

Masked Video and Body-worn IMU Autoencoder for Egocentric Action Recognition

Mingfang Zhang[✉], Yifei Huang^{*✉}, Ruicong Liu[✉], and Yoichi Sato[✉]

Institute of Industrial Science, the University of Tokyo, Tokyo, Japan
{mfzhang,hyf,lruicong,ysato}@iis.u-tokyo.ac.jp

Abstract. Compared with visual signals, Inertial Measurement Units (IMUs) placed on human limbs can capture accurate motion signals while being robust to lighting variation and occlusion. While these characteristics are intuitively valuable to help egocentric action recognition, the potential of IMUs remains under-explored. In this work, we present a novel method for action recognition that integrates motion data from body-worn IMUs with egocentric video. Due to the scarcity of labeled multimodal data, we design an MAE-based self-supervised pretraining method, obtaining strong multi-modal representations via modeling the natural correlation between visual and motion signals. To model the complex relation of multiple IMU devices placed across the body, we exploit the collaborative dynamics in multiple IMU devices and propose to embed the relative motion features of human joints into a graph structure. Experiments show our method can achieve state-of-the-art performance on multiple public datasets. The effectiveness of our MAE-based pretraining and graph-based IMU modeling are further validated by experiments in more challenging scenarios, including partially missing IMU devices and video quality corruption, promoting more flexible usages in the real world.

Keywords: Egocentric action recognition · Inertial Measurement Units · Multimodal Masked Autoencoder

1 Introduction

Advances in wearable cameras and egocentric video datasets bring egocentric action recognition as one pivotal aspect of understanding human behavior [26, 43, 58]. Because of the inherent difficulty in recognizing actions solely from egocentric videos, recent successes embrace multimodal signals, integrating vision, language, and audio to enrich the contextual understanding of actions [18, 29]. However, one affordable alternative, an Inertial Measurement Unit (IMU) that captures acceleration, angular velocity, and orientation of human limbs, remains under-explored. In addition to being cheap and energy-efficient, when affixed to human limbs, IMUs can accurately record 3D human movements which are often hard to observe because of the viewpoint limitation of egocentric cameras [28].

*Corresponding author.

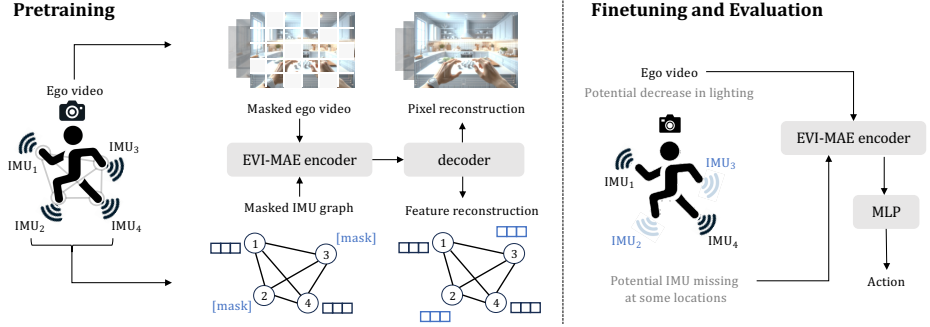


Fig. 1: Overview of our EgoVideoIMU-MAE (EVI-MAE). Because of the scarcity of labeled multimodal data, we propose an MAE-based pretraining approach with unlabeled egocentric video and IMU signals. To exploit the collaborative dynamics in multiple IMU devices, we propose to embed the relative motion features of human joints into a graph. In the finetuning and evaluation phase, we consider potential video and IMU corruption for more flexible usages.

IMUs excel at motion sensing but do not contain visual information [38]. On the other hand, videos provide visual information but are vulnerable to harsh imaging conditions and not optimal for sensing motion [8, 65]. Therefore, joint modeling of video and IMU is a promising direction to leverage their complementarity.

However, two significant challenges prevent us from directly applying existing multimodal action recognition frameworks [17, 19] to this task. The first one is the scarcity of data. While several recent egocentric datasets [18, 20] offer large-scale multimodal egocentric data including IMU, they are measured by camera IMUs and thus can only provide motion clues that can also be inferred from the video. Very few datasets [6, 48] contain synchronized video and separately attached IMU sensor data, but the largest dataset [6] offers only 9 hours of footage with action annotations. The other challenge lies in the complexity of multiple IMU devices across the body. Since each IMU device can provide motion signals of only a single body part. However, a holistic understanding of motions of different body parts is necessary for recognizing various actions. Consequently, modeling these intricate multi-device relations becomes challenging.

In this work, we propose a novel multi-modal egocentric action recognition method using both egocentric video and body-worn IMU signals. To address the challenge of data scarcity, we observe that **egocentric video and IMU signals are correlated**: 1), in egocentric videos, the movements of visible upper limbs show a strong correlation with the IMU signals attached to these limbs. 2), the global motion in the camera view is closely related to the IMU signals from the subject’s lower limbs. To exploit this correlation, we develop a multi-modal Masked Autoencoder (MAE) [22] based pretraining method. Specifically, we use a joint model to take masked video and masked IMU signals as inputs and reconstruct the two modalities simultaneously. With this self-supervised

learning method, we can leverage the large volume of unlabeled video-IMU data to learn the natural alignment and the complementary aspect between the two modalities. This allows us to leverage larger-scale data without manual action annotation, which is a major advantage of our method compared to existing fully supervised methods.

Regarding the second challenge, we observe that **multiple IMU devices are collaborative**. IMU devices positioned at various locations on the human body collect complementary data, providing a holistic picture of body movements. These IMU sensors on human limbs reflect the *relative motion* of human joints. Drawing inspiration from skeleton-based action recognition models [13, 60], we propose a graph-based IMU modeling technique, where we construct a graph with nodes as IMU features. This graph structure can effectively capture motion relationships among different IMU sensors, and the graph neural network can therefore optimize the representation learning.

An overview of our proposed method can be found in Figure 1. During the MAE pretraining, some of the IMU feature nodes are masked together with ego-centric video patches. Our EgoVideoIMU-MAE (EVI-MAE) encoder operates on visible nodes and patches, and the decoder tries to reconstruct the original input. In this way, we obtain an encoder with knowledge of intra-IMU relationships and the video-IMU correlations. The representations learned by this multimodal pre-training technique contain the complementarity of the two modalities and thus can be robust to less restricted scenarios where one modality becomes unreliable, *e.g.*, video quality decline, or partial IMU devices missing.

Experiments on two public datasets [6, 48] show that the representation learned by our EVI-MAE significantly outperforms the performance of baseline models without pretraining or with single-modal pretraining. Our ablation studies confirm the efficacy of constructing the IMU feature graph, regardless of whether pretraining is involved. Considering constraints in real-world applications, we further test our model by designing experiments in various challenging scenarios: 1) partially missing IMU devices, 2) video quality variation, and 3) cross-dataset generalization. Extensive results in these challenging scenarios demonstrate the robustness and adaptability of our model.

In summary, our main contributions are as follows:

- We design the first method for joint egocentric video and body-worn IMU representation learning by multimodal masked autoencoding.
- We propose a graph-based IMU modeling technique to better leverage the collaborative nature of IMU devices on different body locations.
- Experiments show that our method outperforms various baseline methods and continues to exhibit robust performance and adaptability in challenging scenarios, promoting more flexible usages in the real world.

2 Related Works

IMU-based action recognition IMU-based action recognition has seen significant progress. The DeepConvLSTM [39] model merges convolutional and recur-

rent layers to extract features while modeling temporal dependencies. Subsequent efforts [37, 62] have expanded upon this model, introducing novel architectures [1, 56] that further advance the field. Nevertheless, these advancements tend to overlook the intricacies of inter-device relationships among multiple IMUs. SADeepSense [63] and subsequent studies [32, 36] incorporate the attention mechanism [50] to facilitate the fusion of data from heterogeneous sensors, such as accelerometers and gyroscopes, positioned across the body. However, the limited size of datasets with labeled actions constrains their performance. Recent researches [45, 64] introduce the concept of contrastive learning to IMU representation learning. LIMU-BERT [57] adapts BERT [11] for IMU signals, which involves randomly masking a small portion of data and reconstructing them. Compared to previous IMU-based action recognition methods, we propose a representation learning method that utilizes unlabeled IMU signals and also learns multiple IMU device relationships by proposing to construct a graph with IMU features. Moreover, our method incorporates egocentric videos to exploit multimodal relations.

Egocentric multimodal action recognition Most previous action recognition methods are conducted on videos [4, 14, 46, 53]. However, for egocentric action recognition, where the videos often have limited field-of-view and large motion blur [31], incorporating complementary multimodal information becomes crucial [25]. Recent advancements have seen the emergence of many egocentric video datasets [10, 20, 21, 27, 33, 42, 43, 66] that encompass diverse modalities, including video, depth, audio, language, and motion sensors. Prior approaches to multimodal fusion in egocentric activity recognition range from straightforward concatenation [41, 55] to more sophisticated tensor decomposition techniques [35]. MMG [18] aims at addressing the challenge of zero-shot modality mismatch challenge. However, the majority of existing datasets and methods focus on sensing devices attached to the subject’s head, primarily through head-mounted cameras [44, 52], microphones [29, 61], and IMUs [18, 49], while neglecting sensors placed on the human limbs. In contrast to existing works, we propose an egocentric multimodal representation learning approach that integrates visual signals from head-mounted devices with motion signals from limb-mounted sensors.

Multimodal masked autoencoders The concept of data masking has been extensively studied since [51]. Masked Autoencoder (MAE) [22] distinguishes itself by its proficiency in learning meaningful representations. Its applicability spans across various domains, including image processing [22], video analysis [15, 47], audio processing [24], point cloud data interpretation [40], and graph analysis [23]. Additionally, MAE demonstrates robust capabilities in integrating and learning from multiple modalities. This includes image and language data [16, 30], as well as synchronized audio, images [19] and videos [17], showcasing its effectiveness in handling different data representations together. This paper introduces an innovative approach utilizing an MAE-based method to leverage the synergistic and complementary interactions between egocentric video and body-worn IMU motion data for effective action recognition.




Fig. 2: IMU and video data preprocessing and masking.

3 Method

3.1 Problem definition

We use synchronized video and IMU signals of temporal length T as inputs for pretraining and finetuning the model. The video can be represented as $\mathbf{D}_v \in \mathbb{R}^{T \times \mathcal{S}_v \times H \times W \times 3}$, where \mathcal{S}_v is the video sampling rate, and H, W represent the height and width of the video, respectively. For the IMU signal, we consider N_{imu} IMU sensors installed on the person’s limbs. For each IMU, we read the acceleration in the x, y , and z axes while discarding other readings like angular velocity and orientation because acceleration is universally presented across most types of IMUs. Even with the need for including the discarded readings, the proposed method can be easily extended. Overall, the raw IMU signals can be represented as $\mathbf{D}_{raw} \in \mathbb{R}^{N_{imu} \times T \times \mathcal{S}_{imu} \times 3}$, where \mathcal{S}_{imu} is the sampling rate of IMU. Our goal for pretraining is to learn a multimodal representation effective for the downstream action recognition task, which is a classification problem given predefined C classes.

3.2 Data preprocessing and masking

As shown in Fig. 2, the preprocessing of the T -second raw IMU signal involves several steps. Due to varying sample rates across IMU sensors, signals are re-sampled to a consistent rate and normalized to a uniform value range. The STFT (Short Time Fourier Transform) step then processes each 1-dimensional acceleration signal, consisting of T_{imu} resampled temporal points, converting it into a sequence of features of M_{imu} dimensions as a spectrogram $\mathbf{D}_{spec} \in \mathbb{R}^{T_{imu} \times M_{imu}}$. This spectrogram is subsequently divided into P_{imu} square patches, $\mathbf{i} = [i_1, i_2, \dots, i_{P_{imu}}]$, which serve as the input for the model. For the egocentric video, each T -second video segment is temporally down-sampled to T_v frames and then divided into P_v square patches $\mathbf{v} = [v_1, v_2, \dots, v_{P_v}]$.

For the IMU spectrograms, we employ unstructured random masking with a masking ratio of R_{imu} that transforms IMU spectrogram patches \mathbf{i} into \mathbf{i}_m . We also report the performance of structured masking in Tab. 2, including randomly masking a portion of time and frequency of \mathbf{i} . For videos, we employ the tube masking strategy [47] with a high masking ratio of R_v , where the masking locations are the same for all frames, which transforms video patches \mathbf{v} into \mathbf{v}_m .

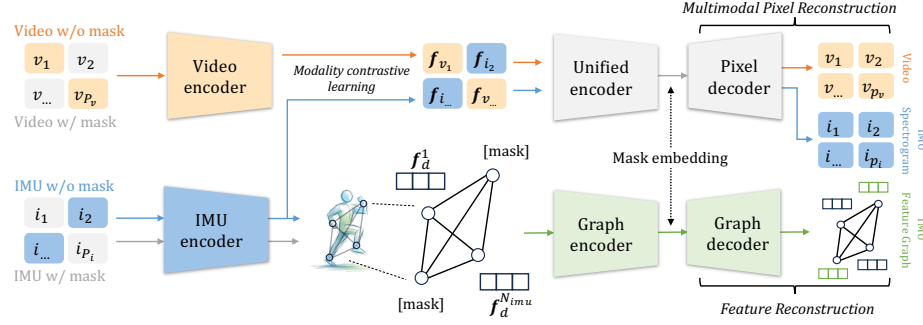


Fig. 3: Our EVI-MAE pretraining network processes video patches \mathbf{v} and IMU spectrogram patches \mathbf{i} and incorporates them into two branches, a multimodal pixel reconstruction branch, and an IMU feature reconstruction branch.

3.3 IMU feature graph

To model the complex relationships among multiple IMU devices positioned across the body, we exploit the collaborative dynamics of these devices and propose to embed the relative motion features of human joints into a graph structure. Specifically, we initiate the process by encoding all the IMU spectrogram patches \mathbf{i} via an IMU encoder. In particular, \mathbf{i} encompasses data from all the N_{imu} IMU devices, $d_1, d_2, \dots, d_{N_{imu}}$, prior to the application of any masking procedures. The encoded features from each IMU device are denoted as $\mathbf{f}_d = [\mathbf{f}_d^1, \mathbf{f}_d^2, \dots, \mathbf{f}_d^{N_{imu}}]$.

Subsequently, we construct an IMU feature graph $\mathcal{G} = (\mathcal{V}, \mathbf{A}, \mathbf{f}_d)$, where $\mathcal{V} = \{\nu_n\}_{n=1}^{N_{imu}}$ denotes the set of nodes corresponding to the IMUs distributed across the body, $\mathbf{A} \in \{0, 1\}^{N_{imu} \times N_{imu}}$ represents the adjacency matrix, and $\mathbf{f}_d \in \mathbb{R}^{N_{imu} \times M_{imu}}$ is IMU node features. The adjacency matrix \mathbf{A} is designed to be fully connected, highlighting that actions usually require coordinated movements across different limbs. The efficiency of this graph-based modeling is guaranteed since in practical applications, individuals typically do not wear a large number of IMU devices.

3.4 Pretraining pipeline

Because of the scarcity of labeled multimodal data, we design an MAE-based self-supervised pretraining method. We aim to effectively capture strong multimodal representations by leveraging the inherent correlation between visual and motion signals. As shown in Fig. 3, the pretraining network structure incorporates two primary branches, the multimodal pixel reconstruction branch and the IMU feature reconstruction branch.

Multimodal pixel reconstruction branch As shown in Fig. 3, this branch involves a video encoder, an IMU encoder, a unified encoder, and a pixel decoder. In this branch, we use the visible patches from a masked video, \mathbf{v}_m , and

the visible patches from masked IMU spectrograms, \mathbf{i}_m , as inputs. Our objective is to reconstruct the masked video and the masked IMU spectrograms. By doing so, we aim to exploit the abundant supply of unlabeled yet synchronized video-IMU data to understand the natural alignment and complementary interaction between these two modalities, thereby achieving a strong multimodal representation.

Specifically, \mathbf{v}_m and \mathbf{i}_m are first incorporated with positional embeddings $\mathbf{p}_v, \mathbf{p}_{imu}$ and modality type embeddings $\mathbf{m}_v, \mathbf{m}_{imu}$. The positional embedding encodes the spatial positions of patches to maintain their spatial relationships and the modality type embedding distinguishes between video and IMU modality. Then, they are inputted into the video encoder and the IMU encoder, respectively. The video encoder utilizes a standard ViT backbone [12] and incorporates joint space-time attention [2, 34], allowing tokens to interact within the multi-head self-attention layer. Meanwhile, the IMU encoder is designed following ImageMAE [22], since we convert the 1D motion signals to 2D spectrogram images. As shown in Fig. 3, the visible video and IMU encoded features, denoted as \mathbf{f}_v and \mathbf{f}_i , are then merged into a unified embedding and processed by a modality-unified encoder, concluding the encoding phase.

For decoding, this multimodal embedding is first padded with trainable masked tokens at the masked positions as $\mathbf{e}_x = [\mathbf{e}_{imu}, \mathbf{e}_v]$. Along with the addition of both modality type embeddings $\mathbf{m}'_{imu}, \mathbf{m}'_v$, and positional embeddings, $\mathbf{p}'_{imu}, \mathbf{p}'_v$, they are passed through a transformer-based pixel decoder, \mathcal{D}_{uni} , i.e.,

$$\mathbf{O}_p = \mathcal{D}_{uni}([\mathbf{e}_{imu}, \mathbf{e}_v] + [\mathbf{m}'_{imu}, \mathbf{m}'_v] + [\mathbf{p}'_{imu}, \mathbf{p}'_v]). \quad (1)$$

Finally, we partition the output pixels \mathbf{O}_p as video and IMU spectrograms and rearrange the pixel patches to align with the original \mathbf{v} and \mathbf{i} .

IMU feature reconstruction branch As shown in Fig. 3, this branch involves the same IMU encoder, a graph encoder, and a graph decoder. In this branch, we shift our focus from patch level to IMU device level. Based on the structure of IMU feature graph, we design an MAE-based pretraining method that uses a corrupted graph with masked nodes to reconstruct the original graph. In this way, we aim to capture the intricate relationships among different IMU sensor features, thereby enhancing the model’s comprehension of motion dynamics and sensor interactions from a holistic perspective.

We first construct an IMU feature graph $\mathcal{G} = (\mathcal{V}, \mathbf{A}, \mathbf{f}_d)$ as described in Sec. 3.3. Inspired by GraphMAE [23], our pretraining output is $\hat{\mathcal{G}}$, whose goal is to reconstruct \mathcal{G} from a corrupted input feature \mathbf{f}_{dc} as:

$$\mathbf{f}_g = \text{GraphEnc}(\mathbf{A}, \mathbf{f}_{dc}), \hat{\mathcal{G}} = \text{GraphDec}(\mathbf{A}, \mathbf{f}_g). \quad (2)$$

Specifically, we employ a masking procedure where a ratio, R_g , of the N_{imu} node features in \mathbf{f}_d are selectively masked by replacing them with mask tokens, resulting in a corrupted graph with features, \mathbf{f}_{dc} . This corrupted graph is then processed using a Graph Isomorphism Network (GIN) [59] encoder, chosen for its

outstanding inductive bias suited for graph-level classification tasks. During the decoding phase, the encoded graph features \mathbf{f}_g , which shares the same shape with \mathbf{f}_d , are once again substituted with learnable mask embeddings at the previously masked positions. At this stage, we reintroduce the GIN network to serve as the decoder to reconstruct the entire graph.

Pretraining criterion For pixel reconstruction, we calculate the mean squared error (MSE) between the pixels of the reconstructed video and those of the original video, as well as the pixels of the IMU spectrogram. In the case of IMU feature graph reconstruction, we use the cosine similarity error to compare the reconstructed IMU device-level features $\hat{\mathbf{f}}_d$ and the original \mathbf{f}_d , within the set of masked node features \mathcal{V}_c . The cosine similarity error is defined as follows:

$$L_{cos} = \frac{1}{|\mathcal{V}_c|} \sum_{\nu_n \in \mathcal{V}_c} \left(1 - \frac{(\hat{\mathbf{f}}_d^{\nu_n})^T \mathbf{f}_d^{\nu_n}}{\|\hat{\mathbf{f}}_d^{\nu_n}\| \cdot \|\mathbf{f}_d^{\nu_n}\|} \right), \quad (3)$$

where $\mathbf{f}_d^{\nu_n}$ denotes the IMU feature of device on node ν_n . Inspired by a multimodal MAE method [19], we further enrich our approach by integrating a contrastive objective, leveraging the output features $\mathbf{f}_v, \mathbf{f}_i$ from both video and IMU encoders:

$$L_{con} = -\frac{1}{2N_b} \left(\sum_{k=1}^{N_b} \log \frac{\exp(s(\mathbf{f}_v^k, \mathbf{f}_i^k)/\tau)}{\sum_{j=1}^{N_b} \exp(s(\mathbf{f}_v^k, \mathbf{f}_i^j)/\tau)} + \sum_{k=1}^{N_b} \log \frac{\exp(s(\mathbf{f}_v^k, \mathbf{f}_i^k)/\tau)}{\sum_{j=1}^{N_b} \exp(s(\mathbf{f}_v^j, \mathbf{f}_i^k)/\tau)} \right), \quad (4)$$

where N_b is the batch size, $s(\mathbf{f}_v^k, \mathbf{f}_i^k) = (\mathbf{f}_v^k)^T (\mathbf{f}_i^k)$, and τ is the temperature. By utilizing supervision formed through non-corresponding video and IMU pairs within the same batch, our model is adept at learning features that are more aligned across modalities. In summary, our total loss is $L = \alpha L_{mse} + \beta L_{cos} + \gamma L_{con}$.

3.5 Finetuning for action recognition

In the fine-tuning stage, we discard the pixel and graph decoders and fine-tune only the encoders. This process involves using all video and IMU patches, denoted as \mathbf{v} and \mathbf{i} . We then concatenate the output features of the unified encoder and the graph encoder, apply global average pooling, and then use a linear classifier on top for the action classification tasks.

4 Experiments

4.1 Datasets

CMU-MMAC The CMU Multi-Modal Activity Database [48] contains *indoor* multimodal data of the human activity involving cooking and food preparation. Our research utilizes egocentric video and four of the 3DM-GX1 IMU sensors

on left and right arms, and left and right legs. The IMU sample rate is roughly 60 Hz. We follow [38] for IMU data preprocessing and get a total of 13 hours of synchronized video and IMUs data, in which 5 hours are annotated with action labels [3] of 32 categories.

WEAR The WEAR [6] dataset is an *outdoor* sports dataset containing visual and IMU data of different workout activities. Our research utilizes egocentric video and all four IMU sensors on left and right wrists, and left and right ankles. The IMU sample rate is 50 Hz. The total duration of the dataset is 15 hours, in which 9 hours are annotated with action labels of 18 categories.

4.2 Implementation details

In this work, we utilize synchronized egocentric video and $N_{imu} = 4$ IMU devices attached to the four limbs with a duration of $T = 2$ seconds. For video processing, we down-sample and crop the clips to $T_v = 16$ frames with a resolution of $H = 224, W = 224$ pixels. These frames are further divided into 16×16 patches. We apply a masking ratio of $R_v = 90\%$, in line with [47]. For the IMU signals, we transform them into spectrograms with a temporal dimension of $T_{imu} = 160$ and a frequency dimension of $M_{imu} = 128$. These spectrograms are similarly divided into 16×16 patches, with a masking ratio of $R_{imu} = 75\%$. For the network architecture, we adapt the designs for the video and IMU encoders from [47] and [22], respectively. The unified encoder consists of a single-layer Transformer. Regarding the training process, we set the loss parameters as follows: $\alpha = 1$, $\beta = 10$, $\gamma = 0.01$, and $\tau = 0.05$.

4.3 Comparison with the state of the art

As presented in Tab. 1, we compare our method with the previous state-of-the-art approaches. The pretraining and finetuning data are from the same dataset and all compared methods share the same data for training. Our experiments are differentiated based on modality, including those that use only IMU or video and those that utilize both modalities for training and testing. For the methods that use Transformer [50] backbones, we all use ViT-Base [12] for a fair comparison. Below, we present detailed analyses of the results in this table.

Within IMU modality For training with IMU signals alone (top block of Tab. 1), we evaluate three prior methods, including DCL [39] and its improved version ADCL [5] that integrates recurrent and convolutional layers to perform supervised pretraining from scratch. Additionally, we evaluate LIMU-BERT [57], which adapts the BERT [11] architecture to process IMU signals, learning general representations from unlabeled IMU data before finetuning with an additional GRU [9] classifier. Compared to these methods, our approach benefits from MAE-based pretraining that learns a stronger IMU representation by reconstructing both raw signals and features. Furthermore, by embedding IMU

Table 1: Comparison with the state of the art. We show action recognition accuracy (mAP \uparrow) on the CMU-MMAC [48] and WEAR [6] datasets.

	Modality	Backbone	Pretrain	CMU-MMAC (mAP)	WEAR (Acc)
DCL [39]	IMU	LSTM	✗	7.46	74.37
ADCL [5]	IMU	LSTM	✗	8.59	77.30
LIMU-BERT [57]	IMU	Transformer	✓	15.30	79.60
Ours	IMU	Transformer	✓	31.68	86.53
SlowFast [14]	Video	ResNet50	✗	82.03	86.13
TimeSformer [4]	Video	Transformer	✗	82.56	86.63
VideoMAE [47]	Video	Transformer	✓	84.63	88.47
Ours	Video	Transformer	✓	85.07	89.78
WEAR+ [6]	IMU-Video	Transformer	✗	83.36	90.14
CAV-MAE [19]	IMU-Image	Transformer	✓	74.69	90.22
AV-MAE [17]	IMU-Video	Transformer	✓	84.75	91.60
Ours	IMU-Video	Transformer	✓	87.96	92.78

features within a graph, our graph network can more effectively leverage the relationships between multiple IMUs. Thanks to these advantages, our method enjoys a significant increase (15.30% \rightarrow 31.68%) against LIMU-BERT [57] in the IMU-only action recognition accuracy. We further show the effect of each component in the following sections.

Within Video modality For training exclusively with video data (mid block of Tab. 1), we assess three methods. SlowFast [14] combines high and low framerate pathways for temporal and spatial semantics with convolutional networks. TimeSformer [4] adapts the Transformer architecture for learning spatiotemporal features from sequences of frame-level patches. In our approach, we design the video branch following the structure proposed in VideoMAE [47]. Our method clearly outperforms SlowFast and TimeSformer, meanwhile slightly surpassing the VideoMAE baseline due to the incorporation of an additional contrastive pretraining objective.

IMU-visual modality For training that incorporates both modalities (bottom block of Tab. 1), we evaluate an IMU-video-based method WEAR+ [6] and upgrade its backbone for a fair comparison. We also adapt and evaluate two audio-visual MAE-based pretraining approaches because audio and IMU signals have similar data format. Given that audio data are preprocessed into one-channel spectrograms, we adapt these methods to accept multi-channel IMU spectrograms by adjusting the network’s channel numbers. CAV-MAE [19] processes a spectrogram and a related image as input, and AV-MAE [17] utilizes audios and the corresponding videos as inputs. Our method surpasses these approaches by employing a graph to model the relationships among multiple collaborative IMU features and by effectively integrating modality fusion during both the pretraining and fine-tuning phases.

Table 2: Ablation studies. The left table outlines the modality selection, pretraining decision, and the employment of the IMU feature graph. The right tables illustrate the effects of different IMU masking ratios, style options, and pre-training objectives.

Modality	Pretrain	Graph	CMU-MMAC	WEAR	IMU mask ratio	60%	75%	90%
IMU	✗	✗	24.87	79.66	IMU	30.38	31.68	31.88
IMU	✗	✓	27.36	83.55	IMU-Video	83.16	87.96	87.62
IMU	✓	✗	28.96	85.29				
IMU	✓	✓	31.68	86.53	IMU mask style	time	freq.	Both
Video	✗	✗	69.82	86.53	IMU	30.51	29.96	31.59
Video	✓	✗	85.07	89.78	IMU-Video	84.63	87.46	86.69
IMU-Video	✗	✗	71.53	90.93				
IMU-Video	✗	✓	73.19	91.54	Supervision	MAE	Contra.	Both
IMU-Video	✓	✗	84.78	91.62	IMU	31.45	14.32	31.68
IMU-Video	✓	✓	87.96	92.78	IMU-Video	86.63	84.65	87.96

4.4 Ablation studies

Modality correlation and pretraining Based on the observation of a strong correlation between ego video and body-worn IMU motion data, we leverage a larger volume of unlabeled data for self-supervised pretraining. As shown in Tab. 2, our experiments validate the efficacy of pretraining, wherein our model successfully learns mutual alignment and completion between the two modalities.

Collaborative IMUs and feature graph Leveraging the collaborative nature of body-worn IMUs for action capture, we embed their features in a graph, employing graph networks to analyze their relationships. As evidenced in Tab. 2, using a graph network boosts action recognition accuracy in a supervised context. Additionally, a node masking strategy in pretraining further strengthens our model’s capability, enhancing its action recognition efficacy.

High masking ratio for both video and IMU VideoMAE [47] adopts a high masking ratio (90%) for pretraining, motivated by the significant correlation and redundancy inherent in video data. We find that a high masking ratio (75%~90%) similarly advantages IMU pretraining, because periodic repetitive motions frequently occur in human activities, as illustrated in Fig. 5.

IMU masking style We experiment with both structured and unstructured random masking techniques. In structured masking, we implement masking across both temporal and frequency-related patches within the IMU spectrogram. We found that unstructured random masking yields marginally better results.

Pretraining supervision We discover that incorporating contrastive loss benefits both single-modality and dual-modality pretraining.

Table 3: IMU device missing challenge. Our method experiences a smaller decrease in accuracy compared to others.

	Modality	Miss	CMU-MMAC	WEAR
ADCL [5]	IMU	✗	8.59	77.30
	IMU	✓	5.25 [↓ 38%]	70.28 [↓ 9%]
LIMU-BERT [57]	IMU	✗	15.30	79.60
	IMU	✓	11.40 [↓ 25%]	73.25 [↓ 8%]
Ours	IMU	✗	31.68	86.53
	IMU	✓	26.47 [↓ 16%]	80.89 [↓ 6%]
Ours	IMU-Video	✗	87.96	92.78
	IMU-Video	✓	85.86 [↓ 3%]	92.42 [↓ 1%]

Table 4: Cross dataset pretraining and finetuning. W : WEAR, C : CMU-MMAC.

Modality	Pretrain	$W \rightarrow C$	$C \rightarrow W$
IMU	✗	27.36	83.55
IMU	✓	28.08	85.27
Video	✗	69.82	86.53
Video	✓	82.32	89.70
IMU-Video	✗	73.19	91.54
IMU-Video	✓	85.67	91.98

Dataset diversity Tab. 2 reveals that the accuracy in the IMU modality of the CMU-MMAC dataset is significantly lower compared to the WEAR dataset. The primary reason is that actions in the CMU-MMAC dataset involve hand-object-environment interactions. However, IMU signals only contain information about hand motion, making it challenging to recognize objects and surroundings.

4.5 Performance under real-world challenges

Since our model employs wearable devices for action recognition, we must account for various challenges in human-centered environments, where individuals might work in diverse environments and have distinct user customizations for their devices. Consequently, we have designed experiments for challenging situations to support more flexible user usages.

IMU device missing challenge When using wearable devices in various scenarios, the number of IMU sensors required may vary due to budget constraints or physical activity limitations. It is assumed that all IMU devices are available during the pretraining phase. Then, we hypothesize a scenario where a user can only afford half the number of IMUs. While finetuning with fewer devices, a robust pretrained model should continue to exhibit advantages. As shown in Tab. 3, we mask IMU signals from ‘left wrist’ and ‘left ankle’ and finetune the models. Due to the loss of information, a natural decrease in action recognition accuracy is observed. However, our method experiences a less significant drop in accuracy compared to other approaches. Because in our pretraining process, we attempt to reconstruct global information from local data, this approach effectively aligns with this challenge.

Feature transferability challenge Datasets that include video and body-worn IMUs may have domain gaps. These gaps arise not only from the potential absence of IMU devices but also from variations in video recording environments and differences in subjects’ actions. For IMUs, discrepancies can occur due to variations in sensor accuracy and range. Particularly, the specific location on

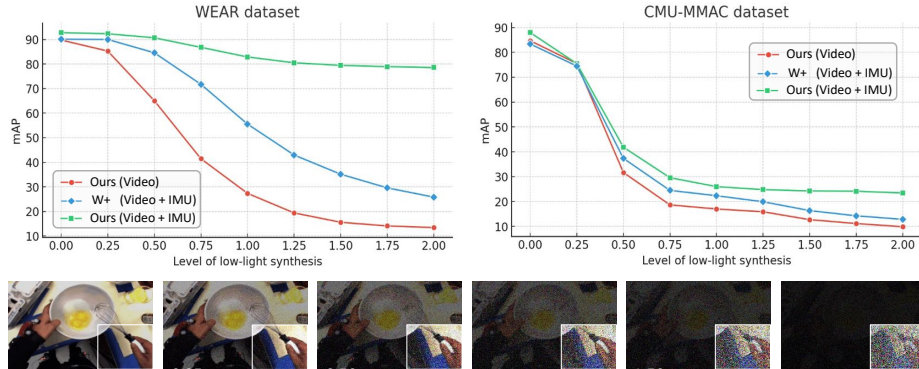


Fig. 4: Visual degradation challenge. We employ a sophisticated method to synthesize low-light effect and degrade the signal-to-noise ratio on the video input. In such cases, our multimodal model switches focus to IMU modality for robust performance. In contrast, simple multimodal feature concatenation (W+ [6]) achieves suboptimal results.

the body where the device is worn and its orientation can result in distinct differences in the data collected. When finetuning on a dataset that differs from the pretraining dataset, a robust pretrained model should continue to exhibit advantages. We verify this with experiments shown in Tab. 4 where we pretrain and finetune in cross dataset settings.

Visual degradation challenge Video models excel at recognizing actions in well-lit environments but struggle in low-light conditions [67], while motion information from IMU signals naturally resists such variations. However, the challenge remains in determining whether a multimodal model can shift reliance to a more reliable modality for action recognition under low-light conditions. Therefore, we devised an experiment in which we degrade the input video quality through a physically-plausible procedure [7, 54], reducing the signal-to-noise ratio and simulating realistic low-light conditions, and then evaluate our multimodal model’s performance. As illustrated in Fig. 4, our (Video + IMU) model demonstrates robustness across various levels of synthesized low-light scenarios. In contrast, using video only or simply concatenating multimodal features [6] achieves suboptimal results. Given that the CMU-MMAC [48] dataset encompasses actions that involve objects and the surroundings, a noticeable performance decline is naturally observed, since IMU signals hardly contain such information.

4.6 Qualitative examples

Fig. 5 shows several examples of challenging cases where current the state-of-the-art video model [47] struggles to handle. However, our multimodal approach, with the aid of pure limb motion information from IMU that is independent of appearance, demonstrates proficient action recognition capabilities. The left

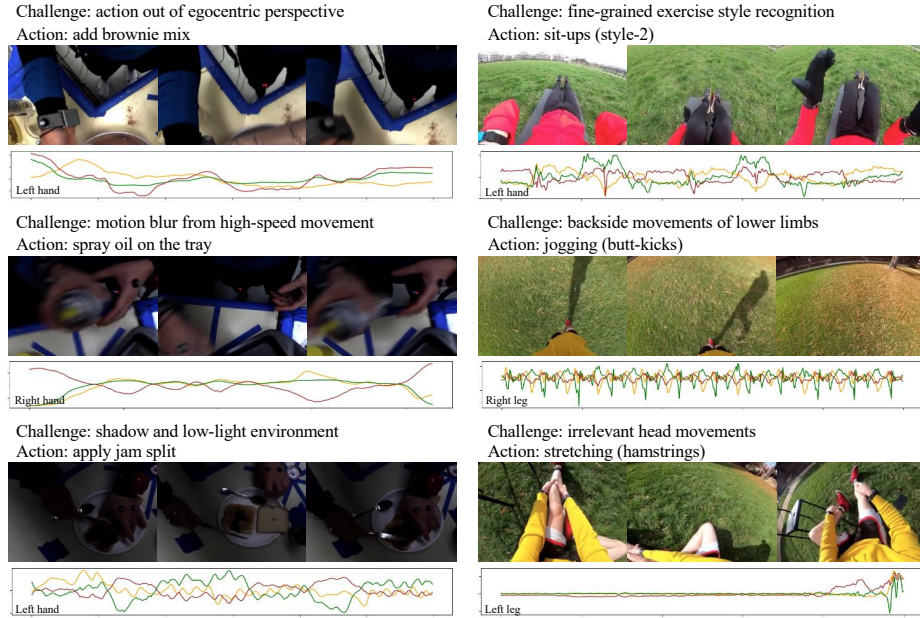


Fig. 5: Visualization instances where our multimodal approach successfully recognizes actions, whereas the VideoMAE [47] model failed.

three challenges are derived from the CMU-MMAC [48] indoor dataset, while the challenges on the right originate from the WEAR [6] outdoor dataset.

5 Conclusion and Future Directions

This paper introduces a novel approach to action recognition leveraging egocentric video and IMU sensors, underpinned by an MAE-based framework enhanced with graph neural networks for more holistic visual and motion representation learning. Our method outperforms existing state-of-the-art models in accuracy and robustness across indoor and outdoor datasets. It enables extensive applications in a variety of fields such as professional athletes, workers, and gamers.

Compared with previous IMU-based action recognition methods [5, 39], our method achieves significant accuracy improvement. Nonetheless, it shares the same level of computational demand and inference time as the latest transformer-based MAE methods. Future research could focus on enhancing the computational efficiency. Another direction is exploring more advanced graph structure designs. We currently utilize a fully connected graph with only four IMU nodes because of the dataset limitation. When more IMUs are available, a more sophisticated graph can be designed to better capture the spatial relations among different body parts.

Acknowledgement

This work was supported by JST ASPIRE Grant Number JPMJAP2303, JST SPRING Grant Number JPMJSP2108, and JSPS KAKENHI Grant Numbers JP22KF0119, JP23H00488 and JP24K02956.

References

1. Abedin, A., Ehsanpour, M., Shi, Q., Rezaatofghi, H., Ranasinghe, D.C.: Attend and discriminate: Beyond the state-of-the-art for human activity recognition using wearable sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* **5**(1), 1–22 (2021) [4](#)
2. Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 6836–6846 (2021) [7](#)
3. Bansal, S., Arora, C., Jawahar, C.: My view is the best view: Procedure learning from egocentric videos. In: *European Conference on Computer Vision*. pp. 657–675. Springer (2022) [9](#)
4. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: *International Conference on Machine Learning*. p. 4 (2021) [4](#), [10](#)
5. Bock, M., Hölzemann, A., Moeller, M., Van Laerhoven, K.: Improving deep learning for har with shallow lstms. In: *Proceedings of the 2021 ACM International Symposium on Wearable Computers*. pp. 7–12 (2021) [9](#), [10](#), [12](#), [14](#)
6. Bock, M., Kuehne, H., Van Laerhoven, K., Moeller, M.: Wear: An outdoor sports for wearable and egocentric activity recognition. *arXiv preprint arXiv:2304.05088* (2023) [2](#), [3](#), [9](#), [10](#), [13](#), [14](#)
7. Brooks, T., Mildenhall, B., Xue, T., Chen, J., Sharlet, D., Barron, J.T.: Unprocessing images for learned raw denoising. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11036–11045 (2019) [13](#)
8. Chen, R., Chen, J., Liang, Z., Gao, H., Lin, S.: Darklight networks for action recognition in the dark. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 846–852 (2021) [2](#)
9. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014) [9](#)
10. Damen, D., Doughty, H., Farinella, G.M., Furnari, A., Kazakos, E., Ma, J., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision* pp. 1–23 (2022) [4](#)
11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018) [4](#), [9](#)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR* (2021) [7](#), [9](#)
13. Du, Y., Wang, W., Wang, L.: Hierarchical recurrent neural network for skeleton based action recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1110–1118 (2015) [3](#)

14. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6202–6211 (2019) 4, 10
15. Feichtenhofer, C., Li, Y., He, K., et al.: Masked autoencoders as spatiotemporal learners. *Advances in Neural Information Processing Systems* **35**, 35946–35958 (2022) 4
16. Geng, X., Liu, H., Lee, L., Schuurmans, D., Levine, S., Abbeel, P.: Multimodal masked autoencoders learn transferable representations. *arXiv preprint arXiv:2205.14204* (2022) 4
17. Georgescu, M.I., Fonseca, E., Ionescu, R.T., Lucic, M., Schmid, C., Arnab, A.: Audiovisual masked autoencoders. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16144–16154 (2023) 2, 4, 10
18. Gong, X., Mohan, S., Dhingra, N., Bazin, J.C., Li, Y., Wang, Z., Ranjan, R.: Mmg-ego4d: Multimodal generalization in egocentric action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6481–6491 (2023) 1, 2, 4
19. Gong, Y., Rouditchenko, A., Liu, A.H., Harwath, D., Karlinsky, L., Kuehne, H., Glass, J.: Contrastive audio-visual masked autoencoder. *arXiv preprint arXiv:2210.07839* (2022) 2, 4, 8, 10
20. Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al.: Ego4d: Around the world in 3,000 hours of egocentric video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18995–19012 (2022) 2, 4
21. Grauman, K., Westbury, A., Torresani, L., Kitani, K., Malik, J., Afouras, T., Ashutosh, K., Baiyya, V., Bansal, S., Boote, B., et al.: Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. *arXiv preprint arXiv:2311.18259* (2023) 4
22. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16000–16009 (2022) 2, 4, 7, 9
23. Hou, Z., Liu, X., Cen, Y., Dong, Y., Yang, H., Wang, C., Tang, J.: Graphmae: Self-supervised masked graph autoencoders. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. pp. 594–604 (2022) 4, 7
24. Huang, P.Y., Xu, H., Li, J., Baevski, A., Auli, M., Galuba, W., Metze, F., Feichtenhofer, C.: Masked autoencoders that listen. *Advances in Neural Information Processing Systems* **35**, 28708–28720 (2022) 4
25. Huang, Y., Cai, M., Li, Z., Lu, F., Sato, Y.: Mutual context network for jointly estimating egocentric gaze and action. *IEEE Transactions on Image Processing* **29**, 7795–7806 (2020) 4
26. Huang, Y., Cai, M., Li, Z., Sato, Y.: Predicting gaze in egocentric video by learning task-dependent attention transition. In: Proceedings of the European conference on computer vision (ECCV). pp. 754–769 (2018) 1
27. Huang, Y., Chen, G., Xu, J., Zhang, M., Yang, L., Pei, B., Zhang, H., Dong, L., Wang, Y., Wang, L., et al.: Egoexolearn: A dataset for bridging asynchronous ego- and exo-centric view of procedural activities in real world. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22072–22086 (2024) 4
28. Huang, Y., Sugano, Y., Sato, Y.: Improving action segmentation via graph-based temporal reasoning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14024–14034 (2020) 1

29. Kazakos, E., Nagrani, A., Zisserman, A., Damen, D.: Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5492–5501 (2019) [1](#), [4](#)
30. Kwon, G., Cai, Z., Ravichandran, A., Bas, E., Bhotika, R., Soatto, S.: Masked vision and language modeling for multi-modal representation learning. arXiv preprint arXiv:2208.02131 (2022) [4](#)
31. Liu, R., Ohkawa, T., Zhang, M., Sato, Y.: Single-to-dual-view adaptation for egocentric 3d hand pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 677–686 (2024) [4](#)
32. Liu, S., Yao, S., Li, J., Liu, D., Wang, T., Shao, H., Abdelzaher, T.: Giobalfusion: A global attentional deep learning framework for multisensor information fusion. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies **4**(1), 1–27 (2020) [4](#)
33. Liu, Y., Liu, Y., Jiang, C., Lyu, K., Wan, W., Shen, H., Liang, B., Fu, Z., Wang, H., Yi, L.: Hoi4d: A 4d egocentric dataset for category-level human-object interaction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21013–21022 (2022) [4](#)
34. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3202–3211 (2022) [7](#)
35. Liu, Z., Shen, Y., Lakshminarasimhan, V.B., Liang, P.P., Zadeh, A., Morency, L.P.: Efficient low-rank multimodal fusion with modality-specific factors. arXiv preprint arXiv:1806.00064 (2018) [4](#)
36. Ma, H., Li, W., Zhang, X., Gao, S., Lu, S.: Attnsense: Multi-level attention mechanism for multimodal human activity recognition. In: IJCAI. pp. 3109–3115 (2019) [4](#)
37. Murahari, V.S., Plötz, T.: On attention models for human activity recognition. In: Proceedings of the 2018 ACM International Symposium on Wearable Computers. pp. 100–103 (2018) [4](#)
38. Nakamura, K., Ohashi, H., Okada, M.: Sensor-augmented egocentric-video captioning with dynamic modal attention. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 4220–4229 (2021) [2](#), [9](#)
39. Ordóñez, F.J., Roggen, D.: Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. Sensors **16**(1), 115 (2016) [3](#), [9](#), [10](#), [14](#)
40. Pang, Y., Wang, W., Tay, F.E., Liu, W., Tian, Y., Yuan, L.: Masked autoencoders for point cloud self-supervised learning. In: European Conference on Computer Vision. pp. 604–621. Springer (2022) [4](#)
41. Poria, S., Chaturvedi, I., Cambria, E., Hussain, A.: Convolutional mkl based multimodal emotion recognition and sentiment analysis. In: 2016 IEEE 16th international conference on data mining (ICDM). pp. 439–448. IEEE (2016) [4](#)
42. Sener, F., Chatterjee, D., Shelepov, D., He, K., Singhania, D., Wang, R., Yao, A.: Assembly101: A large-scale multi-view video dataset for understanding procedural activities. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21096–21106 (2022) [4](#)
43. Somasundaram, K., Dong, J., Tang, H., Straub, J., Yan, M., Goesele, M., Engel, J.J., De Nardi, R., Newcombe, R.: Project aria: A new tool for egocentric multimodal ai research. arXiv preprint arXiv:2308.13561 (2023) [1](#), [4](#)

44. Sudhakaran, S., Escalera, S., Lanz, O.: Lsta: Long short-term attention for egocentric action recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9954–9963 (2019) [4](#)
45. Tang, C.I., Perez-Pozuelo, I., Spathis, D., Mascolo, C.: Exploring contrastive learning in human activity recognition for healthcare. *arXiv preprint arXiv:2011.11542* (2020) [4](#)
46. Tateno, M., Yagi, T., Furuta, R., Sato, Y.: Learning object states from actions via large language models. *arXiv preprint arXiv:2405.01090* (2024) [4](#)
47. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in Neural Information Processing Systems* **35**, 10078–10093 (2022) [4](#), [5](#), [9](#), [10](#), [11](#), [13](#), [14](#)
48. De la Torre, F., Hodgins, J., Bargteil, A., Martin, X., Macey, J., Collado, A., Beltran, P.: Guide to the carnegie mellon university multimodal activity (cmu-mmac) database (2009) [2](#), [3](#), [8](#), [10](#), [13](#), [14](#)
49. Tsutsui, S., Desai, R., Ridgeway, K.: How you move your head tells what you do: Self-supervised video representation learning with egocentric cameras and imu sensors. *arXiv preprint arXiv:2110.01680* (2021) [4](#)
50. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in Neural Information Processing Systems* **30** (2017) [4](#), [9](#)
51. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: *International Conference on Machine Learning*. pp. 1096–1103 (2008) [4](#)
52. Wang, H., Singh, M.K., Torresani, L.: Ego-only: Egocentric action detection without exocentric transferring. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5250–5261 (2023) [4](#)
53. Wang, L., Li, W., Li, W., Van Gool, L.: Appearance-and-relation networks for video classification. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1430–1439 (2018) [4](#)
54. Wei, K., Fu, Y., Yang, J., Huang, H.: A physics-based noise formation model for extreme low-light raw denoising. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2758–2767 (2020) [13](#)
55. Xiao, F., Lee, Y.J., Grauman, K., Malik, J., Feichtenhofer, C.: Audiovisual slowfast networks for video recognition. *arXiv preprint arXiv:2001.08740* (2020) [4](#)
56. Xu, C., Chai, D., He, J., Zhang, X., Duan, S.: Innohar: A deep neural network for complex human activity recognition. *IEEE Access* **7**, 9893–9902 (2019) [4](#)
57. Xu, H., Zhou, P., Tan, R., Li, M., Shen, G.: Limu-bert: Unleashing the potential of unlabeled data for imu sensing applications. In: *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*. pp. 220–233 (2021) [4](#), [9](#), [10](#), [12](#)
58. Xu, J., Huang, Y., Hou, J., Chen, G., Zhang, Y., Feng, R., Xie, W.: Retrieval-augmented egocentric video captioning. *arXiv preprint arXiv:2401.00789* (2024) [1](#)
59. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018) [7](#)
60. Yan, S., Xiong, Y., Lin, D.: Spatial temporal graph convolutional networks for skeleton-based action recognition. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 32 (2018) [3](#)
61. Yang, L., Huang, Y., Sugano, Y., Sato, Y.: Interact before align: Leveraging cross-modal knowledge for domain adaptive action recognition. In: *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14722–14732 (2022) [4](#)
62. Yao, S., Hu, S., Zhao, Y., Zhang, A., Abdelzaher, T.: DeepSense: A unified deep learning framework for time-series mobile sensing data processing. In: Proceedings of the 26th international conference on world wide web. pp. 351–360 (2017) [4](#)
 63. Yao, S., Zhao, Y., Shao, H., Liu, D., Liu, S., Hao, Y., Piao, A., Hu, S., Lu, S., Abdelzaher, T.F.: Sadeepsense: Self-attention deep learning framework for heterogeneous on-device sensors in internet of things applications. In: IEEE INFOCOM 2019-IEEE Conference on Computer Communications. pp. 1243–1251. IEEE (2019) [4](#)
 64. Yuan, H., Chan, S., Creagh, A.P., Tong, C., Clifton, D.A., Doherty, A.: Self-supervised learning for human activity recognition using 700,000 person-days of wearable data. arXiv preprint arXiv:2206.02909 (2022) [4](#)
 65. Zhang, M., Zheng, Y., Lu, F.: Optical flow in the dark. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(12), 9464–9476 (2021) [2](#)
 66. Zhang, S., Ma, Q., Zhang, Y., Qian, Z., Kwon, T., Pollefeys, M., Bogo, F., Tang, S.: Egobody: Human body shape and motion of interacting people from head-mounted devices. In: European Conference on Computer Vision. pp. 180–200. Springer (2022) [4](#)
 67. Zheng, Y., Zhang, M., Lu, F.: Optical flow in the dark. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6749–6757 (2020) [13](#)