the fundamental essence and overarching characteristics of action, ensuring the integrity of the prototypes. 2) **Ambiguous Samples Contrastive Calibration** (§3.3). The identified ambiguous samples are further categorized into false negative (FN) and false positive (FP). We then apply contrastive refinement to enhance the representability of prototypes by pulling FN samples closer to their corresponding prototypes and pushing FP samples away. Additionally, we amplify the differences between distinct prototypes to strengthen the model’s ability to understand inter-class differences. 3) **Prototype-guided Rectification** (§3.4). Considering that prototypes represent the intrinsic characteristics of each category, we propose a prototype-guided prediction rectification module. By evaluating the similarities between input samples and prototypes, this module refines the model’s classification outcomes for MAR. Overall, our contributions are summarized as follows.

- We attempt to address the ambiguity of micro-actions with prototype learning, wherein prototypes serve as effective representations for discerning micro-actions, providing a compact representation for each category.

- We build a prototype-guided rectification approach is implemented to enhance model prediction performance. By evaluating dissimilarities between input samples and prototypes corresponding to different categories, prediction outcomes are rectified.
- Extensive experiments conducted on the public micro-action dataset, MA-52, validate the effectiveness of the proposed method. Comparative analysis highlights the effectiveness of the prototypical calibrating ambiguous network in micro-action recognition.

2 Related Work

Micro-Action Recognition Micro-Actions (Guo et al. 2024a) refer to subtle and rapid movements that occur across various body parts (Liu et al. 2021; Chen et al. 2023; Liu et al. 2024). Existing researches primarily adopt the generic action recognition methods (*e.g.*, 2D CNN based (Lin, Gan, and Han 2019), 3D CNN based (Tran et al. 2015), GCN-based (Yan, Xiong, and Lin 2018), and Transformer-based (Liu et al. 2022)) to realize micro-action recognition. However, addressing the challenges of subtle differences and high similarity remains difficult with generic action recognition methods. Effectively capturing micro-actions demands in-depth understanding and specialized modeling. We conduct experiments on the first human-centered whole-body micro-action dataset named MA-52 (Guo et al. 2024a). It stands out for its extensive range of micro-action categories and abundant samples, encompassing both similar and distinct micro-action categories.

Prototype Learning for Action Recognition Prototype learning aims to obtain representative feature representation, and has demonstrated substantial potential in addressing the challenges of few-shot scenarios (Zhu et al. 2021; Wang et al. 2024), egocentric perspectives (Wang et al. 2021; Dai et al. 2023), and early action prediction in action recognition (Camporese et al. 2023). However, in the domain of micro-action recognition, action ambiguity remains a significant challenge due to the subtle differences between actions, which can lead to misclassification. Previous prototype learning methods can partially mitigate this issue by creating more discriminative and representative prototypes. Despite these advancements, they do not fully resolve the challenges associated with action ambiguity in micro-actions. In this work, we separate the ambiguous samples into sets of false positives and false negatives with the guidance of prototype learning and calibrate these ambiguous samples according to the prototypical distance measurement with their affiliated prototypes.

Ambiguous/Hard Sample Learning The ambiguous samples are difficult to classify or distinguish clearly due to factors such as similar actions, occlusion, variety, or data quality issues (Xia et al. 2023). In recent years, effectively distinguishing ambiguous samples has proven a difficult challenge (Li et al. 2020, 2021; Zhou, Liu, and Wang 2023). For example, (Zhou, Liu, and Wang 2023) utilized prototypical contrastive learning to constrain the distance between confident samples and ambiguous samples to improve the recognition of ambiguous actions that are quite similar and easily

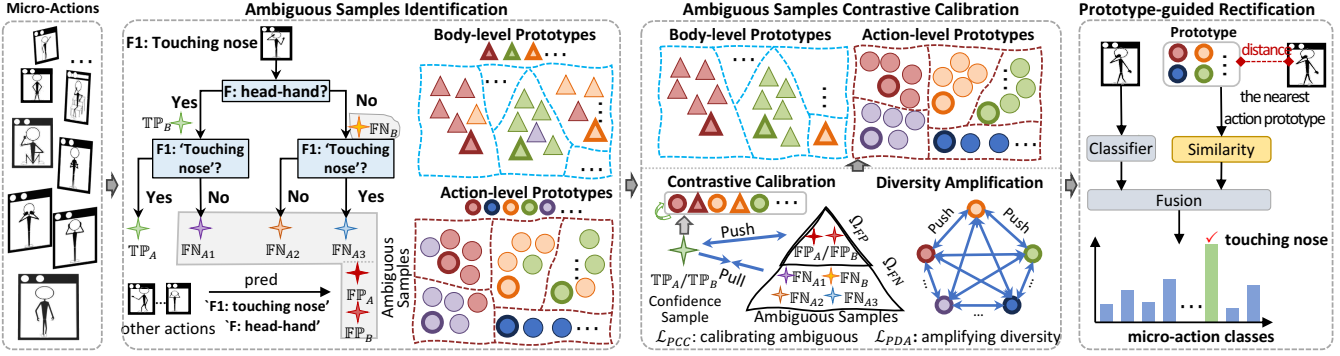


Figure 2: The overview of Prototypical Calibrating Ambiguous Network (PCAN). In **Ambiguous Samples Identification** (§3.2), we discover ambiguous samples (FN and FP) through the preliminary prediction scores. In **Ambiguous Samples Contrastive Calibration** (§3.3), we use contrastive prototype calibration and prototype diversity amplification losses to calibrate the prototypes and eliminate the influence of ambiguous samples. In **Prototype-guided Rectification** (§3.4), we incorporate the similarity between established prototypes and video embedding for action-level category prediction.

misclassified. However, all of the above methods only consider prototypes at a certain grain and lack consideration of cascade coarse- and fine-grained label categories. We propose cascade prototypes that work to discover and calibrate ambiguous samples between these micro-actions in order to reduce the category ambiguity between them.

3 Methodology

3.1 Preliminary and Core Idea

Preliminary. Give a video V , the *Micro-Action Recognition* (MAR) task (Guo et al. 2024a) predicts the micro-action categories at body-level $\mathcal{Y}_{body} \in \{1, 2, \dots, N_B\}$ and at action-level $\mathcal{Y}_{act} \in \{1, 2, \dots, N_A\}$, where N_B and N_A denotes the number of body and action categories, respectively. PoseConv3D (Duan et al. 2022) is a novel framework for RGB-Pose dual-modality action recognition. Concretely, the input RGB data \mathcal{D}_{RGB} and skeleton data \mathcal{D}_{Pose} are fed to separate 3D-CNN based encoder \mathcal{E}_R and \mathcal{E}_P to extract the feature. After that, the predictions of the RGB/Pose are combined through late fusion to compute class scores for action recognition. In this work, we conduct experiments on PoseConv3D to validate the motivation to remove ambiguity in MAR. Considering that RGB and Pose are parallel operations, we take the RGB branch as an example to elaborate.

Core Idea. In this study, we argue that the misclassification in MAR primarily arises from the inherent ambiguity in subtle differences among action-level categories. To address this, we first identify the ambiguous samples at the body and action levels, which are based on the preliminary classification results. Simultaneously, we mine the confidence samples to build body- and action-level prototypes. Then, we utilize a hierarchical contrastive calibration mechanism to reduce ambiguity by regulating the distance between ambiguous samples and their corresponding prototypes. Finally, we leverage the representability of the prototype to perform prototype-guided rectification for action-level classification.

3.2 Ambiguous Samples Identification

Action-tree Hierarchy Modeling. In micro-action recognition, the categories at the body level and action level naturally follow a hierarchical structure. Namely, the Micro-Action existing “**body**→**action**” structural feature, namely, each **action-level** category belongs to a specific primary **body-level** category. For instance, MA “B1: nodding” has a corresponding body label “B: head”. Inspired by this, we propose a hierarchical action-tree structure to model the relation between body- and action-level categories. Let $F_i \in \mathbb{R}^d$ be the learned feature from the backbone of a sample i and two classifiers $\text{CLS}_B(\cdot)$ and $\text{CLS}_A(\cdot)$ to predict the probability of action-level and body-level category of F_i . To inject the hierarchy structure into action prediction, we build the hierarchical action probability q_J :

$$q_J = \lambda \cdot \text{CLS}_B(F_i) \oplus \text{CLS}_A(F_i) = \lambda \cdot q_B \oplus q_A, \quad (1)$$

where λ is a hyper-parameter that is optimally set to 1 and discussed in §4.3, and \oplus is the operation to incorporate the body- and action-level probabilities. q_B and q_A denote the body-level and action-level probability, respectively. To optimize q_J , we define a hierarchical probabilistic loss \mathcal{L}_{HP} to maximize the marginal probability of action-grained predictions by aggregating the structural information of the relevant labels defined as described in the action-tree:

$$\mathcal{L}_{HP} = -\log \frac{\exp(q_J^{GT})}{\sum_{i=1}^{N_B+N_A} \exp(q_J^i)}, \quad (2)$$

where q_J^{GT} is the probability score corresponding to the index of ground truth of the action-level label. Note that q_J is only used in loss \mathcal{L}_{HP} .

Ambiguous Samples. As stated in §1, the ambiguity of micro-actions is caused by the minor differences in action-level, which will hinder the model from achieving optimal results. Based on the classification results from classifiers CLS_B and CLS_A , we can identify the ambiguous samples. (i) **Ambiguous Samples - FN set.** In the classification task,

\mathbb{FN} (False Negative) refers to instances where the model incorrectly predicts the negative class when the true class is positive. If we merely consider the body-grained, instances of body misclassified as other categories are termed as false negatives (\mathbb{FN}_B). We also can consider both body- and action-level predictions. If body-level is correct but action-level is wrong, it is denoted as \mathbb{FN}_{A1} ; if both body-level and action-level are wrong, it is denoted as \mathbb{FN}_{A2} ; if body-level is wrong but action-level is correct, it is denoted as \mathbb{FN}_{A3} . **(ii) Ambiguous Samples - \mathbb{FP} set.** In action-level, if samples from other categories are mistakenly classified as the current category, they are labeled as \mathbb{FP}_A samples. This split is evaluated solely at the action-grained level. Instances from other categories misclassified as the referred body-grained category are referred to as false positives (\mathbb{FP}_B). Please note that the identification of ambiguous samples occurs online during training, and we built ambiguous sample sets for both body-level and action-level categories.

Hierarchical Prototype Initialization. Considering the two-level category in MAR, we define the body-level prototypes as $\mathbf{P}_B = \{p_1, p_2, \dots, p_{N_B}\}$, and the action-level prototypes as $\mathbf{P}_A = \{p_1, p_2, \dots, p_{N_A}\}$. $\mathbf{P}_B \in \mathbb{R}^{N_B \times d}$ and $\mathbf{P}_A \in \mathbb{R}^{N_A \times d}$ serve as the stable estimate of the clustering center for each micro-action category. Initially, these prototypes are randomly initialized and then optimized throughout the training process. In prototype learning, correctly predicted samples capture the fundamental essence and overarching characteristics of an action. These samples are considered confident samples. The body-level confident sample set is defined as \mathbb{TP}_B for samples that correctly predicted their body-level labels. Considering the ambiguity of action-level, the samples predicted at both body-level and action-level are called confident sample sets, denoted as \mathbb{TP}_A . Thus, we have $\Omega_{TP} = \{\mathbb{TP}_B, \mathbb{TP}_A\}$, which represents two kinds of confident samples in the body- and action-level prototypes.

Online Prototype Update. Since the confidence samples (Ω_{TP}) usually have better intra-class consistency, we use it to update the prototype representation of the corresponding category k by exponential moving average (EMA) (Ye et al. 2019; Wu et al. 2018; Wang et al. 2022):

$$p_k = (1 - \rho) \cdot \frac{1}{N_{TP}^k} \sum_{i \in \Omega_{TP}^k} F_i + \rho \cdot p_k^{pre}, \quad (3)$$

where F_i is the feature of i -th sample, p_k^{pre} is the prototype before updating, N_{TP}^k is the size of \mathbb{TP} sample set Ω_{TP} under category k , and ρ is a momentum term and empirically set to 0.9 as in (Zhou, Liu, and Wang 2023).

3.3 Ambiguous Samples Contrastive Calibration

Through the above operations, we obtained ambiguous samples and prototypes. However, the ambiguous samples and prototypes are misaligned in the feature space due to the ambiguity of micro-actions. To address this issue, we utilize a contrastive prototype loss (Zhou, Liu, and Wang 2023; Wang et al. 2022) to regulate the distance between ambiguous samples and their prototypes. Concretely, we aim for \mathbb{FP} samples to be distant from the misassigned prototype and

for \mathbb{FN} samples to be close to the correct prototype. Here, we use the action-level prototype and ambiguous samples as examples to detail the contrastive calibration. The approach of action-level calibration is analogous.

Prototypical Contrastive Calibration For the ambiguous sample sets Ω_{FN} (\mathbb{FN} samples) and Ω_{FP} (\mathbb{FP} samples), we first estimate their mean values as the center representations in the feature space:

$$\mu_\omega^k = \frac{1}{N_\omega^k} \sum_{i \in \Omega_{FN}^{k,\omega}} F_i, \quad \mu_{FP}^k = \frac{1}{N_{FP}^k} \sum_{i \in \Omega_{FP}^k} F_i, \quad (4)$$

where k is action category, ω denotes the type of ambiguous samples in the \mathbb{FN} set $\Omega_{FN} = \{\mathbb{FN}_{A1}, \mathbb{FN}_{A2}, \mathbb{FN}_{A3}\}$ defined in §3.2. N_ω^k and N_{FP}^k are the sizes of sample sets $\Omega_{FN}^{k,\omega}$ and Ω_{FP}^k under category k , respectively.

(i) Ambiguous Samples - \mathbb{FN} set. The auxiliary term ψ_i is used to calibrate the model's predictions for true action categories as,

$$\psi_i^\omega = \begin{cases} 1 - \text{dis}(F_i, \mu_\omega^k), & \text{if } N_\omega^k > 0 \text{ and } i \in \Omega_{FN}^{k,\omega}; \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

where $\text{dis}(\cdot)$ denotes the cosine distance between two features. This operation brings the ambiguous samples closer to the confidence samples (*i.e.*, prototypes) in the feature space. This adjustment helps rectify the predicted logits of the ambiguous samples, aligning them more closely with the correct action categories. By minimizing ψ_i^ω , the model more accurately identifies the ambiguous samples, as a smaller value indicates that the ambiguous samples are in closer proximity to the prototypes. **(ii) Ambiguous Samples - \mathbb{FP} set.** For ambiguous samples \mathbb{FP} in set Ω_{FP} , we define the auxiliary term Ψ_i , which serves to adjust the predictions for non-true action categories as,

$$\Psi_i = \begin{cases} 1 + \text{dis}(F_i, \mu_{FP}^k), & \text{if } N_{FP}^k > 0 \text{ and } i \in \Omega_{FP}^k; \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where Ψ_i^ω is the penalty term for sample sets in Ω_{FP} , which can keep the ambiguous samples in Ω_{FP} away from the confidence samples in the feature space. By minimizing Ψ_i^ω , the distance between ambiguous samples and prototypes can be increased, thus mitigating the risk of the model incorrectly classifying ambiguous samples as action category k .

In order to keep false positive (\mathbb{FP} in Ω_{FP}) samples distant from prototype and false negative (\mathbb{FN} in Ω_{FN}) samples close to the prototype, we define a matrix \mathbf{Z}_i^ω to represent the distance between the ambiguous samples and their prototype as follows:

$$\begin{cases} \mathbf{Z}_i^\omega = \tau \cdot \text{dis}(F_i, p_k) - (1 - q_{ik})\psi_i^\omega, & \text{if } i \in \Omega_{FN}^\omega; \\ \mathbf{Z}_i^{FP} = \tau \cdot \text{dis}(F_i, p_k) - (1 - q_{ik})\Psi_i, & \text{if } i \in \Omega_{FP}, \end{cases} \quad (7)$$

where p_k denotes the prototype feature of action k , τ is a hyper-parameter and is set to 0.125, q_{ik} is the probability of sample i in micro-action k . Thus, the prototype contrastive

calibration loss \mathcal{L}_{PCC} is formulated as follows:

$$\mathcal{L}_{PCC} = - \sum_{i \in \Omega_{FN}^{\omega}} \log \frac{e^{\mathbf{Z}_i^{\omega}}}{e^{\mathbf{Z}_i^{\omega}} + \sum_{l \neq k} e^{dis(F_i, p_l) \cdot \tau}} - \sum_{j \in \Omega_{FP}} \log \frac{e^{\mathbf{Z}_i^{FP}}}{e^{\mathbf{Z}_i^{FP}} + \sum_{j \neq k} e^{dis(F_i, p_j) \cdot \tau}}, \quad (8)$$

where $\omega \in \{\mathbb{F}N_{A1}, \mathbb{F}N_{A2}, \mathbb{F}N_{A3}\}$ denotes the type of ambiguous samples in Ω_{FN} defined in Sec. 3.2. We employ α_1 , α_2 , and α_3 as the weights to balance the set of Ω , and discussed in §4.3.

Prototypical Diversity Amplification As stated in §1, the ambiguity is caused by the high similarity of action-level categories under the same body part. For instance, head movements like “head up” and “nodding” enjoy similar motion patterns, leading to difficulty in differentiation. To mitigate the influence of high similarity among different categories, a prototypical diversity amplification is designed to amplify the diversity among different categories:

$$\mathcal{L}_{PDA} = \sqrt{\sum_{j=1}^{N_A} \sum_{i=1}^{N_A} \frac{p_i \cdot p_j}{\|p_i\| \|p_j\|}}, \quad (9)$$

where p_i and p_j stand for action-level prototypes.

3.4 Prototype-guided Rectification

In the previous section, we established prototypes for each category, providing a clear reference that represents the intrinsic characteristics of different categories. Building on this, we propose a prototype-guided rectification module that rectifies prediction results by analyzing the cosine similarity between input samples and established prototypes. By comparing the cosine similarity between prototypes and output features, we can capture subtle differences missed in the initial prediction. This adjustment refines the final prediction, improving both accuracy and reliability. The final prediction probability of sample F_i is as follows:

$$F_{act,i} = \mathbf{W}_1 F_i + b_1 + \gamma \cdot dis(\mathbf{W}_2 F_i + b_2, \mathbf{P}_A), \quad (10)$$

$$y_{act,i}^* = \arg \max_{y \in \mathcal{Y}_{act}} (\text{SoftMax}(F_{act,i})), \quad (11)$$

where $\mathbf{W}_1 \in \mathbb{R}^{d \times N_A}$ and $\mathbf{W}_2 \in \mathbb{R}^{d \times N_A}$ are two learnable matrices to transform the embedding F_i . b_1 and b_2 are two biases. \mathbf{P}_A is the prototypes of action-level categories. γ serves as a balancing weight. For body-level probability $y_{body,i}^*$, we adopt a body-level classifier to predict it.

3.5 Model Optimization

We use the cross-entropy loss as the basic objective for the micro-action recognition task, denoted as \mathcal{L}_{CE} . Considering that PCAN is RGB-Pose dual-modality based, we implement losses \mathcal{L}_{CE} , \mathcal{L}_{HP} , \mathcal{L}_{PCC} and \mathcal{L}_{PDA} in both RGB and Pose. Finally, the total loss is formulated as follows:

$$\mathcal{L} = \mathcal{L}_{CE} + \mathcal{L}_{HP} + \mathcal{L}_{PCC} + \beta \mathcal{L}_{PDA}, \quad (12)$$

where β is the balance hyperparameters for the loss term.

Method	Body			Body		Action		All
	Top-1	Top-1	Top-5	F1 _{macro}	F1 _{micro}	F1 _{macro}	F1 _{micro}	
TSN [CVPR'18]	59.22	34.46	73.34	52.50	59.22	28.52	34.46	43.67
TIN [AAAI'20]	73.26	52.81	85.37	66.99	73.26	39.82	52.81	58.22
TSM [CVPR'18]	77.64	56.75	87.47	70.98	77.64	40.19	56.75	61.39
MANet [TCSVT'24]	78.95	61.33	88.83	72.87	78.95	49.22	61.33	65.59
C3D [CVPR'14]	74.04	52.22	86.97	66.60	74.04	40.86	52.22	58.43
I3D [CVPR'17]	78.16	57.07	88.67	71.56	78.16	39.84	57.07	61.66
SlowFast [ICCV'19]	77.18	59.60	88.54	70.61	77.18	44.96	59.60	63.09
VSwiT [CVPR'22]	77.95	57.23	87.99	71.25	77.95	38.53	57.23	61.24
TimesFormer [ICML'21]	69.17	40.67	82.67	61.90	69.17	34.38	40.67	51.53
Uniformer [ICLR'22]	79.03	58.89	87.29	71.80	79.03	48.01	58.89	64.43
STGCN [CVPR'18]	69.87	49.61	79.54	61.53	69.87	34.64	49.61	53.91
2s-AGCN [CVPR'19]	70.07	49.48	78.27	61.30	70.07	34.64	49.48	53.87
CTRGCN [ICCV'21]	72.06	52.61	81.22	63.46	72.06	37.79	52.61	56.48
B2C-AFM [TIP'24]	78.34	57.12	88.45	72.53	78.34	44.12	57.12	63.03
PoseConv3D (RGB)	77.77	56.30	86.11	69.87	77.77	35.80	56.30	59.93
PoseConv3D (Pose)	74.60	50.93	84.98	67.31	74.60	31.07	50.93	55.98
PoseConv3D [CVPR'22]	80.95	63.52	90.23	74.96	80.95	47.20	63.52	66.66
Ours (RGB)	79.36	60.03	87.65	72.37	79.36	43.29	60.03	63.76
Ours (Pose)	76.99	56.51	87.13	71.33	76.99	42.28	56.51	61.78
Ours	82.30	66.74	91.75	77.02	82.30	53.83	66.74	69.97

Table 1: Performance comparison of micro-action recognition on the MA-52 dataset. Top-1/5 are Top-1/5 accuracy.

4 Experiments

4.1 Experimental Setup

Micro-Action Dataset. MA-52 (Guo et al. 2024a)¹ is a large-scale human-centered whole-body micro-action dataset that takes an interview format to record unconscious human micro-action behaviors. It contains micro-actions in the whole body, including the body, head, upper, and lower limbs. The dataset contains 22,422 (22.4K) videos interviewed from 205 participants with annotations categorized into two levels: 7 body-level and 52 action-level micro-action categories, making it a large-scale dataset of the largest samples and micro-action categories that we are aware of to date. The dataset consists of 11,250, 5,586, and 5,586 instances in training/validation/test, respectively.

Evaluation Metrics. Following the practice (Guo et al. 2024a,b), we adopt the metrics of Top-1/5 accuracy and macro and micro F1 scores of body and action levels to evaluate the performance of the proposed model. F1_{mean} is used as a general evaluation metric, and its formula is as follows:

$$F1_{mean} = (F1_{macro}^{body} + F1_{micro}^{body} + F1_{macro}^{action} + F1_{micro}^{action})/4. \quad (13)$$

Implementation Details. For each video, each frame is resized into 224 × 224. For model training, we adopt the SGD optimizer with a learning rate of 0.0075, a momentum of 0.9, a weight decay of 1e-4, and a batch size of 10. The learning rate is reduced by a factor of 10 at the 15th and 30th epochs, and the model is trained with 40 epochs. In Eq. 10, we set $\gamma = 1$ for the RGB branch and $\gamma = 5$ for the Pose branch. In Eq. 12, we set $\beta = 5$ for the RGB branch and $\beta = 5$ for the Pose branch.

4.2 Comparison with the State-Of-The-Art

As depicted in Tab. 1, we compare our method with classical methods in 5 groups, namely, **1) 2D CNN based** (i.e., TSN (2018), TIN (2020), TSM (2019) and MANet (2024a)), **2) 3D CNN based** (i.e., C3D (2015), I3D (2017) and SlowFast (2019)), **3) Transformer-based** (i.e., VSwiT (2022)),

¹<https://github.com/VUT-HFUT/Micro-Action>

λ	Body		Action			Body		Action		All
	Top-1	Top-1	Top-5	$F1_{macro}$	$F1_{micro}$	$F1_{macro}$	$F1_{micro}$	$F1_{macro}$	$F1_{micro}$	$F1_{mean}$
0	81.76	66.34	91.80	76.26	81.76	52.69	66.34	69.26		
0.1	81.99	66.27	91.14	76.65	81.99	52.28	66.27	69.30		
10	81.95	66.29	91.48	76.69	81.95	51.82	66.29	69.19		
1	82.30	66.74	91.75	77.02	82.30	53.83	66.74	69.97		

Table 2: Ablation for λ in \mathcal{L}_{HP} on the MA-52 dataset.

α_1	α_2	α_3	Body		Action			Body		Action		All
			Top-1	Top-1	Top-5	$F1_{macro}$	$F1_{micro}$	$F1_{macro}$	$F1_{micro}$	$F1_{macro}$	$F1_{micro}$	$F1_{mean}$
1	1	1	82.04	66.41	91.64	76.59	82.04	53.44	66.61	69.67		
1	0.1	0.1	81.90	66.47	91.51	76.09	81.90	52.79	66.47	69.31		
0.1	1	0.5	81.92	66.13	91.60	76.31	81.92	52.49	66.13	69.21		
1	0.5	1	81.83	66.06	91.34	75.89	81.83	52.23	66.06	69.00		
1	0.1	0.5	82.01	66.67	91.48	76.32	82.01	53.00	66.67	69.50		
1	0.5	0.1	82.30	66.74	91.75	77.02	82.30	53.83	66.74	69.97		

Table 3: Ablation results for different weights of ambiguous sample sets on the MA-52 dataset. α_1 , α_2 and α_3 are weights for \mathbb{FN}_{A1} , \mathbb{FN}_{A2} and \mathbb{FN}_{A3} , respectively.

TimesFormer (2021) and Uniformer (2023c)), **4) GCN based** (STGCN (2018), 2s-AGCN (2019) and CTRGCN (2021)), and **5) Dual-modality based** (B2C-AFM (2023) and PoseConv3D (2022)). In 2D CNN based methods, MANet achieves 65.59% on $F1_{mean}$, while the CTRGCN (Chen et al. 2021) in GCN-based methods only achieves 56.48%, which is 9.11% lower than MANet. This is due to the fact that skeleton information is more susceptible to the quality of key points obtained by the pose estimator, and micro-actions usually have short durations and low action amplitudes, which leads to the difficulty in improving the performance of the GCN-based methods. On the action-level $F1_{micro}$ metric, PoseConv3D (RGB) and PoseConv3D (Pose) achieve 56.30% and 50.93%, and the performance can reach 63.52% with two-modality fusion. Whereas our method achieves 60.03% and 56.51% for RGB and Pose branches, the performance can be improved to 66.74% after fusion. This proves that the proposed prototype calibration method can be effective for multiple modalities. The $F1_{macro}$ in the action-level categories of our method reaches 53.83%, which is 6.63% higher than PoseConv3D, indicating that we are able to distinguish the ambiguous samples effectively. It proves the proposed method can improve the recognition accuracy in micro-action across different action-level categories.

4.3 Ablation Studies

Action-tree Hierarchy Modeling. The body-level micro-actions are represented as parent nodes in the action-tree, while the action-level micro-actions are depicted as child nodes positioned under their corresponding parent nodes, which models the hierarchical structural relationship. In Tab. 2, we conduct ablation experiments with the hyper-parameter λ in Eq. 1 on the RGB branch. When $\lambda = 0$, it signifies that the action-level probabilities are not influenced by the action-tree structure. Regarding the $F1_{macro}$ score for action-level, $\lambda = 1$ yields a 1.14% improvement compared to $\lambda = 0$. This proves that the proposed action-tree structure effectively improves the accuracy of identifying action-level micro-actions.




Figure 3: Top-1 Accuracy (%) on ambiguous micro-actions for the MA-52 dataset.

Ambiguous Samples Contrastive Calibration. Here, we explore the impact of hyper-parameters on ambiguous sample sets \mathbb{FN}_{A1} , \mathbb{FN}_{A2} , and \mathbb{FN}_{A3} as stated in prototype identification. Each of these sets represents distinct types of ambiguous samples. For \mathbb{FN}_{A1} , the model correctly captures most of the general features of the sample, but makes errors in the action-level category. Hence, we assign a higher weight to emphasize the model’s inaccuracies in the finer details. For \mathbb{FN}_{A2} , while the model fails at both scales, errors at the higher are more critical than those at \mathbb{FN}_{A1} . For \mathbb{FN}_{A3} , these ambiguous samples result from the model’s misreading of the overall characteristics, so we assign them lower weights. We experiment with various weight configurations to investigate the importance of different types of ambiguous sample sets, as shown in Tab. 3. The optimal result is achieved when weights for \mathbb{FN}_{A1} , \mathbb{FN}_{A2} , and \mathbb{FN}_{A3} decreased step by step. This indicates that the \mathbb{FN}_{A1} set significantly contributes to the calibration of ambiguous samples. Despite the model accurately classifying overall actions, it struggles with similar micro-action categories, necessitating further calibration. Additionally, when body-level labels are incorrectly predicted, including \mathbb{FN}_{A2} and \mathbb{FN}_{A3} samples, they are often attributed to factors like data noise or sample imbalance. Compared to the \mathbb{FN}_{A1} set, these samples are less reliable and of lower quality. Therefore, they require further calibration with a separate cluster center and a relatively low-weight assignment.

4.4 Ambiguous Micro-actions Statistic Analysis

To evaluate the effectiveness of PCAN on ambiguous micro-actions recognition, we split the action categories in MA-52 into 3 distinct difficulty levels, *i.e.*, “easy”, “medium”, and “hard”. Concretely, based on the recognition results of PoseConv3D (Duan et al. 2022), we gather action with accuracy lower than 50% as the hard level, between 50% and 60% as the medium level, and over 60% as the easy level. The comparison results are shown in Fig. 3. Due to the high similarity among ambiguous micro-actions, many micro-action categories are usually hard to recognize (Guo et al. 2024a). Compared to baseline, our method shows a significant improvement of 11.8% in the hard-level categories, demonstrating its ability to distinguish ambiguous samples. Additionally, our method also exhibits improvements of 5.8% and 0.6% on the medium- and easy-level categories, indicating high recognition accuracy.

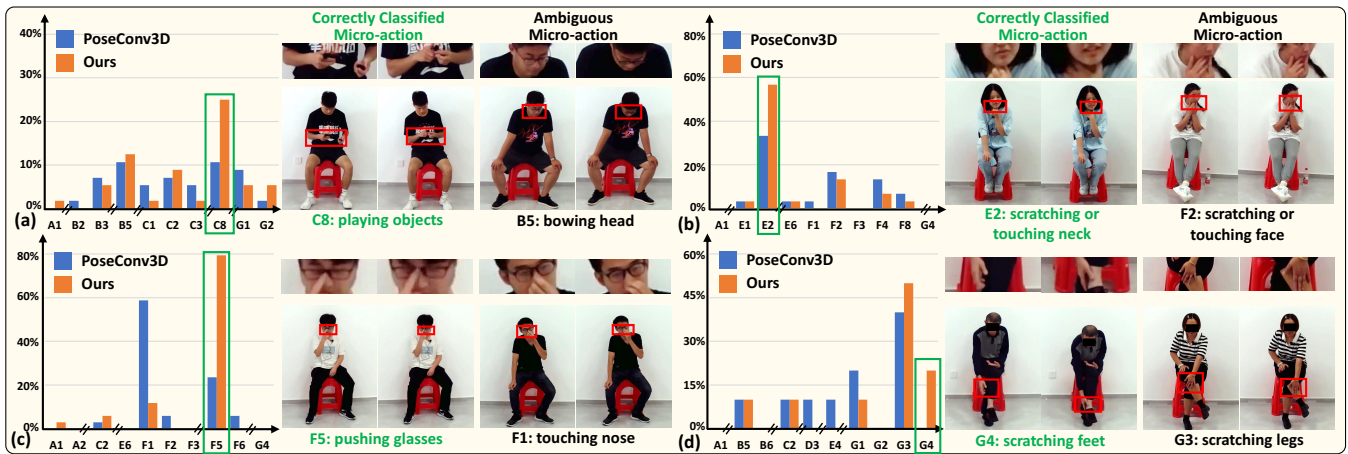


Figure 4: Comparative analysis of the PoseConv3D (Duan et al. 2022) and ours PCAN. The green categories are the micro-action categories being explored. Our PCAN method demonstrates an excellent ability to categorize ambiguous action categories. The proposed PCAN exhibits robust performance at both body-level and action-level.




Figure 5: The group-wise Top-1 accuracy improvement (%) of our method compared to PoseConv3D on the MA-52 dataset.

4.5 Qualitative Analysis

Visualization of Recognition Results. As shown in Fig 4, we visualize the predicted distributions of four ambiguous categories. In case (a), for the micro-action of “playing object”, the focus should be on hand movements, whereas the baseline method places excessive emphasis on the associated head movement of “bowing head”. Cases (b) and (c) are two ambiguous categories with extremely similar appearances. In contrast, our PCAN method of utilizing body- and action-level cascade prototypes can effectively constrain and calibrate the ambiguous samples, prompting the model to distinguish similar samples more effectively. Case (d) shows that although our method failed to classify this particular instance, it still achieves great improvement in all samples. We attribute this misclassification to the high similarity in motion patterns between these two samples. Effectively identifying and calibrating ambiguous samples in micro-action recognition tasks still remains a substantial challenge.

Visualization of Accuracy Distribution. To evaluate the recognition accuracy of PCAN on action-level categories, we also give detailed statistics of the top-1 accuracy improvement of PCAN compared to PoseConv3D. As shown in Fig. 5, there is an improvement of 35 categories in action-level, with the same accuracy in 8 categories and a slight decrease in 9 categories.

All body-level categories in the MA-52 dataset are boosted in subcategories. Specifically, 75% action-level categories are boost under “Body” label, 66.7% in “Head”, 84.6% in “Upper limb”, 60% in “Head-hand” and 75% in “Leg-hand”. It is obvious that PCAN improves by more than 10% on 16 action-level categories compared to PoseConv3D, with the highest improvement of 56% on the “pushing glasses” category.

5 Conclusion

In this paper, we presented a prototypical calibrating ambiguous network that discovers and calibrates ambiguous samples for MAR via the body- and action-level prototypes. The ambiguous samples are categorized into false negatives (FN) and false positives (FP), and we apply a hierarchical contrastive refinement to calibrate these ambiguous samples by adjusting their proximity to prototypes. We also presented to amplify the diversity between prototypes to enhance the model’s discriminativity. Experimental results on the benchmark dataset showed our method corrects a significant number of hard ambiguous samples, underscoring the accuracy gains from prototype calibration. We believe that this work can inspire further research in other recognition tasks with a multitude of intricate samples.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (62272144, 72188101, 62020106007, U20A20183, and 62472381), the Major Project of Anhui Province (202203a05020011), the Fundamental Research Funds for the Zhejiang Provincial Universities (226-2024-00208), and the Earth System Big Data Platform of the School of Earth Sciences, Zhejiang University.

References

- Aviezer, H.; Trope, Y.; and Todorov, A. 2012. Body cues, not facial expressions, discriminate between intense positive and negative emotions. *Science*, 338(6111): 1225–1229.
- Balazia, M.; Müller, P.; Táncoz, Á. L.; Liechtenstein, A. v.; and Bremond, F. 2022. Bodily behaviors in social interaction: Novel annotations and state-of-the-art evaluation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 70–79.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding? In *Proceedings of International Conference on Machine Learning*, 813–824.
- Camporese, G.; Bergamo, A.; Lin, X.; Tighe, J.; and Modolo, D. 2023. Early Action Recognition with Action Prototypes. *arXiv preprint arXiv:2312.06598*.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6299–6308.
- Chen, H.; Shi, H.; Liu, X.; Li, X.; and Zhao, G. 2023. SMG: A micro-gesture dataset towards spontaneous body gestures for emotional stress state analysis. *International Journal of Computer Vision*, 131(6): 1346–1366.
- Chen, Y.; Zhang, Z.; Yuan, C.; Li, B.; Deng, Y.; and Hu, W. 2021. Channel-wise topology refinement graph convolution for skeleton-based action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13359–13368.
- Dai, G.; Shu, X.; Yan, R.; Huang, P.; and Tang, J. 2023. Slowfast Diversity-aware Prototype Learning for Egocentric Action Recognition. In *Proceedings of the 31st ACM International Conference on Multimedia*, 7549–7558.
- Duan, H.; Zhao, Y.; Chen, K.; Lin, D.; and Dai, B. 2022. Revisiting skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2969–2978.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2019. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6202–6211.
- Guo, D.; Li, K.; Hu, B.; Zhang, Y.; and Wang, M. 2024a. Benchmarking Micro-action Recognition: Dataset, Methods, and Applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(7): 6238–6252.
- Guo, D.; Li, X.; Li, K.; Chen, H.; Hu, J.; Zhao, G.; Yang, Y.; and Wang, M. 2024b. MAC 2024: Micro-Action Analysis Grand Challenge. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 11304–11305.
- Guo, F.; Jin, T.; Zhu, S.; Xi, X.; Wang, W.; Meng, Q.; Song, W.; and Zhu, J. 2023. B2c-afm: Bi-directional co-temporal and cross-spatial attention fusion model for human action recognition. *IEEE Transactions on Image Processing*, 32: 4989–5003.
- Hao, X.; Li, J.; Guo, Y.; Jiang, T.; and Yu, M. 2021. Hypergraph neural network for skeleton-based action recognition. *IEEE Transactions on Image Processing*, 30: 2263–2275.
- Li, K.; Guo, D.; Chen, G.; Liu, F.; and Wang, M. 2023a. Data Augmentation for Human Behavior Analysis in Multi-Person Conversations. In *Proceedings of the 31st ACM International Conference on Multimedia*, 9516–9520.
- Li, K.; Guo, D.; Chen, G.; Peng, X.; and Wang, M. 2023b. Joint skeletal and semantic embedding loss for micro-gesture classification. *arXiv preprint arXiv:2307.10624*.
- Li, K.; Guo, D.; Liu, P.; Chen, G.; and Wang, M. 2024a. Mmad: Multi-label micro-action detection in videos. *arXiv preprint arXiv:2407.05311*.
- Li, K.; Peng, X.; Guo, D.; Yang, X.; and Wang, M. 2024b. Repetitive Action Counting with Hybrid Temporal Relation Modeling. *arXiv preprint arXiv:2412.07233*.
- Li, K.; Wang, Y.; Zhang, J.; Gao, P.; Song, G.; Liu, Y.; Li, H.; and Qiao, Y. 2023c. Uniformer: Unifying convolution and self-attention for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10): 12581–12600.
- Li, S.; Xu, J.; Xu, X.; Shen, P.; Li, S.; and Hooi, B. 2021. Spherical confidence learning for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15629–15637.
- Li, Y.-L.; Liu, X.; Lu, H.; Wang, S.; Liu, J.; Li, J.; and Lu, C. 2020. Detailed 2d-3d joint representation for human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10166–10175.
- Lin, J.; Gan, C.; and Han, S. 2019. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7083–7093.
- Liu, P.; Wang, F.; Li, K.; Chen, G.; Wei, Y.; Tang, S.; Wu, Z.; and Guo, D. 2024. Micro-gesture Online Recognition using Learnable Query Points. *arXiv preprint arXiv:2407.04490*.
- Liu, W.; Jia, X.; Zhong, X.; Jiang, K.; Yu, X.; and Ye, M. 2025. Dynamic and static mutual fitting for action recognition. *Pattern Recognition*, 157: 110948.
- Liu, X.; Shi, H.; Chen, H.; Yu, Z.; Li, X.; and Zhao, G. 2021. imigue: An identity-free video dataset for micro-gesture understanding and emotion analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10631–10642.
- Liu, Z.; Ning, J.; Cao, Y.; Wei, Y.; Zhang, Z.; Lin, S.; and Hu, H. 2022. Video swin transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3202–3211.

- Mou, Y.; Jiang, X.; Xu, K.; Sun, T.; and Wang, Z. 2024. Compressed video action recognition with dual-stream and dual-modal transformer. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(5): 3299–3312.
- Shao, H.; Qian, S.; and Liu, Y. 2020. Temporal interlacing network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 11966–11973.
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12026–12035.
- Shiffrar, M.; Kaiser, M. D.; and Chouhourelou, A. 2011. Seeing human movement as inherently social. *The Science of Social Vision*, 248–264.
- Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 4489–4497.
- Wang, H.; Chu, H.; Siming, F.; Liu, Z.; and Hu, H. 2022. Renovate yourself: Calibrating feature representation of misclassified pixels for semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2450–2458.
- Wang, L.; and Koniusz, P. 2023. 3mformer: Multi-order multi-mode transformer for skeletal action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5620–5631.
- Wang, L.; Xiong, Y.; Wang, Z.; Qiao, Y.; Lin, D.; Tang, X.; and Van Gool, L. 2018. Temporal segment networks for action recognition in videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(11): 2740–2755.
- Wang, X.; Zhang, S.; Cen, J.; Gao, C.; Zhang, Y.; Zhao, D.; and Sang, N. 2024. CLIP-guided prototype modulating for few-shot action recognition. *International Journal of Computer Vision*, 132(6): 1899–1912.
- Wang, X.; Zhu, L.; Wang, H.; and Yang, Y. 2021. Interactive prototype learning for egocentric action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8168–8177.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3733–3742.
- Xia, S.; Lv, J.; Xu, N.; Niu, G.; and Geng, X. 2023. Towards effective visual representations for partial-label learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15589–15598.
- Yan, S.; Xiong, Y.; and Lin, D. 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Ye, M.; Zhang, X.; Yuen, P. C.; and Chang, S.-F. 2019. Unsupervised embedding learning via invariant and spreading instance feature. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6210–6219.
- Zhou, H.; Liu, Q.; and Wang, Y. 2023. Learning discriminative representations for skeleton based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10608–10617.
- Zhu, X.; Toisoul, A.; Perez-Rua, J.-M.; Zhang, L.; Martinez, B.; and Xiang, T. 2021. Few-shot action recognition with prototype-centered attentive learning. *arXiv preprint arXiv:2101.08085*.
- Zhu, Y.; Huang, G.; Xu, X.; Ji, Y.; and Shen, F. 2022. Selective hypergraph convolutional networks for skeleton-based action recognition. In *Proceedings of the 2022 International Conference on Multimedia Retrieval*, 518–526.