

Towards Fine-Grained Emotion Understanding via Skeleton-Based Micro-Gesture Recognition

Hao Xu¹, Lechao Cheng^{1,*}, Yaxiong Wang¹, Shengeng Tang¹ and Zhun Zhong¹

¹Hefei University of Technology, No. 485, Danxia Road, Shushan District, Hefei, 230601, China

Abstract

We present our solution to the MiGA Challenge at IJCAI 2025, which aims to recognize micro-gestures (MGs) from skeleton sequences for the purpose of hidden emotion understanding. MGs are characterized by their subtlety, short duration, and low motion amplitude, making them particularly challenging to model and classify. We adopt PoseC3D as the baseline framework and introduce three key enhancements: (1) a topology-aware skeleton representation specifically designed for the iMiGUE dataset to better capture fine-grained motion patterns; (2) an improved temporal processing strategy that facilitates smoother and more temporally consistent motion modeling; and (3) the incorporation of semantic label embeddings as auxiliary supervision to improve the model's generalization ability. Our method achieves a Top-1 accuracy of 67.01% on the iMiGUE test set. As a result of these contributions, our approach ranks third on the official MiGA Challenge leaderboard. The source code is available at https://github.com/EGO-False-Sleep/Miga25_track1.

Keywords

Micro-gesture, action classification, data preprocessing, skeleton-based action recognition

1. Introduction

The Micro-Gesture Analysis for Hidden Emotion Understanding (MiGA) Challenge at IJCAI 2025 focuses on advancing the recognition of micro-gestures (MGs)—brief, involuntary body movements that convey subtle emotional cues. Unlike conventional action [1, 2] or gesture recognition [3, 4], MG recognition presents unique challenges due to the subtlety [5], short duration, and low amplitude of the gestures, as well as the strong contextual dependency and significant class imbalance [6] in the datasets. This challenge builds upon two publicly available datasets, iMiGUE [7] and SMG [8], with iMiGUE featuring spontaneous gestures collected from real-world post-match press conferences. To address these challenges, we formulate the MG classification task as a skeleton-based action recognition problem and adopt PoseC3D [9] as our baseline model. We further introduce three key improvements: (1) a joint skeleton topology with a customized connectivity structure inspired by ST-GCN [10], tailored to better capture localized motion patterns in MGs; (2) a refined temporal processing strategy that ensures smoother and more coherent skeletal motion representation; and (3) the integration of semantic label embeddings [11, 12, 13, 14] to provide auxiliary supervision for improved generalization across imbalanced and semantically overlapping categories. Our final model achieves a Top-1 accuracy of 67.01% on the iMiGUE test set, demonstrating the effectiveness of our topology-aware, semantically guided modeling framework for micro-gesture recognition. The main contributions of our method are summarized as follows.

- We propose a joint graph structure inspired by ST-GCN [10], along with a novel joint connectivity topology that more accurately captures the fine-grained motion patterns specific to the iMiGUE dataset. We also perform a cross-dataset analysis to investigate how different skeletal configurations influence recognition performance across architectures.

MiGA@IJCAI25: International IJCAI Workshop on 3rd Human Behavior Analysis for Emotion Understanding, August 29, 2025, Guangzhou, China.

*Corresponding author.

✉ 2865346358@qq.com (H. Xu); chenglc@hfut.edu.cn (L. Cheng); wangyx15@stu.xjtu.edu.cn (Y. Wang); tangsg@hfut.edu.cn (S. Tang); zhunzhong007@gmail.com (Z. Zhong)

>ID 0009-0003-8156-1765 (H. Xu); 0000-0002-7546-9052 (L. Cheng); 0000-0001-6596-8117 (Y. Wang); 0000-0001-6313-2543 (S. Tang); 0000-0002-8202-0544 (Z. Zhong)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

- We introduce an improved temporal sampling and alignment strategy that departs from the original ST-GCN formulation. This approach enhances motion continuity and enables more coherent representation of raw skeleton sequences.
- Our complete pipeline, including the proposed topological and temporal enhancements, achieves a Top-1 accuracy of 67.01% on the iMiGUE test set.

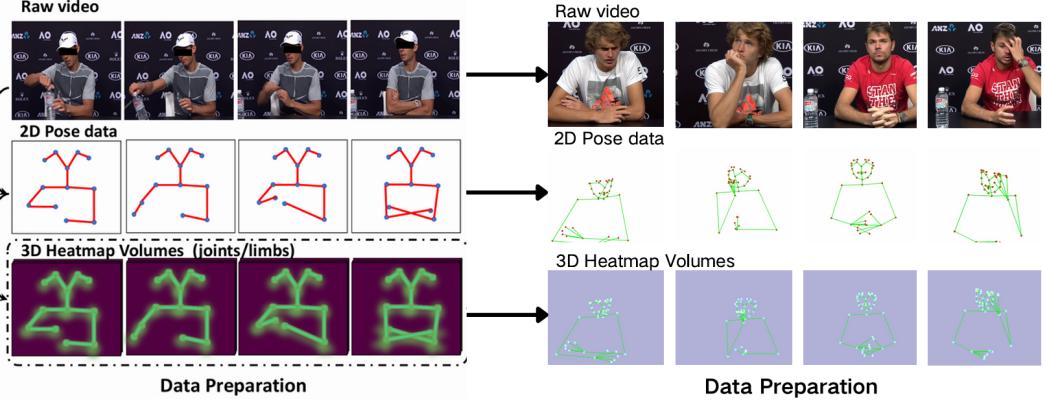


Figure 1: Visualization of the proposed new skeletal joint connection diagram for micro-gesture classification.

2. Methodology

In this challenge, we experiment with two representative skeleton-based action recognition frameworks: **ST-GCN** and **PoseC3D**. This section details our design choices, empirical findings, and analysis based on both architectures.

2.1. Skeleton Augmentation with Facial Keypoints

Human skeletal connectivity is spatially consistent and relatively easy for graph-based models to learn. However, in the iMiGUE micro-gesture recognition task, many action categories—such as touching the face, adjusting a hat, or biting lips—are localized in the facial region. To enhance facial motion perception, we extend the standard 22-joint OpenPose skeleton to a 41-joint structure by incorporating additional facial landmarks (e.g., cheeks, eyebrows, and lips). This augmentation provides finer spatial resolution in regions critical to emotion-related micro-gestures.

Figure 2 compares the skeletal connectivity diagrams under different keypoint configurations. While this modification benefits representation, it also diverges from the original graph topology assumptions of ST-GCN, as discussed in the next subsection.

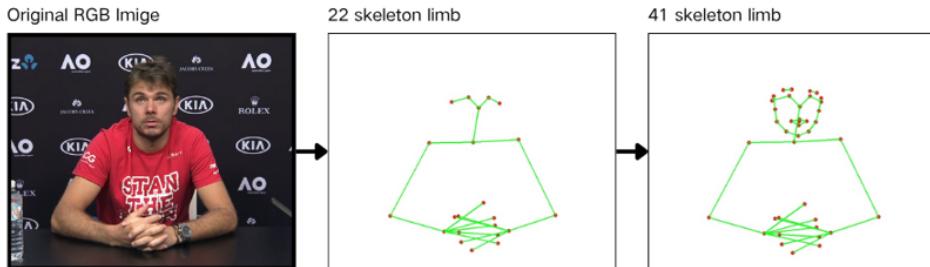


Figure 2: Comparison of skeletal connectivity with different numbers of keypoints.

2.2. Partitioning and Graph Reasoning in ST-GCN

Effective action recognition requires modeling meaningful spatiotemporal motion patterns. ST-GCN achieves this by partitioning the skeleton graph into sub-regions, such as centripetal, centrifugal, and stationary limbs [10]. While effective for coarse-scale action categories, we find this partitioning suboptimal for micro-motion classification, likely due to the lack of distinguishable limb dynamics and limited data scale.

To compensate, we first enrich the model input with additional facial keypoints. These regions, while less dynamic in raw motion, are semantically aligned with many micro-gesture classes. However, empirical results reveal that these added keypoints degrade performance in ST-GCN. We hypothesize this is due to:

- Overfitting from increased graph complexity and limited training data.
- The added nodes not contributing salient temporal or relational motion patterns that ST-GCN is designed to exploit.
- The peripheral nature and low motion magnitude of facial keypoints reducing their relative attention in the learned graph features.

Thus, despite their semantic relevance, these keypoints are possibly treated as noise within the ST-GCN’s fixed topology and partitioning scheme¹.

2.3. PoseC3D with Extended Keypoints

Given the limitations of ST-GCN, we shift our focus to PoseC3D, a 3D-CNN based method that processes skeletons as spatiotemporal heatmaps. Even in its baseline form, PoseC3D significantly outperforms ST-GCN. More importantly, when extended with facial keypoints, PoseC3D benefits substantially in performance. We attribute this to two key properties:

- **Heatmap-based Representation:** PoseC3D encodes keypoints as dense heatmaps, which preserve richer spatial information and allow the network to infer latent movement patterns—even in low-motion regions. This representation has higher information entropy than raw joint coordinates, enabling stronger generalization.
- **Flexible 3D Convolutions:** The spatiotemporal convolutions in PoseC3D operate over the entire motion volume with uniform treatment of all locations. Unlike GCN’s, the receptive field and feature propagation are not constrained by predefined skeletal graphs, granting PoseC3D greater expressivity and robustness to irrelevant noise.

2.4. Temporal Frame Stream Processing

Temporal modeling is a critical component in micro-gesture recognition, given the subtlety and brevity of such actions. The default temporal sampling strategy in ST-GCN adopts simple rule-based heuristics:

- When the number of frames exceeds the target length, a continuous subsequence is randomly cropped from the original sequence;
- When the number of frames is insufficient, zero-padding is applied to extend the sequence to the required length.

However, these strategies often fail to preserve the complete temporal structure of micro-gestures. Random cropping may exclude key motion cues, while zero-padding introduces artificial discontinuities that disrupt temporal coherence. These limitations are particularly detrimental in micro-motion scenarios, where discriminative features are both sparse and temporally localized. To address these issues, we propose a structure-preserving temporal alignment strategy as follows:

¹We acknowledge that this interpretation may be influenced by our limited experience with ST-GCN and time constraints during the challenge. We welcome future improvements from the community in this direction.

- For over-length sequences, we perform uniform interval sampling, ensuring that both the first and last frames are retained. This guarantees that the sampled sequence spans the full temporal range of the original gesture;
- For under-length sequences, we apply linear interpolation to generate intermediate frames, thereby expanding the sequence to the target length while maintaining temporal smoothness and continuity.

Compared to conventional approaches, the proposed strategy offers better coverage of the gesture trajectory and preserves fine-grained motion dynamics. Empirical results further confirm that this refinement contributes to improved model stability and recognition accuracy in the micro-gesture classification task.

3. Experiments

3.1. Dataset: iMiGUE [7]

The Micro-Gesture Understanding and Emotion Analysis (iMiGUE) dataset consists of 32 micro-gesture (MG) categories and one additional non-MG class. All data are collected from post-match press conference videos of professional tennis players. The dataset comprises a total of 18,499 annotated MG samples, which are labeled from 359 long video sequences ranging in duration from 0.5 to 26 minutes, totaling approximately 3,765,600 frames. iMiGUE provides two modalities for each sample: (1) RGB videos, and (2) 2D skeletal joint coordinates extracted using the OpenPose pose estimation framework. This multi-modal design enables both appearance-based and skeleton-based gesture analysis, supporting the development of robust models for emotion understanding based on subtle behavioral cues.

3.2. Evaluation Metrics and Implementation Details

We evaluate micro-gesture classification performance using the Top-1 Accuracy, which measures the percentage of samples for which the predicted label exactly matches the ground truth. Our method is implemented based on the open-source PySk1 toolbox [15], and the training pipeline incorporates a loss function inspired by the winning solution of the MiGA 2023 challenge. The model is trained using Stochastic Gradient Descent (SGD) with a momentum of 0.9, a weight decay of 3×10^{-4} , and a batch size of 24. The initial learning rate is set to 0.1/3, and we adopt a Cosine Annealing learning rate schedule. We use ResNet3D-SlowOnly as the feature extraction backbone and I3D as the classification head. For multi-stream ensemble modeling, which integrates joint and limb modalities, we apply a weighted fusion scheme with a ratio of 1:1, ensuring equal contribution from both sources of motion information.

3.3. Experiments

3.4. Quantitative Results

Table 1 presents the top-3 results on the iMiGUE test set as reported on the official CodaLab competition leaderboard². Our method achieves a Top-1 Accuracy of 67.01%, ranking third overall among all participants.

Table 2 compares our method against standard PoseC3D and the winning solution of MiGA2023 [16] across different input modalities on the iMiGUE dataset. All models use the same backbone (ResNet3D-SlowOnly) and classification head (I3D). Our method consistently outperforms existing approaches, achieving improvements across all modality models compared to the 2023 champion solution that also uses semantic embeddings. 4

These results demonstrate the effectiveness of our approach in modeling subtle motion cues through enhanced skeleton topology, refined temporal processing, and semantic label supervision.

²Available at: <https://www.kaggle.com/competitions/the-3rd-mi-ga-ijcai-challenge-track-1/leaderboard>

Table 1

Top-3 entries on the iMiGUE test set (Top-1 Accuracy).

Rank	Team	Top-1 Accuracy (%)
1	gkdx2	73.213
2	awuniverse	68.697
3	Ours	67.010

Table 2

Performance comparison under different modality settings on iMiGUE.

Method	Modality	Top-1 Accuracy (%)
ST-GCN [10]	Joint	46.38
PoseC3D [9]	Joint	59.54
PoseC3D [9]	Limb	60.74
MiGA2023 1st [16]	Joint	62.28
MiGA2023 1st [16]	Limb	63.48
MiGA2023 1st [16]	Joint & Limb	64.12
Ours	Joint	65.43 (+3.15)
Ours	Limb	64.40 (+0.92)
Ours	Joint & Limb	67.01 (+2.89)

4. Conclusions

In this paper, we present the solution developed for the MiGA Challenge held at IJCAI 2025. Throughout the process, we employed both the ST-GCN and PoseC3D models, comparing their similarities and differences to explore the relationship between convolutional approaches and sequential data. Ultimately, by leveraging joint and limb modality data and adopting PoseC3D as the backbone—combined with the semantic embedding loss[16] proposed in 2023—our method achieved third place with a Top-1 accuracy of 67.01%. For this task, we recognize that there remains ample room for further research. Moving forward, we plan to address the challenges from additional perspectives, such as improved denoising techniques, strategies for handling imbalanced data, and the integration of RGB video streams, among others.

5. Acknowledgements

This work has been supported by the National Natural Science Foundation of China (Grant No. 62472139), by the Anhui Provincial Natural Science Foundation, China (Grant No. 2408085QF191).

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: paraphrase and reword sentences, and check for grammar and spelling errors. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

- [1] X. Lu, S. Zhao, L. Cheng, Y. Zheng, X. Fan, M. Song, Mixed resolution network with hierarchical motion modeling for efficient action recognition, *Knowledge-Based Systems* 294 (2024) 111686.
- [2] X. Lu, Y. Hao, L. Cheng, S. Zhao, Y. Liu, M. Song, Mixed attention and channel shift transformer

for efficient action recognition, *ACM Transactions on Multimedia Computing, Communications and Applications* 21 (2025) 1–20.

- [3] S. Tang, J. He, D. Guo, Y. Wei, F. Li, R. Hong, Sign-idd: Iconicity disentangled diffusion for sign language production, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 2025, pp. 7266–7274.
- [4] S. Tang, J. He, L. Cheng, J. Wu, D. Guo, R. Hong, Discrete to continuous: Generating smooth transition poses from sign language observations, in: *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 3481–3491.
- [5] Y. Zhang, L. Cheng, Y. Wang, Z. Zhong, M. Wang, Towards micro-action recognition with limited annotations: An asynchronous pseudo labeling and training approach, *arXiv preprint arXiv:2504.07785* (2025).
- [6] C. Fang, L. Cheng, Y. Mao, D. Zhang, Y. Fang, G. Li, H. Qi, L. Jiao, Separating noisy samples from tail classes for long-tailed image classification with label noise, *IEEE Transactions on Neural Networks and Learning Systems* (2023).
- [7] X. Liu, H. Shi, H. Chen, Z. Yu, X. Li, G. Zhao, 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*, 2021, pp. 10631–10642.
- [8] H. Chen, H. Shi, X. Liu, X. Li, G. Zhao, Smg: A micro-gesture dataset towards spontaneous body gestures for emotional stress state analysis, *International Journal of Computer Vision* 131 (2023) 1346–1366.
- [9] H. Duan, Y. Zhao, K. Chen, D. Lin, B. Dai, Revisiting skeleton-based action recognition, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 2969–2978.
- [10] S. Yan, Y. Xiong, D. Lin, Spatial temporal graph convolutional networks for skeleton-based action recognition, in: *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [11] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, M. Ranzato, T. Mikolov, Devise: A deep visual-semantic embedding model, *Advances in neural information processing systems* 26 (2013).
- [12] M.-C. Yeh, Y.-N. Li, Multilabel deep visual-semantic embedding, *IEEE transactions on pattern analysis and machine intelligence* 42 (2019) 1530–1536.
- [13] Z. Wei, J. Zhang, Z. Lin, J.-Y. Lee, N. Balasubramanian, M. Hoai, D. Samaras, Learning visual emotion representations from web data, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 13106–13115.
- [14] P. P. Filntisis, N. Efthymiou, G. Potamianos, P. Maragos, Emotion understanding in videos through body, context, and visual-semantic embedding loss, in: *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part I* 16, Springer, 2020, pp. 747–755.
- [15] H. Duan, J. Wang, K. Chen, D. Lin, Pyskl: Towards good practices for skeleton action recognition, in: *Proceedings of the 30th ACM International Conference on Multimedia*, 2022, pp. 7351–7354.
- [16] K. Li, D. Guo, G. Chen, X. Peng, M. Wang, Joint skeletal and semantic embedding loss for micro-gesture classification, *arXiv preprint arXiv:2307.10624* (2023).