

# HAMLET: A Hierarchical Multimodal Attention-based Human Activity Recognition Algorithm

Md Mofijul Islam<sup>1</sup> and Tariq Iqbal<sup>1</sup>

**Abstract**—To fluently collaborate with people, robots need the ability to recognize human activities accurately. Although modern robots are equipped with various sensors, robust human activity recognition (HAR) still remains a challenging task for robots due to difficulties related to multimodal data fusion. To address these challenges, in this work, we introduce a deep neural network-based multimodal HAR algorithm, HAMLET. HAMLET incorporates a hierarchical architecture, where the lower layer encodes spatio-temporal features from unimodal data by adopting a multi-head self-attention mechanism. We develop a novel multimodal attention mechanism for disentangling and fusing the salient unimodal features to compute the multimodal features in the upper layer. Finally, multimodal features are used in a fully connect neural-network to recognize human activities. We evaluated our algorithm by comparing its performance to several state-of-the-art activity recognition algorithms on three human activity datasets. The results suggest that HAMLET outperformed all other evaluated baselines across all datasets and metrics tested, with the highest top-1 accuracy of 95.12% and 97.45% on the UTD-MHAD [1] and the UT-Kinect [2] datasets respectively, and F1-score of 81.52% on the UCSD-MIT [3] dataset. We further visualize the unimodal and multimodal attention maps, which provide us with a tool to interpret the impact of attention mechanisms concerning HAR.

## I. INTRODUCTION

Robots are sharing physical spaces with humans in various collaborative environments, from manufacturing to assisted living to healthcare [4]–[6], to improve productivity and to reduce human cognitive and physical workload [7]. To be effective in close proximity to people, collaborative robotic systems (CRS) need the ability to automatically and accurately recognize human activities [8]. This capability will enable CRS to operate safely and autonomously to work alongside human teammates [9].

To fluently and fluidly collaborate with people, CRS needs to recognize the activities performed by their human teammates robustly [3], [10], [11]. Although modern robots are equipped with various sensors, robust human activity recognition (HAR) remains a fundamental problem for CRS [5]. This is partly because fusing multimodal sensor data efficiently for HAR is challenging. Therefore, to date, many researchers have focused on recognizing human activities by leveraging on a single modality, such as visual, pose or wearable sensors [7], [12]–[15]. However, HAR models reliant on unimodal data often suffer a single point feature representation failure. For example, visual occlusion, poor lighting, shadows, or complex background can adversely

<sup>1</sup>The authors are with the Dept. of Engineering Systems and Environment, Univ. of Virginia, USA. {mi8uu, tiqbal}@virginia.edu.



Fig. 1: Example of two activities (*Sit-Down* and *Carry*) from the UT-Kinect dataset (the first row). The second row presents the temporal-attention weights on the corresponding RGB frames using HAMLET. For these sequences, HAMLET pays more attention to the third RGB image segment for the *Sit-Down* activity (top) and on the fourth RGB image segment for the *Carry* activity (bottom). Here, a lighter color represents a lower attention.

affect only visual sensor-based HAR methods. Similarly, noisy data from accelerometer or gyroscope sensors can reduce the performance of HAR methods solely depending on these sensors [3], [16].

Several approaches have been proposed to overcome the weaknesses of the unimodal methods by fusing multimodal sensor data that can provide complementary strengths to achieve a robust HAR [3], [16]–[20]. Although many of these approaches exhibit robust performances than unimodal HAR approaches, there remain several challenges that prevent these methods from efficiently working on CRSs [16]. For example, while fusing data from multiple modalities, these methods rely on a fixed-fusion approach, e.g., concatenate, average, or sum. Although one type of fusion approach works for a specific activity, these approaches can not provide any guaranty that the same performance can be achieved on a different activity class using the same merging method. Moreover, these proposed approaches provide uniform weightage on the data from all modalities. However, depending on the environment, one sensor modality may provide more enhanced information than the other sensor modality. For example, a visual sensor may provide valuable information about a gross human activity than a gyroscope sensor data, which a robot needs to learn from data automatically. Thus, these approaches can not provide robust HAR for CRSs.

To address these challenges, in this work, we introduce a novel multimodal human activity recognition algorithm, called HAMLET: **H**ierarchical **M**ultimodal **S**elf-attention based HAR algorithm for CRS. HAMLET first extracts the spatio-temporal salient features from the unimodal data for each modality. HAMLET then employs a novel multimodal attention mechanism, called MAT: **M**ultimodal **A**ttention based Feature Fusion, for disentangling and fusing the unimodal features. These fused multimodal features enable

HAMLET to achieve higher HAR accuracies (see Sec. III).

The modular approach to extract spatial-temporal salient features from unimodal data allows HAMLET to incorporate pre-trained feature encoders for some modalities, such as pre-trained ImageNet models for RGB and depth modalities. This flexibility enables HAMLET to incorporate deep neural network-based transfer learning approaches. Additionally, the proposed novel multimodal fusion approach (MAT) utilizes a multi-head self-attention mechanism, which allows HAMLET to be robust in learning weights of different modalities based on their relative importance in HAR from data.

We evaluated HAMLET by assessing its performance on three human activity datasets (UCSD-MIT [3], UTD-MHAD [1] and UT-Kinect [2]) compared with several state-of-the-art activity recognition algorithms from prior literature ([1], [3], [18], [21]–[27]) and two baseline methods (see Sec. IV). In our empirical evaluation, HAMLET outperformed all other evaluated baselines across all datasets and metrics tested, with the highest top-1 accuracy of 95.12% and 97.45% on the UTD-MHAD [1] and the UT-Kinect [2] datasets respectively, and F1-score of 81.52% on the UCSD-MIT [3] dataset (see Sec. V). We visualize an attention map representing how the unimodal and the multimodal attention mechanism impacts multimodal feature fusion for HAR (see Sec. V-D).

## II. RELATED WORKS

**Unimodal HAR:** Human activity recognition has been extensively studied by analyzing and employing the unimodal sensor data, such as skeleton, wearable sensors, and visual (RGB or depth) modalities [28]. As generating hand-crafted features is found to be a difficult task, and these features are often highly domain-specific, many researchers are now utilizing the deep neural network-based approaches for human activity recognition.

Deep learning-based feature representation architectures, especially convolutional neural networks (CNNs) and long-short-term memory (LSTM), have been widely adopted to encode the spatio-temporal features from visual (i.e., RGB and depth) [12], [29]–[33] and non-visual (i.e., sEMG and IMUs) sensors data [3], [7], [34]. For example, Li et al. [29] developed a CNN-based learning method to capture the spatio-temporal co-occurrences of skeletal joints. To recognizing human activities from video data, Wang et al. proposed a 3D-CNN and LSTM-based hybrid model to detect compute salient features [35]. Recently, the graphical convolutional network has been adopted to find spatial-temporal patterns in unimodal data [13].

Although these deep-learning-based HAR methods have shown promising performances in many cases, these approaches rely significantly on modality-specific feature embeddings. If such an encoder fails to encode the feature properly because of noisy data (e.g., visual occlusion or missing or low-quality sensor data), then these activity recognition methods suffer to perform correctly.

**Multimodal HAR:** Many researchers have started working on designing multimodal learning methods by utilizing

the complementary features from different modalities effectively to overcome the dependencies on a single modality data of modality-specific HAR models [17], [18], [36], [37]. One crucial challenge that remains in developing a multimodal learning model is to fuse the various unimodal features efficiently.

Several approaches have been proposed to fuse data from similar modalities [38]–[42]. For example, Simonyan et al. proposed a two-stream CNN-based architecture, where they incorporated a spatial CNN network to capture the spatial features, and another CNN-based temporal network to learn the temporal features from visual data [38]. As CNN-based two-stream network architecture allows to appropriately combine the spatio-temporal features, it has been studied in several recent works, e.g., residual connection in streams [39], convolutional fusion [41] and slow-fast network [33].

Other works have focused on fusing features from various modalities, i.e., fusing features from visual (RGB), pose, and wearable sensor modalities simultaneously [16], [37], [43]. Münzner et al. [19] studied four types of feature fusion approaches: early fusion, sensor and channel-based late fusion, and shared filters hybrid fusion. They found that the late and hybrid fusion outperformed early fusion. Other approaches have focused on fusing modality-specific features at a different level of a neural network architecture [43]. For example, Joze et al. [37] designed an incremental feature fusion method, where the features are merged at different levels of the architecture. Although these approaches have been proposed in the literature, generating multimodal features by dynamically selecting the unimodal features is still an open challenge.

**Attention mechanism for HAR:** Attention mechanism has been adopted in various learning architectures to improve the feature representation as it allows the feature encoder to focus on specific parts of the representation while extracting the salient features [18], [44]–[50]. Recently, several multi-head self-attention based methods have been proposed, which permit to disentangle the feature embedding into multiple features (multi-head) and to fuse the salient features to produce a robust feature embedding [51].

Many researchers have started adopting the attention mechanism in human activity recognition [17], [18]. For example, Xiang et al. proposed a multimodal video classification network, where they utilized an attention-based spatio-temporal feature encoder to infer modality-specific feature representation [18]. The authors explored the different types of multimodal feature fusion approaches (feature concatenation, LSTM fusion, attention fusion, and probabilistic fusion), and found that the concatenated features showed the best performance among the other fusion methods. To date, most of the HAR approaches have utilized attention-based methods for encoding the unimodal features. However, the attention mechanism has not been used for extracting and fusing salient features from multiple modalities.

To address these challenges, in our proposed multimodal HAR algorithm (HAMLET), we have designed a modular way to encode unimodal spatio-temporal features by adopt-




Fig. 2: HAMLET: Hierarchical Multimodal Self-Attention based HAR.

ing a multi-head self-attention approach. Additionally, we have developed a novel multimodal attention mechanism, MAT, for disentangling and fusing the salient unimodal features to compute the multimodal features.

### III. PROPOSED MODULAR LEARNING METHOD

In this section, we present our proposed multimodal human-activity recognition method, called HAMLET: **H**ierarchical **M**ultimodal **S**elf-attention based **H**AR. We present the overall architecture in Fig. 2. In HAMLET, the multimodal features are encoded into two steps, and those features are then used for activity recognition as follows:

- At first, the Unimodal Feature Encoder encodes the spatial-temporal features for each modality by employing a modality-specific feature encoder and a multi-head self-attention mechanism (UAT).
- In the second step, the Multimodal Feature Fusion module (MAT) fuses the extracted unimodal features by applying our proposed novel multimodal self-attention method.
- These computed multimodal features are then utilized by a fully connected neural network to calculate the probability of each activity class.

#### A. Unimodal Feature Encoder

The first step of HAMLET is to compute a feature representation for data from every modality. To achieve that, we have designed modality-specific feature encoders to encode data from different modalities. The main reasoning behind this type of modality-specific modular feature encoder architecture is threefold. First, each of the modalities has different feature distribution and thus needs to have a different feature encoder architecture. For example, the distribution and representation of visual data differ from the skeleton and inertial sensor data. Second, the modular architecture allows incorporating unimodal feature encoders without interrupting the performance of the encoders of other modalities. This capability enables the modality-specific transfer learning. Thus we can employ a pre-trained feature encoder to produce robust feature representation for each modality. Third, the unimodal feature encoders can be trained and executed in parallel, which reduces the computation time during the training and inference phases.

Each of the unimodal feature encoders is divided into three separate sequential sub-modules: spatial feature encoder, temporal feature encoder, and unimodal attention module (UAT). Before applying a spatial feature encoder, at first the whole sequence of data  $D^m = (d_1^m, d_2^m, \dots, d_T^m)$  from

modality  $m$  is converted into segmented sequence  $X^m = (x_1^m, x_2^m, \dots, x_{S^m}^m)$  of size  $B \times S^m \times E^m$ , where  $B$  is the batch size,  $S^m$  and  $E^m$  are the number of segments and feature dimension for modality  $m$  respectively. In this work, we represent the feature dimension  $E^m$  for RGB and depth modality as  $(channel(C^m) \times height(H^m) \times width(W^m))$ , where  $C^m$  is the number of channels in an image.

1) *Spatial Feature Encoder*: We used a temporal pooling method to encode segment-level features instead of extracting the frame-level features, similar to [18]. We have implemented the temporal pooling for two reasons: first, as the successive frames represent similar features, it is redundant to apply spatial feature encoder on each frame, which increases the training and testing time. By Utilizing the temporal pooling, HAMLET reduces its computational time. Moreover, this polling approach is necessary to implement HAMLET on a real-time robotic system. Second, the application of recurrent neural networks for each frame is computationally expensive for a long sequence of data. We used adaptive temporal max-pool to pool the encoded segment level features.

As our proposed modular architecture allows modality-specific transfer learning, we have incorporated the available state-of-the-art pre-trained unimodal feature encoders. For example, we have incorporated ResNet50 to encode the RGB modality. We extend the convolutional co-occurrence feature learning method [29] to hierarchically encode segmented skeleton and inertial sensor data. In this work, we used two stacked 2D-CNNs architecture to encode co-occurrence features: first 2D-CNN encodes the intra-frame point-level information and second 2D-CNN extract the inter-frame features in a segment. Finally, spatial feature encoder for modality  $m$  produces a spatial feature representation  $F_m^S$  of size  $(B \times S^m \times E^{S,m})$  from segmented  $X^m$ , where  $E^{S,m}$  is the spatial feature embedding dimension.

2) *Temporal Feature Encoder*: After encoding the segment level unimodal features, we employ recurrent neural networks, specifically unidirectional LSTM, to extract the temporal feature features  $H^m = (h_1^m, h_2^m, \dots, h_s^m)$  of size  $(B \times S^m \times E^{H,m})$  from  $F_m^S$ , where  $E^{H,m}$  is the LSTM hidden feature dimension. Our choice of unidirectional LSTM over other recurrent neural network architectures (such as gated recurrent units) was based on the ability of LSTM units to capture long-term temporal relationships among the features. Besides, we need our model to detect human activities in real-time, which motivated our choice of unidirectional LSTMs over bi-directional LSTMs.

3) *Unimodal Self-Attention (UAT) Mechanism*: The spatial and temporal feature encoder sequentially encodes the long-range features. However, it cannot extract salient features by employing sparse attention to the different parts of the spatial-temporal feature sequence. Self-attention allows the feature encoder to pay attention to the sequential features sparsely and thus produce a robust unimodal feature encoding. Taking inspiration from the Transformer-based multi-head self-attention methods [51], UAT combines the temporal sequential salient features for each modality. As

each modality has its unique feature representation, the multi-head self-attention enables the UAT to disentangle and attend salient unimodal features.

To compute the attended modality-specific feature embedding  $F_m^a$  for modality  $m$  using unimodal multi-head self-attention method, at first we need to linearly project the spatial-temporal hidden feature embedding  $H^m$  to create query ( $Q_i^m$ ), key ( $K_i^m$ ) and value ( $V_i^m$ ) for head  $i$  in the following way,

$$Q_i^m = H^m W_i^{Q,m} \quad (1)$$

$$K_i^m = H^m W_i^{K,m} \quad (2)$$

$$V_i^m = H^m W_i^{V,m} \quad (3)$$

Here, each modality  $m$  has its own projection parameters,  $W_i^{Q,m} \in \mathbb{R}^{E^{H,m} \times E^K}$ ,  $W_i^{K,m} \in \mathbb{R}^{E^{H,m} \times E^K}$ , and  $W_i^{V,m} \in \mathbb{R}^{E^{H,m} \times E^V}$ , where  $E^K$  and  $E^V$  are projection dimensions,  $E^K = E^V = E^{H,m}/h^m$ , and  $h$  is the total number of heads for modality  $m$ . After that we used scaled dot-product softmax approach to compute the attention score for head  $i$  as:

$$\text{Attn}(Q_i^m, K_i^m, V_i^m) = \sigma \left( \frac{Q_i^m K_i^{mT}}{\sqrt{d_k^m}} \right) V_i^m \quad (4)$$

$$\text{head}_i^m = \text{Attn}(Q_i^m, K_i^m, V_i^m) \quad (5)$$

After that, all the head feature representation is concatenated and projected to produce the attended feature representation,  $F_m^a$  in the following way,

$$F_m^a = [\text{head}_1^m; \dots; \text{head}_h^m] W^{O,m} \quad (6)$$

Here,  $W^{O,m}$  is the projection parameters of size  $E^{H,m} \times E^H$ , and the shape of  $F_m^a$  is  $(B \times S^m \times E^H)$ , where  $E^H$  is the attended feature embedding size. We used the same feature embedding size  $E^H$  for all modalities to simplify the application of multimodal attention MAT for fusing all the modality-specific feature representation, which is presented in the next section III-B. However, our proposed multimodal attention based feature fusion method can handle different unimodal feature dimensions. Finally, we fused the attended segmented sequential feature representation  $F_m^a$  to produce the local unimodal feature representation  $F_m$  of size  $(B \times E^H)$ . We can use different types of fusion to combine the spatio-temporal segmented feature encodings, such as sum, max, or concatenation. However, the concatenation fusion method is not a suitable approach to fuse large sequences, whereas max fusion may lose the temporal feature embedding information. As the sequential feature representations produced from the same modality, we have used the sum fusion approach to fuse attended unimodal spatial-temporal feature embedding  $F_m^a$ ,

$$F_m = \sum_{s \in S^m} F_{m,s}^a \quad (7)$$




Fig. 3: MAT: Multimodal Attention-based Feature Fusion Architecture.

## B. Multimodal Feature Fusion

In this work, we developed a novel multimodal feature fusion architecture based on our proposed multi-head self-attention model, MAT: **M**ultimodal **A**ttention based Feature Fusion, which is depicted in Fig. 3. After encoding the unimodal features using the modular feature encoders, we combine these feature embeddings  $F_m$  in an unordered multimodal feature embedding set  $F^{G^u} = (F_1, F_2, \dots, F_M)$  of size  $(B \times M \times D^H)$ , where  $M$  is the total number of modalities. After that, we fed the set of unimodal feature representations  $F^{G^u}$  into MAT, which produces the attended fused multimodal feature representation  $F^{G^a}$ .

The multimodal multi-head self-attention computation is almost similar to the self-attention method described in Section III-A.3. However, there are two key differences. First, unlike encoding the positional information using LSTM to produce the sequential spatial-temporal feature embedding before applying the multi-head self-attention, in MAT, we combine all the modalities feature embeddings without encoding any positional information. Also, MAT and UAT modules have separate multi-head self-attention parameters. Second, after applying the multimodal attention method on the extracted unimodal features, we used two fusion approaches to fuse the multimodal features:

- MAT-SUM: extracted unimodal features are summed after applying the multimodal attention

$$F^G = \sum_{m=1}^M F_m^{G^a} \quad (8)$$

- MAT-CONCAT: in this approach the attended multimodal features are concatenated

$$F^G = [F_1^{G^a}; F_2^{G^a}; \dots; F_M^{G^a}] \quad (9)$$

## C. Activity Recognition

Finally, the fused multimodal feature representation  $F^G$  is passed through a couple of fully-connected layers to compute the probability for each activity class. For aiding the learning process, we applied activation, dropout, batch normalization in different parts of the learning architecture (see the section IV-B for the implementation details). As all the tasks of human-activity recognition, which we addressed in this work, are multiclass classification, we trained the model using cross-entropy loss function, mini-batch stochastic gradient optimization with weight decay regularization [52].

$$\text{loss}(y, \hat{y}) = \frac{1}{B} \sum_{i=1}^B y_i \log \hat{y}_i \quad (10)$$

TABLE I: Performance comparison (mean top-1 accuracy) of multimodal fusion methods in HAMLET on UT-Kinect dataset [2]

Number of Heads		Fusion Method	
UAT	MAT	MAT-SUM	MAT-CONCAT
1	1	87.97	88.50
<b>1</b>	<b>2</b>	<b>93.50</b>	<b>97.45</b>
2	2	92.50	93.00
2	4	93.50	94.50

#### IV. EXPERIMENTAL SETUP

##### A. Datasets

We evaluated the performance of our proposed multimodal HAR method, HAMLET, using three human-activity datasets: UTD-MHAD [1], UT-Kinect [2], UCSD-MIT [3].

**UTD-MHAD** [1] human activity dataset consists of a total of 27 human actions covering from sports, to hand gestures, to training exercises and daily activities. Eight people repeated each action for four times. After removing the corrupted sequences, this dataset contains a total of 861 data samples.

**UT-Kinect** [2] dataset contains a total of ten indoor daily life activities (e.g., walking, standing up, etc.) with three modalities: RGB, depth, and 3D skeleton. Each activity was performed two times by each person. Thus there were a total of 200 activity samples in this dataset.

**UCSD-MIT** [3] human activity dataset consists of eleven sequential activities in an automotive assembly task. Each assembly task was performed five people, and each person performed the task for five times. This dataset contains three modalities: 3D skeleton data from a motion capture system, and sEMG and IMUs data from a wearable sensor.

##### B. Implementation Details

**Spatial-temporal feature encoder:** We incorporated pre-trained ResNet50 for encoding the RGB and depth data [53]. We applied max pooling with a kernel size of five and stride of three for pooling segment level features. We extended the co-occurrence [29] feature extraction network to encode segmented skeleton and inertial sensor features. Finally, for capturing the temporal features, we used a two-layer unidirectional LSTM. We used embedding size 128 and 256 for UCSD-MIT [3] and UT-Kinect [2] spatial-temporal features embedding respectively.

**Hyper-parameters and optimizer:** We utilized the pre-trained ResNet architecture for encoding RGB and depth modality. However, in the case of a co-occurrence feature encoder (skeleton and inertial sensor), we applied BatchNorm-2D, ReLu activation, and Dropout layers sequentially. After encoding each unimodal features, we applied ReLu activation and Dropout. Finally, in MAT, after fusing the multimodal features, we used BatchNorm-1D, ReLu activation, and Dropout sequentially. We varied the dropout probability between 0.2 – 0.4 in different layers. In multi-head self-attention for both unimodal and multimodal feature encoders, we varied the number of heads from one to eight. We train the learning model using Adam optimizer with weight decay regularization option [52] and cosine annealing warm restarts [54] with an initial learning rate set to  $3e^{-4}$ .

TABLE II: Performance comparison (mean top-1 accuracy) of multimodal HAR methods on UT-Kinect dataset [2]

Method	Fusion Type	Top-1 Accuracy (%)
NSA	SUM	54.34
	CONCAT	52.31
USA	SUM	55.82
	CONCAT	54.34
KEYLESS [18] (2018)	CONCAT	94.50
<b>HAMLET</b>	MAT-SUM	95.56
	<b>MAT-CONCAT</b>	<b>97.45</b>

**Training environment:** We implemented all the parts of the learning model using Pytorch-1.4 deep learning framework [55]. We trained our model in different types of GPU-based computing environments (GPUs: P100, V100, K80, and RTX6000).

##### C. State-of-the-art Methods and Baselines

We designed two baseline HAR methods and reproduce a state-of-art HAR method to evaluate the impact of attention method in encoding and fusing multimodal features:

- **Baseline-1 (NSA)** does not use the attention mechanism for encoding unimodal or fusing multimodal features.
- **Baseline-2 (USA)** only applies multi-head self-attention to encode unimodal features but fuses the multimodal embedding without applying attention. This baseline method is similar to the self-attention based multimodal HAR proposed in [17].
- **Keyless Attention** [18] employed an attention mechanism to encode the modality-specific features. However, it did not utilize attention methods to fuse the multimodal features, instead those were concatenated.

##### D. Evaluation metrics

To evaluate the accuracy of HAMLET, the Keyless Attention model [18], the NSA, and the USA algorithms, we performed leave-one-actor-out cross-validation across all the trials for each person on each dataset. Similar to the original evaluation schemes, we reported activity recognition accuracy for the UT-Kinect [2] and the UTD-MHAD datasets [1], and F1-score (in %) for the UCSD-MIT dataset [3].

To evaluate HAMLET, the Keyless attention method, and baseline methods on UT-Kinect and UTD-MHAD datasets, we used RGB and skeleton data. We leveraged skeleton, IMUs, and sEMG modalities on the UCSD-MIT dataset.

#### V. RESULTS AND DISCUSSION

##### A. Multimodal Attention-based Fusion Approaches

We first evaluated the accuracy of two multimodal attention-based feature fusion approaches of HAMLET: MAT-SUM and MAT-CONCAT. We also varied the number of heads used in UAT and MAT steps to determine the optimal configuration of these values.

**Results:** We evaluated UAT and MAT attention methods as well as the fusion approaches (MAT-SUM and MAT-CONCAT) on the UT-Kinect dataset [2], presented in Table I. We used the RGB and skeleton modalities and reported top-1 accuracy by following the original evaluation scheme. The results suggest that the MAT-CONCAT fusion method

TABLE III: Performance comparison (mean top-1 accuracy) of multimodal fusion methods on UTD-MHAD dataset [1]

Method	Year	Top-1 Accuracy (%)
Kinect & Inertial [1]	2015	79.10
DMM-MFF [27]	2015	88.40
DCNN [26]	2016	91.2
JDM-CNN [25]	2017	88.10
S <sup>2</sup> DDI [22]	2017	89.04
SOS [24]	2018	86.97
MCRL [23]	2018	93.02
PoseMap [21]	2018	94.51
<b>HAMLET (MAT-CONCAT)</b>	-	<b>95.12</b>

showed the highest top-1 accuracy (97.45%), with one and two heads in UAT and MAT methods, respectively.

**Discussion:** The results suggest the concatenation-based fusion approach (MAT-CONCAT) performed better than the summation-based fusion approach (MAT-SUM). Because the MAT-CONCAT allows MAT to disentangle and apply attention mechanisms on the unimodal features to generate robust multimodal features for activity classification. On the other hand, the sum-based fusion method merged the unimodal features into a single representation, which makes it difficult for MAT to disentangle and apply appropriate attention to unimodal features.

The results from Table I also indicate an improvement in activity recognition accuracy with the increment of the number of heads in the MAT when keeping the number of heads fixed in the UAT. However, this relationship does not hold when the number of heads was changed in the UAT. As a large number of heads reduce the size of feature embedding, increasing the number of heads in the UAT may result in an inadequate feature representation. Thus, based on the size of the features used in this work, the results suggest that one head in the UAT and two heads in the MAT methods display the best accuracy. Thus, we utilized these values for further evaluations.

### B. Comparison with Multimodal HAR Methods

As HAMLET takes a multimodal approach, it is reasonable to evaluate the accuracy against the state-of-the-art multimodal approaches. Thus, we compare the performance of HAMLET with two baseline methods (the USA and the NSA, see Sec. IV-C) and several state-of-the-art multimodal approaches. We presented the results in Tables II (UT-Kinect), III (UTD-MHAD) & IV (UCSD-MIT).

**Results:** In the UT-Kinect dataset, RGB and skeleton modalities have been used to train the learning models. Following the original evaluation scheme, we report the top-1 accuracy in Table II. The results indicate that HAMLET achieved the highest 97.45% top-1 accuracy across all other methods.

We also evaluate the performance of HAMLET on the UTD-MHAD [1] dataset. We train and test HAMLET on RGB and Skeleton data and report the top-1 accuracy while using MAT-CONCAT in Table III. The results suggest that HAMLET outperformed all the evaluated state-of-the-art baselines and achieved the highest accuracy of 95.12%.

For the UCSD-MIT dataset, all the learning methods are trained on the skeleton, inertial, and sEMG data. All the

TABLE IV: Performance comparison (mean F1-scores in %) of multimodal HAR methods on UCSD-MIT dataset [3]

Method	Fusion Type	F1-Score (%)
NSA	SUM	59.61
	CONCAT	45.10
USA	SUM	60.78
	CONCAT	69.85
KEYLESS [18] (2018)	CONCAT	74.40
Best of UCSD-MIT [3] (2019)	Early Fusion	59.0
<b>HAMLET</b>	<b>MAT-SUM</b>	<b>81.52</b>
	MAT-CONCAT	76.86

training models have been used late or intermediate fusion except for the results presented from [3], which used an early feature fusion approach. In Table IV, the results suggest that HAMLET with MAT-SUM fusion method outperformed the baselines and state-of-the-art works by achieving the highest 81.52% F1-score (in %).

**Discussion:** HAMLET outperformed all other evaluated baselines across all datasets and metrics tested. The results on the UTD-MHAD dataset suggest that HAMLET outperformed all the state-of-the-art multimodal HAR methods. These methods didn't leverage the attention-based approaches to dynamically weighting the unimodal features to generate multimodal features. The results also suggest that, the other attention-based approaches, such as USA and Keyless [18], also showed better performance compared to the non-attention based approaches on UT-Kinect (Table II) and UCSD-MIT (Table II) datasets. The overall results support that our proposed approach is robust in finding appropriate multimodal features, hence it has achieved the highest HAR accuracies.

The results indicate that the MAT-CONCAT approach achieved higher accuracy on the UT-Kinect dataset; however, the MAT-SUM approach delivered higher accuracy on the UCSD-MIT dataset. One explanation behind this variation is that the modalities (skeleton, sEMG, and IMUs) in the UCSD-MIT dataset represent similar physical body features, thus summing up the feature vectors work well. However, as the UT-Kinect dataset modalities have different characteristics, the visual (RGB) and the physical body (skeleton) features, MAT-CONCAT works better than MAT-SUM.

Finally, the overall results suggest that HAMLET achieved the mean F-1 score of 81.52% on the UCSD-MIT dataset, which is lower compared to the highest accuracy on other datasets (please note that the top-1 accuracies were presented for other datasets). The main reason behind this performance degradation in UCSD-MIT is that this dataset contains missing data, especially sEMG, and IMUs data are missing in many instances. However, in the presence of the missing information, HAMLET showed the best performance compared to all other approaches.

### C. Combined Impact of Unimodal and Multimodal Attention

We evaluated the comparative importance of unimodal and multimodal attention mechanism (presented in Fig. 4). We can observe that the incorporation of unimodal attention (Fig. 4-b) can help to reduce the miss-classification error in comparison to the non-attention based feature learning

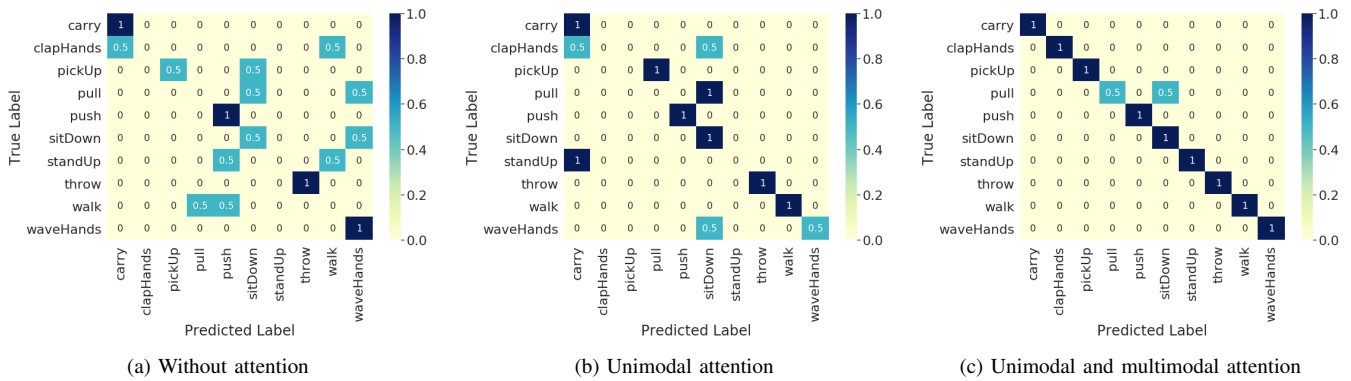


Fig. 4: Comparative impact of multimodal and unimodal attention in HAMLET for different activities on UT-Kinect dataset.




Fig. 5: Multimodal and unimodal attention visualization for different activities on UT-Kinect Dataset.

method (Fig. 4-a). This is because unimodal attention can able to extract the sparse salient spatio-temporal features. We also can observe an improved accuracy in activity classification when the multimodal attention based unimodal feature fusion approach was incorporated (Fig. 4-c vs. a, b). The results indicate that HAMLET can reduce the number of miss-classification, especially in the cases of similar activities, such as *sitDown* and *pickUp*, which is depicted in the confusion matrix in Fig. 4-c.

#### D. Visualizing Impact of Multimodal Attention: MAT

We visualize the attention map of the unimodal and multimodal feature encoders to gauge the impact of attention in local (unimodal) and global (multimodal) feature representation in Fig 5. We used the data of the eighth performer from the UT-Kinect dataset [2] as a sample data to produce the attention map for different activities, as shown in Fig. 5, where we observe that the unimodal attention is able to detect salient segments of RGB (Fig 5-a) and skeleton (Fig 5-b) modalities. For example, the unimodal attention method focuses on the beginning parts of the *sitDown* and the *pull* activities, as these activities have distinguishable actions in the beginning parts of the activity. On the other hand, the unimodal attention method needs to pay attention to the full sequence to differentiate the *carry* and the *push* activities, as a specific part of these activities are not more informative than the other parts.

Moreover, we evaluate the impact of MAT by observing the multimodal attention map in Fig. 5-c, which represents the relative attention given to unimodal features. For example, the *pickUp* and *sitDown* may involve similar skeleton

joint movements, and thus if we concentrate only on the skeleton data, it may be challenging to differentiate between these two activities. However, if we incorporate the complementary modalities, such as RGB and skeleton, it may be easier to differentiate between similar activities. Thus, MAT pays equal attention to the RGB and skeleton data while recognizing the *sitDown* activity, whereas solely pay attention to the skeleton data while identifying the *pickUp* activity (Fig. 5-c).

## VI. CONCLUSION

In this paper, we presented HAMLET, a novel multimodal human activity recognition algorithm, for collaborative robotic systems. HAMLET first extracts the spatio-temporal salient features from the unimodal data and then employs a novel multimodal attention mechanism for disentangling and fusing the unimodal features for activity recognition. The experimental results suggest that HAMLET outperformed all other evaluated baselines across all datasets and metrics tested for human activity recognition.

In the future, we plan to implement HAMLET on a robotic system to enable it to perform collaborative activities in close proximity with people in an industrial environment. We also plan to extend HAMLET so that it can appropriately learn the relationship among the data from the modalities to address the missing data problem.

## REFERENCES

- [1] C. Chen, R. Jafari, and N. Kehtarnavaz, "Utd-mhad: A multimodal dataset for human action recognition utilizing a depth camera and a wearable inertial sensor," in *2015 IEEE ICIP*, Sep. 2015, pp. 168–172.

- [2] L. Xia, C. Chen, and J. Aggarwal, "View invariant human action recognition using histograms of 3d joints," in *CVPRW*. IEEE, 2012, pp. 20–27.
- [3] A. Kubota, T. Iqbal, J. A. Shah, and L. D. Riek, "Activity recognition in manufacturing: The roles of motion capture and semg+ inertial wearables in detecting fine vs. gross motion," in *2019 ICRA*. IEEE, 2019, pp. 6533–6539.
- [4] L. Riek, "Healthcare robotics," *Communications of the ACM*, 2017.
- [5] T. Iqbal and L. D. Riek, "Human-robot teaming: Approaches from joint action and dynamical systems," *Humanoid robotics: A reference*, pp. 2293–2312, 2019.
- [6] T. Iqbal, S. Rack, and L. D. Riek, "Movement coordination in human-robot teams: A dynamical systems approach," *IEEE Transactions on Robotics*, vol. 32, no. 4, pp. 909–919, 2016.
- [7] A. E. Frank, A. Kubota, and L. D. Riek, "Wearable activity recognition for robust human-robot teaming in safety-critical environments via hybrid neural networks," in *IEEE/RSJ IROS*, 2019, pp. 449–454.
- [8] T. Iqbal, S. Li, C. Fourie, B. Hayes, and J. A. Shah, "Fast online segmentation of activities from partial trajectories," in *2019 ICRA*, May 2019, pp. 5019–5025.
- [9] T. Iqbal and L. D. Riek, "Coordination dynamics in multihuman multirobot teams," *IEEE RA-L*, vol. 2, no. 3, pp. 1712–1717, 2017.
- [10] T. Iqbal, M. J. Gonzales, and L. D. Riek, "Joint action perception to enable fluent human-robot teamwork," in *2015 24th IEEE RO-MAN*, Aug 2015, pp. 400–406.
- [11] T. Iqbal and L. D. Riek, "A Method for Automatic Detection of Psychomotor Entrainment," *IEEE Transactions on Affective Computing*, vol. 7, no. 1, pp. 3–16, 2016.
- [12] Q. Ke, M. Bennamoun, S. An, F. Sohel, and F. Boussaid, "A new representation of skeleton sequences for 3d action recognition," in *CVPR*, 2017, pp. 3288–3297.
- [13] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *Thirty-second AAAI conference on artificial intelligence*, 2018.
- [14] F. Han, B. Reily, W. Hoff, and H. Zhang, "Space-time representation of people based on 3d skeletal data: A review," *Computer Vision and Image Understanding*, vol. 158, pp. 85–105, 2017.
- [15] T. Iqbal, M. Moosaie, and L. D. Riek, "Tempo adaptation and anticipation methods for human-robot teams," in *RSS, Planning HRI: Shared Autonomy Collab. Robot. Workshop*, 2016.
- [16] T. Baltrušaitis, C. Ahuja, and L. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 2, pp. 423–443, 2019.
- [17] G. Liu, J. Qian, F. Wen, X. Zhu, R. Ying, and P. Liu, "Action recognition based on 3d skeleton and rgb frame fusion," in *2019 IEEE/RSJ IROS*, Nov 2019, pp. 258–264.
- [18] X. Long, C. Gan, G. De Melo, X. Liu, Y. Li, F. Li, and S. Wen, "Multimodal keyless attention fusion for video classification," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [19] S. Münzner, P. Schmidt, A. Reiss, M. Hanselmann, R. Stiefelhagen, and R. Dürichen, "Cnn-based sensor fusion techniques for multimodal human activity recognition," in *Proceedings of the 2017 ACM ISWC*, 2017, p. 158–165.
- [20] M. K. Hasan, W. Rahman, A. Bagher Zadeh, J. Zhong, M. I. Tanveer, L.-P. Morency, and M. E. Hoque, "Ur-funny: A multimodal language dataset for understanding humor," *EMNLP-IJCNLP*, 2019.
- [21] M. Liu and J. Yuan, "Recognizing human actions as the evolution of pose estimation maps," in *CVPR*, 2018, pp. 1159–1168.
- [22] P. Wang, S. Wang, Z. Gao, Y. Hou, and W. Li, "Structured images for rgb-d action recognition," in *CVPRW*, 2017, pp. 1005–1014.
- [23] T. Liu, J. Kong, and M. Jiang, "Rgb-d action recognition using multimodal correlative representation learning model," *IEEE Sensors Journal*, vol. 19, no. 5, pp. 1862–1872, 2019.
- [24] Y. Hou, Z. Li, P. Wang, and W. Li, "Skeleton optical spectra-based action recognition using convolutional neural networks," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 28, no. 3, pp. 807–811, 2016.
- [25] C. Li, Y. Hou, P. Wang, and W. Li, "Joint distance maps based action recognition with convolutional neural networks," *IEEE Signal Processing Letters*, vol. 24, no. 5, pp. 624–628, 2017.
- [26] J. Imran and P. Kumar, "Human action recognition using rgb-d sensor and deep convolutional neural networks," in *ICACCI*, 2016.
- [27] M. F. Bulbul, Y. Jiang, and J. Ma, "Dmms-based multiple features fusion for human action recognition," *IJMDEM*, vol. 6, no. 4, pp. 23–39, 2015.
- [28] M. E. Hussein, M. Torki, M. A. Gawayyed, and M. El-Saban, "Human action recognition using a temporal hierarchy of covariance descriptors on 3d joint locations," in *Twenty-Third AAAI*, 2013.
- [29] C. Li, Q. Zhong, D. Xie, and S. Pu, "Co-occurrence feature learning from skeleton data for action recognition and detection with hierarchical aggregation," in *IJCAI*, 2018, p. 786–792.
- [30] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, "A closer look at spatiotemporal convolutions for action recognition," in *CVPR*, 2018, pp. 6450–6459.
- [31] B. Zhou, A. Andonian, A. Oliva, and A. Torralba, "Temporal relational reasoning in videos," in *ECCV*, 2018, pp. 803–818.
- [32] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *ICCV*, 2015, pp. 4489–4497.
- [33] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," in *2019*, 2019.
- [34] M. S. Totty and E. Wade, "Muscle activation and inertial motion data for noninvasive classification of activities of daily living," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 5, pp. 1069–1076, 2017.
- [35] X. Wang, L. Gao, J. Song, and H. Shen, "Beyond frame-level cnn: saliency-aware 3-d cnn with lstm for video action recognition," *IEEE Signal Processing Letters*, vol. 24, no. 4, pp. 510–514, 2016.
- [36] N. C. Garcia, P. Morerio, and V. Murino, "Modality distillation with multiple stream networks for action recognition," in *ECCV*, 2018.
- [37] H. R. V. Joze, A. Shaban, M. L. Iuzzolino, and K. Koishida, "MMTM: Multimodal transfer module for cnn fusion," in *CVPR*, 2020.
- [38] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *NeurIPS*, 2014, pp. 568–576.
- [39] C. Feichtenhofer, A. Pinz, and R. P. Wildes, "Spatiotemporal residual networks for video action recognition," in *Proceedings of the 30th NeurIPS'16*. Red Hook, NY, USA: Curran Associates Inc., 2016, p. 3476–3484.
- [40] —, "Spatiotemporal multiplier networks for video action recognition," in *CVPR*, 2017, pp. 4768–4777.
- [41] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Convolutional two-stream network fusion for video action recognition," in *CVPR*, 2016, pp. 1933–1941.
- [42] S. Zhang, Y. Yang, J. Xiao, X. Liu, Y. Yang, D. Xie, and Y. Zhuang, "Fusing geometric features for skeleton-based action recognition using multilayer lstm networks," *IEEE Transactions on Multimedia*, vol. 20, no. 9, pp. 2330–2343, 2018.
- [43] J.-M. Perez-Rua, V. Vielzeuf, S. Pateux, M. Baccouche, and F. Jurie, "Mfas: Multimodal fusion architecture search," in *CVPR*, June 2019.
- [44] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *ICLR*, 2015.
- [45] M.-T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," *arXiv preprint arXiv:1508.04025*, 2015.
- [46] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *ICML*, 2015, pp. 2048–2057.
- [47] J. Lu, C. Xiong, D. Parikh, and R. Socher, "Knowing when to look: Adaptive attention via a visual sentinel for image captioning," in *CVPR*, 2017, pp. 375–383.
- [48] V. Mnih, N. Heess, A. Graves, et al., "Recurrent models of visual attention," in *NeurIPS*, 2014, pp. 2204–2212.
- [49] J. Lu, D. Batra, D. Parikh, and S. Lee, "Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks," in *NeurIPS*, 2019.
- [50] P. Gao, H. You, Z. Zhang, X. Wang, and H. Li, "Multi-modality latent interaction network for visual question answering," in *ICCV*, 2019.
- [51] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, vol. 2017-December, no. NeurIPS, pp. 5999–6009, 2017.
- [52] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.
- [53] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.
- [54] I. Loshchilov and F. Hutter, "Sgdr: Stochastic gradient descent with warm restarts," in *ICLR*, 2017.
- [55] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., "Pytorch: An imperative style, high-performance deep learning library," in *NeurIPS*, 2019, pp. 8024–8035.