

# Skeleton-Based Human Action Recognition with Noisy Labels

Yi Xu<sup>1</sup>, Kunyu Peng<sup>1,\*</sup>, Di Wen<sup>1</sup>, Ruiping Liu<sup>1</sup>, Junwei Zheng<sup>1</sup>, Yufan Chen<sup>1</sup>, Jiaming Zhang<sup>1</sup>,  
Alina Roitberg<sup>2</sup>, Kailun Yang<sup>3,4</sup>, and Rainer Stiefelhagen<sup>1</sup>

**Abstract**—Understanding human actions from body poses is critical for assistive robots sharing space with humans in order to make informed and safe decisions about the next interaction. However, precise temporal localization and annotation of activity sequences is time-consuming and the resulting labels are often noisy. If not effectively addressed, label noise negatively affects the model’s training, resulting in lower recognition quality. Despite its importance, addressing label noise for skeleton-based action recognition has been overlooked so far. In this study, we bridge this gap by implementing a framework that augments well-established skeleton-based human action recognition methods with label-denoising strategies from various research areas to serve as the initial benchmark. Observations reveal that these baselines yield only marginal performance when dealing with sparse skeleton data. Consequently, we introduce a novel methodology, NoiseEraSAR, which integrates global sample selection, co-teaching, and Cross-Modal Mixture-of-Experts (CM-MOE) strategies, aimed at mitigating the adverse impacts of label noise. Our proposed approach demonstrates better performance on the established benchmark, setting new state-of-the-art standards. The source code for this study will be made accessible at <https://github.com/xuyizdy/NoiseEraSAR>.

## I. INTRODUCTION

Skeleton-based human action recognition is vital in robotics, particularly in human-robot interaction, enabling more natural and intuitive communication [1], [2]. It plays a crucial role in surveillance [3], [4], allowing robots to identify and respond to emergencies or unsafe behaviors. Skeleton-based human action recognition also enables personalized services in areas like healthcare and fitness by analyzing and adapting to individual human movements [5], [6]. Additionally, it improves efficiency and safety in collaborative industrial environments by enabling robots to understand and predict human actions [7], [8]. The success of deep learning for human action recognition heavily relies on data

This work was supported in part by the SmartAge project sponsored by the Carl Zeiss Stiftung (P2019-01-003; 2021-2026), by the project served to prepare the SFB 1574 Circular Factory for the Perpetual Product (project ID: 471687386), approved by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) with a start date of April 1, 2024, and in part by the BMBF through a fellowship within the IFI program of the German Academic Exchange Service (DAAD), in part by the HoreKA@KIT supercomputer partition, and in part by Hangzhou Surlmage Technology Company Ltd. The authors also acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant INST 35/1597-1 FUGG.

\*Corresponding author. (Email: kunyu.peng@kit.edu.)

<sup>1</sup>The authors are the Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Germany.

<sup>2</sup>The author is with the Institute for Artificial Intelligence, University of Stuttgart, Germany.

<sup>3</sup>The author is with the School of Robotics, Hunan University, China.

<sup>4</sup>The author is also with the National Engineering Research Center of Robot Visual Perception and Control Technology, Hunan University, China.




Fig. 1. An overview of our task setting. We randomly inject asymmetric label noise into the training set according to a predefined noise ratio. On the right-hand side, we deliver the comparison of the performances on the test set with correct labels in terms of the Cross-Subject (X-Sub) and the Cross-View (X-View) settings, where our approach shows the best performance.

annotation for supervised training. However, the cost-effective and decentralized manual labeling work on crowdsourcing platforms introduces significant concerns related to the quality of the annotation [9]. Additionally, the development of large models necessitates diverse and extensive datasets with large amounts of labels. The quality of the labels will directly affect the model training [10]–[12]. Due to the lack of visual appearance, skeleton data is hard to annotate compared with video data, which may result in more label noise. Therefore, designing models that can alleviate the negative effect brought by label noise is a challenge for skeleton-based action recognition in real-world scenarios. In recent literature concerning the mitigation of label noise, various methodologies have been proposed, encompassing techniques such as sample selection and noise modeling [13]–[15]. These methodologies have been further extended by some researchers who integrate sample selection within a semi-supervised learning paradigm [16], categorizing clean annotated data as labeled and unclean annotated data as unlabeled, thereby leveraging the strengths of semi-supervised learning techniques. Moreover, a distinct subset of studies has focused on the explicit modeling of data labels, primarily through the construction of transition matrices, to systematically address the issue of label noise [17]. Nevertheless, the literature lacks investigations specifically targeting skeleton-based human action recognition in the context of noisy label scenarios, whereas the majority of existing studies are concentrated on image classification tasks.

The objective of this paper is to conduct the first work on skeleton-based action recognition under a noisy label setting and introduce a novel approach for mitigating label noise within this new task. To conduct the first testbed, two well-established label denoising approaches are implemented into skeleton-based action recognition pipeline, which are SOP [16] and NPC [17], where we firstly choose CTR-

GCN [18] as backbone for skeleton feature learning and conduct experiments on different label noise ratios. We find these two label-denoising methods can not well address the label noise scenario for the skeleton-based action recognition task, which is mostly due to the sparsity of the skeleton data. We thereby propose a new approach in this field to handle the aforementioned challenges, which is named as NoiseEraSAR. This approach synergizes methodologies pertaining to sample selection and multi-modality fusion. The sample selection mechanism reduces the propensity for model overfitting to noisy data by employing a co-selection strategy between two models. Concurrently, the multi-modality fusion aspect addresses the distinctive attributes of skeletal data, harnessing complementary insights from multi-modal sources to fortify training resilience against label noise. This is achieved through the new proposed Cross-Modal Mixture-of-Experts (CM-MoE) framework, leveraging a spatio-temporal graph gate network. Our proposed method achieves state-of-the-art performances on the constructed testbed, as shown in Figure 1.

Our contributions are summarized as follows:

- We open the vistas for Skeleton-based Human Action Recognition (SHAR) under noisy labels and create a new benchmark by implementing the well-established methods from other domains, which focus on the label noise problem, into the skeleton-based action recognition field.
- We introduce a new method, NoiseEraSAR, to address label noise in SHAR. It combines the principles of cross-training and the Cross-Modal Mixture-of-Experts (CM-MoE) to formulate an effective training framework for SHAR under noisy labels. We also conduct comprehensive ablation experiments to verify the effectiveness of different components of the method.
- Experimental results demonstrate the effectiveness of the proposed method across different label noise ratios. NoiseEraSAR clearly surpasses state-of-the-art methods in various label noise settings. Our method achieves 74.9% and 79.5% of accuracy on cross-subject and cross-view evaluations on the NTU-60 dataset under 80% label noise.

## II. RELATED WORK

### A. Skeleton-based Human Action Recognition

Skeleton-based human action recognition is exceptionally suitable for various applications that demand high reliability and efficiency, due to its strong environmental robustness and extensive capability for subject generalization.

Early studies leaned towards CNN-based methods [19]–[21], leveraging their superior hierarchical feature learning, alongside RNN-based methods [22]–[24], recognized for modeling temporal dynamical behavior in sequences. Recent approaches have shifted to transformer-based methods [25]–[27], acclaimed in capturing long-range dependencies and facilitating parallel processing for efficiency. To adeptly utilize skeletal geometric information, GCN-based methods [18],

[28]–[30] focus on topology modeling as a fundamental design principle. ST-GCN [28], which initially utilized fixed graph convolution, effectively models dynamic spatial-temporal relationships using a predefined skeleton topology. Despite its efficiency, the reliance on a fixed graph structure introduces limitations in adaptability and precision for recognizing dynamic and diverse human actions. Dynamic GCN [29] introduces adaptive graph topologies that allow for more flexible and accurate modeling of complex human actions. CTR-GCN [18] innovates by refining the graph topology at a channel-wise aspect, enabling more precise and adaptive capture of spatial-temporal relationships in human actions. We construct our main benchmark on CTR-GCN [18] to achieve the backbone unification while conducting ablation for different skeleton-based action recognition backbones on HD-GCN [26] and ST-GCN [28].

### B. Noisy Labels Learning

Noisy labels learning focuses on efficiently training models on datasets with inaccurate labels, aiming to enhance model robustness and accuracy despite the presence of label noise.

Strategies to enhance the robustness of models can be roughly classified into five categories [31], *i.e.*, (i) Modifying loss for noisy labels, which can be further achieved through four aspects: (a) Estimating the noise label transition matrix to adjust the loss function [32]–[36]; (b) Re-weighting individual sample losses by reducing the weight of samples likely mislabeled [37]; (c) Refurbishing the labels that are presumed noisy [10]–[12]; (d) Adjusting with the optimal rule deriving from meta-learning [38]–[40]. (ii) Developing loss functions robustness to noise [41]–[43]: These are aiming to inherently tolerate label inaccuracies without requiring a noise transition matrix. (iii) Applying regularization methods [44]–[46]: These methods aim to leverage regularization techniques to prevent over-fitting to noisy labels, thereby enhancing the overall robustness and generalization of the models. (iv) Enhancing robustness with architecture [47]: Integrating a specialized noise adaptation structure atop the base deep neural network (DNN) facilitates understanding of the label transition mechanism, thereby enhancing training robustness against the label noise. (v) Implementing dynamic sample selection [13]–[15]: This approach starts with selecting a subset of clean samples and progressively incorporates incorrectly labeled samples in a semi-supervised learning framework. Notably, our model uses a cross-training scheme inspired by [13] to identify small-loss (and thus likely “clean”) samples for further training.

## III. BENCHMARK

### A. Label Noise

In our benchmark, we establish an initial testing framework specifically designed for skeleton-based human action recognition, incorporating the challenge of noisy labels. Our goal is to assess how well a model can perform under various levels of label inaccuracies. This involves deliberately mislabeling a certain percentage of the samples while making sure these incorrect labels are still among those available in the dataset

utilized. We distinguish between two main types of noise: **Symmetric noise** and **Asymmetric noise**. **Symmetric noise** is characterized by the random alteration of each label to any other class with a uniform probability, as described in [48]. In contrast, **Asymmetric noise** involves changing labels to different classes based on distinct probabilities. For example, if categories A and B are typically more confusable, we might specify that a mislabeled sample from class A has a 0.6 probability of being wrongly assigned to class B and a 0.4 chance of being attributed to any other category.

Implementing **Asymmetric noise** necessitates a pre-established understanding of the likelihood of confusion between specific categories, informed by real-world data labeling experiences. In this study, our focus is restricted to scenarios with **Symmetric noise**. Studying symmetric label noise is crucial as it provides a foundational understanding of how random errors in labeling affect model training, setting the stage for tackling the more complex and realistic scenarios of asymmetric label noise. We have introduced symmetric noise into the training set at various levels: 20%, 40%, 50%, and 80%. To ensure the integrity of our evaluation, we do not introduce label noise into the testing set.

### B. Baseline

A crucial aspect of our benchmark is the selection of the baseline methods. We opt for two established techniques as our baselines due to their proven effectiveness in handling noisy labels across different domains. These methods serve as a standard for comparing the performance of our model, and we conduct experiments based on them. We choose CTR-GCN [18] as the feature extraction backbone to construct our main test bed for skeleton-based human action recognition under different label noise ratios while leveraging the other two backbones, *i.e.*, HD-GCN [26] and ST-GCN [28], for the evaluation of the cross-backbone generalizability of different denoising methods.

**Sparse Over-Parameterization (SOP)** [16]. SOP is an approach for robust training against label noise, which models the label noise via another sparse over-parameterization term and exploits implicit algorithmic regularizations to separate the noise. **Noisy Prediction Calibration (NPC)** [17]. NPC estimates the explicit transition from a noisy prediction to a true latent class via utilizing a deep generative model.

### C. Dataset

**NTU RGB+D.** NTU RGB+D (NTU-60) [49] is a large 3D human action recognition dataset, which includes 56,880 samples covering a total of 60 action classes. Each sample is captured from three different viewpoints, including one action type and up to two subjects. The dataset comprises various modalities, including skeleton, IR, and RGB. For our experiments, we select the skeleton temporal modality, which involves a total of 25 joint points. The 3D coordinates of each joint point are recorded for each frame as the original features. The NTU-60 dataset is a popular choice for skeleton-based human action recognition problems and is therefore selected as the primary dataset for our experiments.

## IV. METHODOLOGY

### A. Overview

We first deliver an overview of the contributed new method in this subsection. As illustrated in Figure 2, our method, NoiseEraSAR, contains two main phases: The pretraining phase and the fine-tuning phase. Especially, three components are involved: cross-training, global sample selection, and a Cross-Modal Mixture of Expert (CM-MoE) technique. Firstly, data from three modalities (joint, bone, and motion represented in green, blue, and purple) are extracted for each sample. For each modality, two models with the same structure are simultaneously trained using the co-teaching method, where these models select samples whose labels are worthy of belief according to the losses for each other in every epoch (*e.g.*,  $Model_{j_1}$  and  $Model_{j_2}$  in Figure 2 are the peer networks of the joint modality). Next, in the global sample selection process, models from the three modalities individually choose a defined percentage of clean labeled samples from the whole training dataset and the final sample set is constructed by taking the union of these three selected sets considering the samples. The final step involves employing the CM-MoE system while relying on the union of the samples from the selected three clean labeled sets. The triple networks pre-trained in the first step will be combined with a gate network constructed by spatiotemporal graph convolution that connects the starting input layer and the SoftMax layer (as shown on the right side of Figure 2). The gate network dynamically adjusts the weights of the outputs from different models across modalities. The weighted average of the three scores will be used as the prediction of our final method. Below, we present a detailed overview of the different components comprising our method.

### B. Modalities

We leverage three skeletal modalities: joint, bone, and motion. Both bone and motion modalities are derived from the original joint coordinate data.

**Joint modality.** Joint modality is the raw form of skeletal data, which mainly records the 3D coordinates ( $\mathbf{x}$ ,  $\mathbf{y}$ , and  $\mathbf{z}$ ) of each joint point at different times (frames). In particular, the NTU-60 dataset contains a total of 25 joints with their coordinates at each frame. Then we could describe the joint data of joint index  $i$  at time  $t$  as:

$$j_{i,t} = (x_{i,t}, y_{i,t}, z_{i,t}) \quad (1)$$

**Bone modality.** A bone vector can be described by the head and tail joint points. For example, while processing with the NTU-60 dataset, we calculate a total of 25 bones (*i.e.* 25 joint pairs). Therefore, we can simply calculate the bone data ( $b_{i,j}$ ) based on the joint pairs ( $x_{j,t}$  and  $x_{i,t}$ ):

$$b_{i,j} = (x_{j,t} - x_{i,t}, y_{j,t} - y_{i,t}, z_{j,t} - z_{i,t}) \quad (2)$$

**Motion modality.** The motion modality is described by the difference between the same joints in two consecutive time

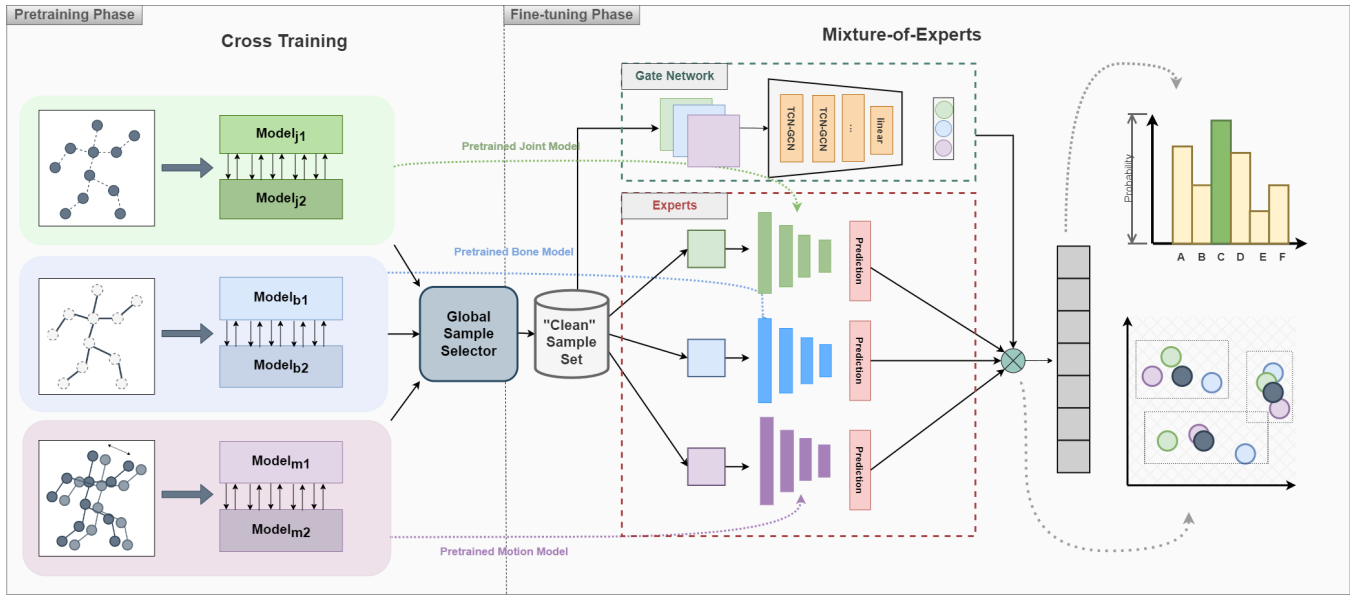


Fig. 2. Overview of the method, NoiseEraSAR: In the pre-training phase, the proposed method first trains special models for joint, bone, and motion modalities by using a cross training method. The small clean dataset is generated from the pre-trained models by evaluating the loss value, and it is fed into the Cross-Modal Mixture-of-Experts (CM-MoE). In the fine-tuning phase, the gate network is added to control the weights of each expert and assists the CM-MoE.

points (frames). The motion data of joint  $i$  between the time  $t$  and  $t + 1$  can be represented as follows:

$$m_{i,t,t+1} = (x_{i,t+1} - x_{i,t}, y_{i,t+1} - y_{i,t}, z_{i,t+1} - z_{i,t}) \quad (3)$$

### C. Cross Training

The first part of the method is the cross-training process under three different modalities (joint, bone, and motion). Inspired by the method used in co-training [13], we simultaneously train two networks ( $\mathbf{N}_1$  and  $\mathbf{N}_2$ ) with identical structure. The small-loss method is applied here, where a certain proportion of samples with smaller losses within each batch is identified as clean samples at each epoch, and only these samples participate in the back-propagation process. Consequently, at each epoch, both networks select samples for their counterpart based on the small-loss criterion.

If in each batch of the epoch ( $T$ ), we have the mini-batch dataset  $\bar{D}$ , The process of sample selection in each epoch can be described as follows:

$$R(T) = 1 - \min \left\{ \frac{T}{T_{in}} r, r \right\} \quad (4)$$

$$\bar{D}_1 = \arg \min_{D': |D'| \geq R(T)|\bar{D}|} \ell(\mathbf{N}_1, D') \quad (5)$$

$$\bar{D}_2 = \arg \min_{D': |D'| \geq R(T)|\bar{D}|} \ell(\mathbf{N}_2, D') \quad (6)$$

where  $R(T)$  is the the ratio of selected samples and  $r$  is the noise ratio. And the  $T_{in}$  represents the fixed number of epochs and  $\ell$  is the cross-entropy loss function. Finally, we obtain two set of samples ( $\bar{D}_1$  and  $\bar{D}_2$ )

Then, each network updates its parameters by the selected samples from the peer as follows:

$$w_1 = w_1 - \eta \nabla \ell(\mathbf{N}_1, \bar{D}_2) \quad (7)$$

$$w_2 = w_2 - \eta \nabla \ell(\mathbf{N}_2, \bar{D}_1) \quad (8)$$

where  $w_1$  and  $w_2$  represent the trainable parameters of the  $\mathbf{N}_1$  and  $\mathbf{N}_2$ .  $\mathbf{N}_1$  represents  $Model_{\phi_1}$  and  $\mathbf{N}_2$  represents  $Model_{\phi_2}$ , while  $\phi \in [j, b, m]$ , as shown in Figure 2.

Throughout the cross-training process, the tendency to memorize noisy labels is mitigated, since the peer networks have different learning abilities and can filter out different types of introduced errors by noisy labels.

After the last epoch, the network with superior accuracy is selected as the final model for this modality. Therefore, we obtain three pre-trained networks ( $\mathbf{N}_{joint}$ ,  $\mathbf{N}_{bone}$ , and  $\mathbf{N}_{motion}$ ).

### D. Global Sample Selection

After the pre-training process with the cross-training method for each modality, a global sample selector is proposed to extract more representative samples for the next modality fusion process. Given that the entire dataset includes a significant proportion of noisy labels, sample selection proves beneficial for further training, both in terms of denoising and improving calculation speed. If the ratio of instances to be selected is set to  $p$  and training set  $D$  is given, then we have the selected samples from three networks as follows:

$$D_{joint} = \arg \min_{D': |D'| \geq p|D|} \ell(\mathbf{N}_{joint}, D') \quad (9)$$

$$D_{bone} = \arg \min_{D': |D'| \geq p|D|} \ell(\mathbf{N}_{bone}, D') \quad (10)$$

$$D_{motion} = \arg \min_{D': |D'| \geq p|D|} \ell(\mathbf{N}_{motion}, D') \quad (11)$$

To obtain the benefits of triple modalities and to avoid the oversampling of the examples with lower learning difficulty from single networks, the final clean set  $D_c$  is the global union set of samples from three networks:

$$D_c = D_{joint} \cup D_{bone} \cup D_{motion} \quad (12)$$

TABLE I  
THE COMPARISON RESULTS OF SKELETON-BASED HUMAN ACTION RECOGNITION ON THE NTU-60 [49] DATASET.

Method/Noise ratio	Symmetric Noise							
	20%		40%		50%		80%	
	X-Sub	X-View	X-Sub	X-View	X-Sub	X-View	X-Sub	X-View
CTR-GCN [50]	86.8%	90.9%	83.7%	88.2%	81.7%	85.7%	64.8%	64.6%
CTR-GCN [50] + SOP [16]	83.7%	88.3%	82.0%	87.3%	81.0%	85.6%	69.2%	68.2%
CTR-GCN [50] + NPC [17]	86.6%	90.8%	83.2%	87.4%	82.4%	85.5%	68.9%	70.2%
NoiseEraSAR(ours)	<b>90.6%</b>	<b>95.3%</b>	<b>89.0%</b>	<b>93.1%</b>	<b>88.5%</b>	<b>90.5%</b>	<b>74.9%</b>	<b>79.5%</b>

TABLE II  
ABLATION STUDY OF THE PROPOSED COMPONENTS OF OUR METHOD, INCLUDING CROSS-TRAINING, GLOBAL SAMPLE SELECTOR, AND CROSS-MODAL MIXTURE-OF-EXPERTS TECHNIQUE (CM-MoE).

Methods	Modality	Symmetric Noise 80%	
		X-Sub	X-View
CTR-GCN [50]+Cross-training	joint	70.8%	74.9%
CTR-GCN [50]+Cross-training	bone	65.2%	69.0%
CTR-GCN [50]+Cross-training	motion	67.5%	69.6%
CTR-GCN [50]+Cross-training+ensemble	joint+bone+motion	72.5%	78.9%
<b>CTR-GCN [50]+Cross-training+CM-MoE(ours)</b>	joint+bone+motion	<b>74.9%</b>	<b>79.5%</b>

### E. Cross-Modal Mixture of Expert (CM-MoE)

In the domain of skeleton-based human action recognition, models trained on joint, bone, and motion modalities learn complementary patterns from the data, even when using identical the same network structures. Therefore, they are expected to provide complementary information to achieve label denoising. Consequently, we introduce the concept of the CM-MoE technique for fusing models of different modalities to acquire robust embeddings against label noise. The gate structure in such a system can dynamically adjust the importance of each modality for a specific action, thereby enhancing the robustness of the model in noisy environments.

After obtaining three pre-trained models and a dataset with labels worthy of belief, we construct a spatiotemporal graph-based gate network structure to control the varying weights of the SoftMax layer of the three models. The gate network receives raw data compiled from the three modalities, with the three pre-trained models serving as experts to process their respective modalities.

Let  $S_c^{joint} = \mathbf{N}_{joint}(\mathbf{j}_c)$ ,  $S_c^{bone} = \mathbf{N}_{bone}(\mathbf{b}_c)$  and  $S_c^{motion} = \mathbf{N}_{motion}(\mathbf{m}_c)$  as the SoftMax score of each expert network, where  $\mathbf{j}_c$ ,  $\mathbf{b}_c$ , and  $\mathbf{m}_c$  are derived from one sample in  $D_c$ . The final prediction result is calculated as follows:

$$\mathbf{W} = \mathbf{G}_\theta(\mathbf{j}_c || \mathbf{b}_c || \mathbf{m}_c) \quad (13)$$

$$S = \mathbf{W} \left( S_c^{joint} || S_c^{bone} || S_c^{motion} \right) \quad (14)$$

where  $\mathbf{G}_\theta$ ,  $\mathbf{W}$ , and  $S$  represent the gate network, the output (weights for three expert networks) of the gate network, and the final prediction score.  $||$  represents the concatenation operation. Specifically, The gate network  $\mathbf{G}_\theta$  comprises two basic blocks in CTR-GCN [18], followed by global average pooling and a SoftMax classifier to generate the predicted weights for each expert.

## V. EXPERIMENTS

### A. Experimental Settings

All experiments in this section are implemented on RTX 6000 GPU with the PyTorch deep learning framework. And all data pre-processing methods follow CTR-GCN [50].

**Backbone.** For a fairer performance comparison, we implement the CTR-GCN [18] model as the backbone of the NPC and SOP methods in our testbed. Additionally, we select HD-GCN [26] and ST-GCN [28] for ablation of the backbones.

**Hyperparameters.** In the cross-training part, the hyperparameters setting follows CTR-GCN. each model is trained using SGD with a momentum of 0.9 and weight decay of 0.0004. The basic learning rate is set to 0.1 and the training epoch is set to 65. In the part of fine-tuning with the CM-MoE network, and weight decay is set to 0.0005, the basic learning rate and the training epoch are set to 0.1 and 10.

### B. Results

In this section, we compare the performance of our method with two typical existing methods (NPC [17] and SOP [16]), which focus on the noisy label problem in the field of image classification. Table I showcases our method’s performance, subjected to symmetric noise levels ranging from 20% to 80% on the NTU-60 [49] dataset while using CTR-GCN [18] as the backbone. We randomly choose 4 different noise ratios among 0.0% to 100.0% to set up this benchmark, encompassing easy, middle, and hard levels.

**Comparison with the State-of-the-Art.** SOP and NPC demonstrate varying performance improvements over CTR-GCN without label denoising across different noise levels. In high-noise situations, both SOP and NPC achieve similar accuracies around 70%. At noise ratio from 20% to 50%, NPC shows better performance than SOP, especially at 20% where it can reach 92% accuracy.

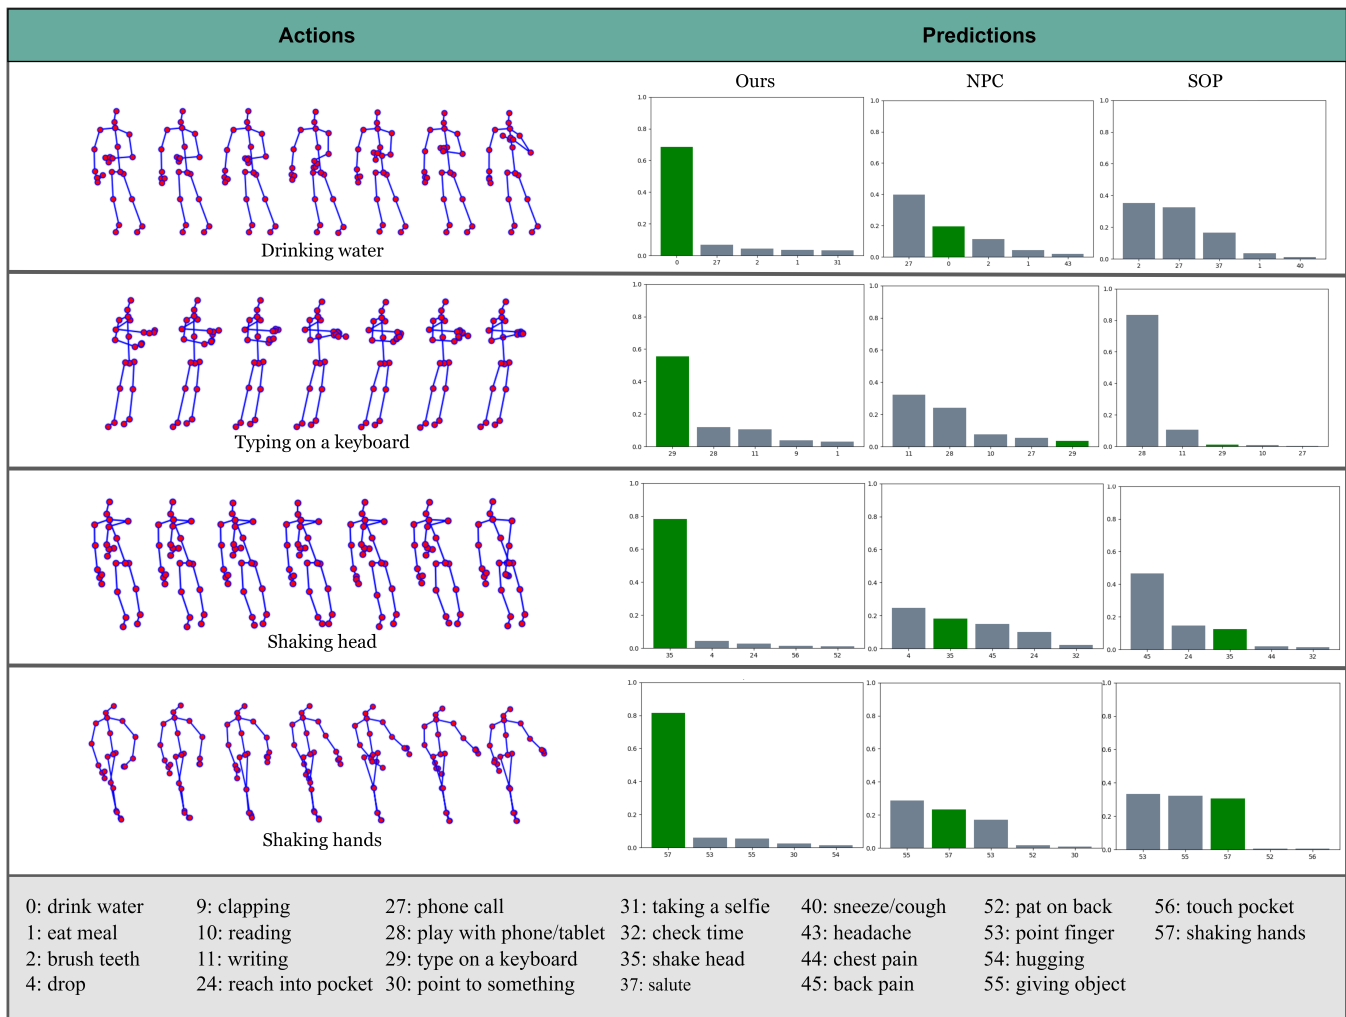


Fig. 3. Prediction results of the three methods. Four samples from different action classes are visualized on the left. The Top-5 SoftMax scores are drawn on the right with the ground truth (green color). All these predictions are generated from the models under the 80% Cross-View setting.

TABLE III  
 THE EXPERIMENT WITH DIFFERENT BACKBONES UNDER 80% SYMMETRIC NOISE, COMPARING WITH SOP [16] METHOD.

Method/Backbone	Symmetric Noise 80%					
	CTR-GCN [50]		HD-GCN [26]		ST-GCN [28]	
	X-Sub	X-View	X-Sub	X-View	X-Sub	X-View
SOP	69.2%	68.2%	71.1%	70.5%	47.5%	46.2%
NoiseEraSAR(ours)	<b>74.9%</b>	<b>79.5%</b>	<b>76.2%</b>	<b>76.5%</b>	<b>54.6%</b>	<b>47.2%</b>

Our method consistently outperforms SOP and NPC across various noise levels. At the noise level of 80%, our model shows an improvement from 5.7% to 10.3% over the baselines. It indicates that, while SOP and NPC play a limited role in enhancing robustness at high noise levels, our method obtains further significant improvement. Even in scenarios with lower noise levels, our method surpasses SOP and NPC. For example, at 20% noise, our method achieves an enhancement of approximately 5%.

The main idea behind SOP and NPC is to model noisy labels, establishing a relationship between noisy labels and true labels. However, this approach is challenging to accurately estimate, particularly when dealing with a large

number of classes [13], and it is highly influenced by the modeling method. Regarding SHAR, datasets often have many categories, with high similarity between many of them, making the applicability of such methods relatively limited.

**Robustness of backbone.** Then, we analyze the CTR-GCN model’s robustness without any additional anti-noise methods. Our experiments indicate that the CTR-GCN model remains sufficient robustness under 20% to 50% noise ratio, with test accuracy dropping from 86.8% to 81.7% (cross-subject) and from 90.9% to 85.7% (cross-view). However, at 80% noise, a significant accuracy decline is evident, with 64.8% (cross-subject) and 64.6% (cross-view). This highlights that the SHAR algorithm, exemplified by CTR-GCN, is sensitive

to noisy labels since the powerful learning ability of these models comes with a great tendency to memorize random noise from the labels.

**Performance under different noise ratios.** Our method demonstrates superior performance in all noise environments. Notably, our model attains over 90% top-1 accuracy in low-noise scenarios. For instance, at 20%, the method achieves 95.3% accuracy at cross-view and 90.6% accuracy at cross-subject. Furthermore, in high-noise environments (at 80% noise ratio), the model performs well, reaching 74.9% accuracy at cross-subject and 79.5% at cross-view. Additionally, at 40% and 50% noise, the model maintains an accuracy of approximately 90%. With the model’s improved expressive capability during training, its denoising ability consistently strengthens through cross-training. Consequently, the model can adaptively obtain suitable denoising capabilities. Furthermore, our multi-modality fusion approach enhances the model’s robustness by leveraging the complementarity among different modalities, thus mitigating excessive memorization of noisy labels within a single modality. These findings underscore our method’s adaptability to diverse noise environments, highlighting its potential for addressing real-world noise challenges.

### C. Qualitative Results

To further assess our method against two other approaches, we select samples from four action classes that are easily confused by other categories when recognized by the human eye: *Drinking water*, *Typing on a keyboard*, *Shaking hands*, and *Shaking head*, as shown in Figure 3. Both SOP and NPC methods display less discriminative power, often ranking the correct class within the top 5 but confusing it with similar classes—for example, NPC mixes up *phone call* with *drinking water*.

In contrast, our model excels in distinguishing different actions, consistently ranking the correct class higher than others, and showcasing improved confidence and accuracy. This enhancement is credited to our model’s ability to effectively reduce noise in training and to utilize the complementary strengths of various modalities, such as joint and bone movements, to better understand and differentiate actions.

### D. Ablation Study

In this section, we first conduct an ablation study to demonstrate the effectiveness of the key components of our NoiseEraSAR method. We ablate each component to quantify its contribution to the overall performance on NTU-60 with 80% noisy labels. The result of the ablation study is shown in Table II. Finally, we conduct the ablation to evaluate backbone generalizability.

**Cross Training.** First, comparing the CTR-GCN with other state-of-the-art methods in Sec. V-B, it is observed that the cross-training method can effectively alleviate the issue of model performance degradation under a high proportion of label noise. With up to 80% symmetric noise, the accuracy of model recognition under X-Sub improves from 64.8% to 70.8% (X-View improves from 64.6% to 74.9%). The

effectiveness of the cross-training approach for the SHAR under the noisy label problem is evident. The models cross-select samples with each other, which weakens the effect of over-memorization of noise.

**Multi-Modality Fusion.** In our ablation experiments, we verify that the multi-modality fusion approach enhances model performance under label noise. We fuse the models after cross-training in three modalities, using a simple ensemble approach similar to CTR-GCN. Specifically, we calculate the weighted sum of the SoftMax layer outputs from the three model outputs. The weights selected for joint, bone, and motion are 0.6, 0.6, and 0.4. The results show that although the recognition accuracy of the bone and motion modalities after cross-training is lower than that of joint modality (70.8% and 74.9%), the model’s performance can be improved by 2% to 4% (72.5% and 78.9%) from the single modality. Therefore, we believe that models with different modalities have complementarity in human action recognition tasks after cross-training, and that modal fusion is significant for performance improvement.

**CM-MoE.** We then conduct experiments to evaluate the effectiveness of the CM-MoE system. If the global sample selector and the CM-MoE system are not beneficial to address the issue of noisy labels, we would expect the utilization of these components can not improve the accuracy, compared with the simple ensemble method. Specifically, the model achieved a performance of 74.9% under X-Sub and 79.5% under X-View, surpassing the performance of the simple ensemble method (summation with fixed weights as presented in the previous section). In cases of a high ratio of noise, the MoE system can optimize the fusion of different modalities by assigning different weights to different features of the sample through the gate network.

**Backbone Generalizability.** We also conduct experiments with various backbones to assess the cross-backbone generalizability of our method. As shown in Table III, our model consistently outperforms the SOP method on CTR-GCN [50], HD-GCN [26], and ST-GCN [28]. For instance, when employing HD-GCN as the backbone, the model achieves a recognition accuracy exceeding 76.0%, surpassing SOP by more than 5%. Even with ST-GCN, where both methods exhibit lower accuracy due to ST-GCN’s weaker robustness in noisy environments, our approach still maintains an advantage over SOP. This result demonstrates the promising adaptability of our method to current state-of-the-art human action recognition algorithms when using them as backbones and validates the effectiveness of our proposed NoiseEraSAR independently of the backbone.

## VI. CONCLUSION

In this paper, we explore skeleton-based human action recognition with noisy labels. We created a benchmark using three label-denoising methods applied to GCN backbones. Our new method, NoiseEraSAR, employs co-teaching for multi-modal streams, global sample selection, and CM-MOE. NoiseEraSAR shows promising results on the NTU-60 dataset across various noisy label ratios and evaluation methods. This

work aims to enhance robust skeleton-based action recognition models for robot-assisted human activities, even with low-quality training data, benefiting the community.

## REFERENCES

- [1] M. Dallel, V. Havard, D. Baudry, and X. Savatier, "Inhard-industrial human action recognition dataset in the context of industrial collaborative robotics," in *Proc. ICHMS*, 2020, pp. 1–6.
- [2] I. Rodomagoulakis *et al.*, "Multimodal human action recognition in assistive human-robot interaction," in *Proc. ICASSP*, 2016, pp. 2702–2706.
- [3] J. Lee and B. Ahn, "Real-time human action recognition with a low-cost RGB camera and mobile robot platform," *Sensors*, vol. 20, no. 10, p. 2886, 2020.
- [4] S. C. Akkaladevi and C. Heindl, "Action recognition for human robot interaction in industrial applications," in *Proc. CGVIS*, 2015, pp. 94–99.
- [5] J. Yu, H. Gao, Y. Chen, D. Zhou, J. Liu, and Z. Ju, "Adaptive spatiotemporal representation learning for skeleton-based human action recognition," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 14, no. 4, pp. 1654–1665, 2022.
- [6] S. Cho, M. Maqbool, F. Liu, and H. Foroosh, "Self-attention network for skeleton-based human action recognition," in *Proc. WACV*, 2020, pp. 624–633.
- [7] D.-T. Pham, T.-N. Nguyen, T.-L. Le, and H. Vu, "Spatio-temporal representation for skeleton-based human action recognition," in *Proc. MAPR*, 2020, pp. 1–6.
- [8] J. Koch, L. Büsch, M. Gomse, and T. Schüppstuhl, "A methods-time-measurement based approach to enable action recognition for multi-variant assembly in human-robot collaboration," *Procedia CIRP*, vol. 106, pp. 233–238, 2022.
- [9] A. Althnian *et al.*, "Impact of dataset size on classification performance: An empirical evaluation in the medical domain," *Applied Sciences*, vol. 11, no. 2, p. 796, 2021.
- [10] H. Song, M. Kim, and J.-G. Lee, "SELFIE: Refurbishing unclean samples for robust deep learning," in *Proc. ICML*, 2019, pp. 5907–5915.
- [11] J. Wei, H. Liu, T. Liu, G. Niu, M. Sugiyama, and Y. Liu, "To smooth or not? When label smoothing meets noisy labels," in *Proc. ICML*, 2022, pp. 23 589–23 614.
- [12] P. Chen, J. Ye, G. Chen, J. Zhao, and P.-A. Heng, "Beyond class-conditional assumption: A primary attempt to combat instance-dependent label noise," in *Proc. AAAI*, 2021, pp. 11 442–11 450.
- [13] B. Han *et al.*, "Co-teaching: Robust training of deep neural networks with extremely noisy labels," in *Proc. NeurIPS*, 2018, pp. 8536–8546.
- [14] P. Chen, B. B. Liao, G. Chen, and S. Zhang, "Understanding and utilizing deep neural networks trained with noisy labels," in *Proc. ICML*, 2019, pp. 1062–1070.
- [15] J. Li, R. Socher, and S. C. H. Hoi, "DivideMix: Learning with noisy labels as semi-supervised learning," in *Proc. ICLR*, 2020.
- [16] S. Liu, Z. Zhu, Q. Qu, and C. You, "Robust training under label noise by over-parameterization," in *Proc. ICML*, 2022, pp. 14 153–14 172.
- [17] H. Bae, S. Shin, B. Na, J. Jang, K. Song, and I.-C. Moon, "From noisy prediction to true label: Noisy prediction calibration via generative model," in *Proc. ICML*, 2022, pp. 1277–1297.
- [18] Y. Chen, Z. Zhang, C. Yuan, B. Li, Y. Deng, and W. Hu, "Channel-wise topology refinement graph convolution for skeleton-based action recognition," in *Proc. ICCV*, 2021, pp. 13 339–13 348.
- [19] P. Wang, Z. Li, Y. Hou, and W. Li, "Action recognition based on joint trajectory maps using convolutional neural networks," in *Proc. MM*, 2016, pp. 102–106.
- [20] T. S. Kim and A. Reiter, "Interpretable 3D human action analysis with temporal convolutional networks," in *Proc. CVPRW*, 2017, pp. 1623–1631.
- [21] J. Tu, M. Liu, and H. Liu, "Skeleton-based human action recognition using spatial temporal 3D convolutional neural networks," in *Proc. ICME*, 2018, pp. 1–6.
- [22] Y. Du, W. Wang, and L. Wang, "Hierarchical recurrent neural network for skeleton based action recognition," in *Proc. CVPR*, 2015, pp. 1110–1118.
- [23] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng, "View adaptive neural networks for high performance skeleton-based human action recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 8, pp. 1963–1978, 2019.
- [24] C. Si, W. Chen, W. Wang, L. Wang, and T. Tan, "An attention enhanced graph convolutional LSTM network for skeleton-based action recognition," in *Proc. CVPR*, 2019, pp. 1227–1236.
- [25] Y. Zhang, B. Wu, W. Li, L. Duan, and C. Gan, "STST: Spatial-temporal specialized transformer for skeleton-based action recognition," in *Proc. MM*, 2021, pp. 3229–3237.
- [26] J. Lee, M. Lee, D. Lee, and S. Lee, "Hierarchically decomposed graph convolutional networks for skeleton-based action recognition," in *Proc. ICCV*, 2023, pp. 10 410–10 419.
- [27] D. Ahn, S. Kim, H. Hong, and B. C. Ko, "STAR-Transformer: A spatio-temporal cross attention transformer for human action recognition," in *Proc. WACV*, 2023, pp. 3330–3339.
- [28] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *Proc. AAAI*, 2018, pp. 7444–7452.
- [29] F. Ye, S. Pu, Q. Zhong, C. Li, D. Xie, and H. Tang, "Dynamic GCN: Context-enriched topology learning for skeleton-based action recognition," in *Proc. MM*, 2020, pp. 55–63.
- [30] K. Peng *et al.*, "Navigating open set scenarios for skeleton-based action recognition," in *Proc. AAAI*, 2024.
- [31] H. Song, M. Kim, D. Park, Y. Shin, and J.-G. Lee, "Learning from noisy labels with deep neural networks: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 11, pp. 8135–8153, 2023.
- [32] X. Xia *et al.*, "Are anchor points really indispensable in label-noise learning?" in *Proc. NeurIPS*, 2019, pp. 6838–6849.
- [33] R. Tanno, A. Saedi, S. Sankaranarayanan, D. C. Alexander, and N. Silberman, "Learning from noisy labels by regularized estimation of annotator confusion," in *Proc. CVPR*, 2019, pp. 11 236–11 245.
- [34] Z. Zhu, Y. Song, and Y. Liu, "Clusterability as an alternative to anchor points when learning with noisy labels," in *Proc. ICML*, 2021, pp. 12 912–12 923.
- [35] Z. Zhu, J. Wang, and Y. Liu, "Beyond images: Label noise transition matrix estimation for tasks with lower-quality features," in *Proc. ICML*, 2022, pp. 27 633–27 653.
- [36] S. Li, X. Xia, H. Zhang, Y. Zhan, S. Ge, and T. Liu, "Estimating noise transition matrix with label correlations for noisy multi-label learning," in *Proc. NeurIPS*, 2022, pp. 24 184–24 198.
- [37] T. Liu and D. Tao, "Classification with noisy labels by importance reweighting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 3, pp. 447–461, 2016.
- [38] J. Shu *et al.*, "Meta-weight-net: Learning an explicit mapping for sample weighting," in *Proc. NeurIPS*, 2019, pp. 1919–1930.
- [39] Z. Wang, G. Hu, and Q. Hu, "Training noise-robust deep neural networks via meta-learning," in *Proc. CVPR*, 2020, pp. 4523–4532.
- [40] G. Zheng, A. H. Awadallah, and S. Dumais, "Meta label correction for noisy label learning," in *Proc. AAAI*, 2021, pp. 11 053–11 061.
- [41] Y. Liu and H. Guo, "Peer loss functions: Learning from noisy labels without knowing noise rates," in *Proc. ICML*, 2020, pp. 6226–6236.
- [42] X. Ma, H. Huang, Y. Wang, S. Romano, S. Erfani, and J. Bailey, "Normalized loss functions for deep learning with noisy labels," in *Proc. ICML*, 2020, pp. 6543–6553.
- [43] Z. Zhu, T. Liu, and Y. Liu, "A second-order approach to learning with instance-dependent label noise," in *Proc. CVPR*, 2021, pp. 10 108–10 118.
- [44] H. Wei, L. Tao, R. Xie, and B. An, "Open-set label noise can improve robustness against inherent label noise," in *Proc. NeurIPS*, 2021, pp. 7978–7992.
- [45] H. Cheng, Z. Zhu, X. Sun, and Y. Liu, "Mitigating memorization of noisy labels via regularization between representations," in *Proc. ICLR*, 2023.
- [46] S. Liu, Z. Zhu, Q. Qu, and C. You, "Robust training under label noise by over-parameterization," in *Proc. ICML*, 2022, pp. 14 153–14 172.
- [47] L. Cheng *et al.*, "Weakly supervised learning with side information for noisy labeled images," in *Proc. ECCV*, 2020, pp. 306–321.
- [48] D. Patel and P. S. Sastry, "Adaptive sample selection for robust learning under label noise," in *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2023, pp. 3921–3931.
- [49] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "NTU RGB+ D: A large scale dataset for 3D human activity analysis," in *Proc. CVPR*, 2016, pp. 1010–1019.
- [50] Y. Chen, Z. Zhang, C. Yuan, B. Li, Y. Deng, and W. Hu, "Channel-wise topology refinement graph convolution for skeleton-based action recognition," in *Proc. ICCV*, 2021, pp. 13 339–13 348.