Sensor-Augmented Egocentric-Video Captioning with Dynamic Modal Attention

Katsuyuki Nakamura
Hitachi, Ltd.
Tokyo, Japan
katsuyuki.nakamura.xv@hitachi.com

Hiroki Ohashi
Hitachi, Ltd.
Tokyo, Japan
hiroki.ohashi.uo@hitachi.com

Mitsuhiro Okada
Hitachi, Ltd.
Tokyo, Japan
mitsuhiro.okada.uf@hitachi.com

Figure 1: (a)–(b) First-person view is effective for observing more detailed hand operations compared with third-person view. (c)–(d) First-person view can be affected by harmful situations. (e) Sensor data contain noise caused by electromagnetic cooker. (f) The subject reaches a hand for the fridge while walking, where “walking” is the ground-truth activity. ©Ekaterina H. Spriggs.

ABSTRACT

Automatically describing video, or video captioning, has been widely studied in the multimedia field. This paper proposes a new task of sensor-augmented egocentric-video captioning, a newly constructed dataset for it called MMAC Captions, and a method for the newly proposed task that effectively utilizes multi-modal data of video and motion sensors, or inertial measurement units (IMUs). While conventional video captioning tasks have difficulty in dealing with detailed descriptions of human activities due to the limited view of a fixed camera, egocentric vision has greater potential to be used for generating the finer-grained descriptions of human activities on the basis of a much closer view. In addition, we utilize wearable-sensor data as auxiliary information to mitigate the inherent problems in egocentric vision: motion blur, self-occlusion, and out-of-camera-range activities. We propose a method for effectively utilizing the sensor data in combination with the video data on the basis of an attention mechanism that dynamically determines the modality that requires more attention, taking the contextual information into account. We compared the proposed sensor-fusion

method with strong baselines on the MMAC Captions dataset and found that using sensor data as supplementary information to the egocentric-video data was beneficial, and that our proposed method outperformed the strong baselines, demonstrating the effectiveness of the proposed method.

CCS CONCEPTS

• Computing methodologies → Activity recognition and understanding.

KEYWORDS

Video captioning, egocentric vision, sensor fusion

ACM Reference Format:

Katsuyuki Nakamura, Hiroki Ohashi, and Mitsuhiro Okada. 2021. Sensor-Augmented Egocentric-Video Captioning with Dynamic Modal Attention. In *Proceedings of the 29th ACM Int'l Conference on Multimedia (MM '21)*, Oct. 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3474085.3475557>

1 INTRODUCTION

Video captioning is an active research topic in the multimedia field. Most studies on this topic have attempted to describe events using third-person video data [9], but these studies have demonstrated the difficulty in describing finer-grained activities such as cooking because these activities require more detailed information than third-person video data. Egocentric video, which is recorded by using a wearable camera, is useful for this purpose because it can

This work is licensed under a Creative Commons Attribution International 4.0 License.

MM '21, October 20–24, 2021, Virtual Event, China.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8651-7/21/10.

<https://doi.org/10.1145/3474085.3475557>

produce much finer-grained information on the activities of the camera wearer on the basis of the closer view (Figures 1 (a)–(b)). However, video captioning based on the egocentric video is not trivial because it, in nature, often suffers from motion blur, self-occlusion, and out-of-camera-range activities (Figures 1 (c)–(d)). It is therefore better to utilize supplementary information to overcome these difficulties.

This paper proposes a new task of sensor-augmented egocentric-video captioning for the finer-grained description of human activities. It utilizes egocentric-video data along with motion data from wearable sensors, or inertial measurement unit (IMU) sensors. For this task, we constructed a new dataset called MMAC Captions¹ by extending the CMU-Multimodal Activity (CMU-MMAC) dataset [70]. MMAC Captions contains 16 hours of egocentric-video and wearable-sensor data with more than 5,000 egocentric activity descriptions.

One of the difficulties of the proposed task is how to fuse the differing video and sensor modalities. Although both of these complementary modalities are useful in many situations, one modality could have a negative impact on performance. For example, sensor data may occasionally contain undesirable noise caused by surrounding objects or by irrelevant motions to the target activity (Figures 1 (e)–(f)). In such cases, not using sensor data or assigning less weights to it is preferable. To this end, we propose a dynamic modal attention (DMA) module, which determines the modalities to be emphasized more (or less) depending on the context.

We compared the proposed sensor-fusion method with strong baselines on the MMAC Captions dataset. The experimental results showed that the sensor data was beneficial as supplementary information to the egocentric video data, and that our DMA module effectively determined the modalities to be emphasized.

The key contributions of this study are summarized as follows. (1) We propose a new task of sensor-augmented egocentric-video captioning for the finer-grained description of human activities, and provide a newly constructed MMAC Captions dataset that contains 16 hours of egocentric video and sensor data with more than 5,000 egocentric activity descriptions. (2) We propose a DMA module that determines the modalities to be emphasized more (or less) depending on the contexts. (3) We experimentally demonstrate the effectiveness of using sensor data and that of the DMA module through comparisons with strong baseline methods.

This paper is structured as follows. Section 2 reviews existing research on activity recognition. Section 3 introduces our MMAC Captions dataset. Sections 4 and 5 describe the method for sensor-augmented video captioning and its performance evaluation, respectively.

2 RELATED WORK

Image and Video Captioning. There has been much research on generating textual descriptions from images and videos [9]. One approach is phrase-based captioning, in which an image is first converted into several phrases using object recognition and action recognition, then connecting them with textual modeling [24, 38–40, 58, 66, 72, 74, 75]. This approach can avoid critical mistakes because it uses a prefixed word sequence; however, it is difficult to

generate natural sentences. Another approach is to use an encoder-decoder model for image and video captioning [4, 12, 17, 26, 33, 37, 45, 49, 60, 76, 77, 79, 81, 83, 84]. The pioneering work by Venugopalan *et al.* [77] uses a sequence-to-sequence model that combines a convolutional neural network (CNN)-based image encoder and long short-term memory (LSTM)-based textual decoder. Many studies have proposed extensions to this work from various perspectives to improve performance; selective encoding [13], boundary-aware encoding [4], adaptive attention [46], word-transition molding [11, 34], and object relational modeling [1, 59, 87, 88]. Transformer [63] has been proven to be more effective for video captioning [8, 23, 32, 42, 59, 71, 89]. Zhou *et al.* [89] first proposed the end-to-end transformer model for dense video captioning. Lei *et al.* [42] introduced a memory-augmented recurrent transformer based on bidirectional encoder representations from transformers (BERT) [16].

Multi-modal Captioning. Several works have explored captioning based on multi-modal information such as that from video and audio information [65], and multivariate well logs [73]. One of the typical methods for combining multi-modal information is to use multi-modal attention [28, 64, 82]. The attention mechanism has been further explored with the transformer architecture for multi-modal video captioning. Iashin *et al.* proposed a feature transformer for video and audio input [30] and bi-modal attention for both modalities [29].

Our work expands upon these prior works and exploits the attention-based modality-integration module for dynamically changing emphasis on different video and sensor modalities depending on the contexts.

Activity Recognition using Wearable Camera and Sensors.

There is a substantial body of work on egocentric vision, ranging from activity recognition to video summarization [6, 54]. One of the following cues is commonly used in activity recognition; a motion cue, which uses coherent motion patterns for egocentric video [35, 62, 67, 69, 85], an object cue, which uses objects in video sequences [14, 19, 21, 22, 25, 41, 80], or an integrated cue, which uses both motion and object information [2, 20, 43, 44, 47, 57]. Much effort has been put into constructing applications [7, 18, 50, 53, 57, 86] and datasets [15, 36, 51, 68, 70], too.

Several studies have used non-visual wearable sensors for activity recognition, including accelerometers and heart rate sensors [51, 52, 55, 56]. Bao *et al.* [3] used multiple accelerometers on the hip, wrist, arm, ankle, and thigh to classify 20 classes of household activities. Spriggs *et al.* [70] used an egocentric camera, IMUs, and other sensors to classify 29 classes of kitchen activities. Maekawa *et al.* [48] used a wrist-mounted camera and sensors to detect activities in daily life. The present study extends these prior works to tackle a newly proposed task of sensor-augmented egocentric-video captioning.

3 MMAC CAPTIONS DATASET

The CMU-MMAC dataset [70] is the first dataset that augments egocentric video with wearable IMU data. The sampling rate is 30 fps with a resolution of both 800×600 and 1024×768 for egocentric

¹The dataset will be released to facilitate further research.

Figure 2: Statistics of MMAC Captions: (a) distribution of verb classes, (b) distribution of noun classes, (c) distribution of word-types at each word step, and (d) the scatter plot for standard deviations of gyroscope in both hands.

video and a maximum of 125 Hz for 9-axes IMUs (a 3-axes gyroscope, 3-axes accelerometer, and 3-axes compass). 9 IMUs were attached to the subject’s body: both forearms/upper arms, both thighs/lower limbs, and back. A Point Grey FL2-08S2C-C camera was used for capturing egocentric video. To construct this dataset, 43 participants wore a wearable camera and IMU sensors and performed five types of cooking activities in a kitchen (*i.e.*, *Brownie*: making brownies, *Salad*: preparing a salad, *Pizza*: making pizza, *Eggs*: frying eggs, and *Sandwich*: making a sandwich). Although several types of annotations including object and action classes are provided², detailed activity descriptions are not.

We therefore introduce a new dataset called **MMAC Captions**, which contains 5,002 activity descriptions for 16 hours of egocentric data. Two annotators were independently asked to first define a “segment”, which corresponds to a short period that contains only one activity, by determining the timestamps of the start and end of the segment. Second, they were asked to provide one sentence per segment. Finally, cross checking was carried out to integrate multiple captions into one consistent sentence after each annotator finished their own annotations. We omit the subject of the sentences in the annotations because the subject is always the person who wears the sensors and camera. The dataset will be made publicly available.

A number of activity description examples are shown as follows.

- *Spreading tomato sauce on pizza crust with a spoon.*
- *Taking a fork, knife, and peeler from a drawer.*
- *Cutting a zucchini in half with a kitchen knife.*

²<http://kitchen.cs.cmu.edu/>.

- *Moving a paper plate slightly to the left.*

- *Stirring brownie batter with a fork.*

We found that the provided annotations were much more diverse than those of the action or object categories in a sense that there exist multiple objectives, adverbs, and a wide distribution of sentence lengths.

We used the Natural Language Toolkit (NLTK)³ and TreeTagger⁴ to analyze the dataset, resulting in observing 131 noun classes and 89 verb classes as shown in Figures 2 (a) and (b). Figure 2 (c) shows the distribution of word-types at each word step. All the sentences start with a verb, and a determiner and noun pair follows in many cases. The average sentence length is seven words, and the maximum length is 14 words. The average duration of a segment is 6.7 seconds, and the median is 3.0 seconds. Figure 2 (d) shows that sensor data from an IMU, such as the gyroscope, are useful to predict activities. For example, *shaking* and *turning* take a significantly higher variance than other activities, suggesting that IMU signals are effective for activity description.

4 METHOD

The flow of our approach is shown in Figure 3. Our goal is to describe activities using egocentric-video and wearable-sensor data. To achieve this goal, we extend the boundary-aware neural encoding model [4] for multi-modal activity descriptions.

The core idea includes two key aspects. First, we model a sensor encoder with recurrent neural networks and introduce a learnable

³<https://www.nltk.org/>

⁴<https://www.cis.uni-muenchen.de/~schmid/tools/TreeTagger/>

Figure 3: Flow of proposed approach, which involves inputting egocentric video and sensor signals, and outputting the activity description from an egocentric perspective. Visual and sensor encoders individually output their representations. The AMMT module adaptively integrates these representations. A DMA module dynamically determines the best fusion of different representations for generating each word in a sentence.

transformation module to effectively integrate multi-modal features. Second, we propose a step-wise modality-attention module to dynamically determine the best fusion of different representations for generating each word in a sentence.

4.1 Multi-modal Encoder

The input to our model is a sequence of video frames $\mathcal{V} = \{v_1, v_2, \dots, v_K\}$ and wearable-sensor signals $\mathcal{S} = \{s_1, s_2, \dots, s_T\}$, where v_t is a video frame at time t , and s_t is a 63-dimensional vector at time t from 9-axis IMUs attached at seven locations on the body⁵.

Visual Encoder. We developed a visual encoder ϕ_V based on the boundary-aware neural encoder [4] that can identify discontinuity points between frames such as appearance change and motion change. Our intuition is that such discontinuity frequently occurs in egocentric videos because there are many undesirable imaging situations such as motion blur, frame out, illuminance change, and self-occlusion. Therefore, we assumed a boundary detector would work efficiently. Given a video sequence, we compute a visual representation as $h_V = \phi_V(\mathcal{V})$, where h_V represents the final hidden state of the LSTM layer.

Sensor Encoder. Unlike in ϕ_V , we do not use the boundary detector in sensor encoder ϕ_S . In many cases, boundary detection occurs too often during the sensor-encoding step, which results in an enormous state initialization. We therefore developed a neural encoder based on the LSTM layers without boundary-aware encoding. Raw sensor signals are first resampled at 30 Hz using a piecewise linear interpolation and directly fed into the LSTM input layer. With the sensor signals, we compute sensor representation as $h_S = \phi_S(\mathcal{S})$, where h_S is also the final hidden state of the LSTM layer.

Asymmetric Multi-modal Transformation (AMMT). Once the representations of video and sensor have been encoded, the decoder can generate a textual description by simply concatenating both representations (*i.e.*, $h_{V+S} = h_V \oplus h_S$).

However, the simple concatenation may not be optimal because the amount of useful information contained in each representation vector may be significantly different due to the difference in their dimensions, or the length of the representation vectors. To mitigate the possible imbalance, we introduce a learnable transformation module called the asymmetric multi-modal transformation (AMMT), which is inspired by feature-wise linear modulation [61]. AMMT is defined as follows:

$$h_{V+S} = h_V \oplus (W_c h_S + b_c), \quad (1)$$

where W_c is the weight matrix representing a linear transformation and b_c is a bias. The W_c and b_c are learnable vectors, which are initially set to identity transformation (*i.e.*, $W_c = I$ and $b_c = 0$). In other words, we use a standard concatenation in the initial phase and update parameters during the training phase.

The intuition behind the design of asymmetry, or applying linear transformation only to h_S , is two-fold. First, the mitigation of the aforementioned possible imbalance can be achieved by applying the transformation to one of the representations rather than both; the latter has a risk of over-fitting due to the redundant parameters. Second, because the sensor data sometimes contain undesirable noise, which is explained in Section 1, adjusting the sensor representation so that it does not adversely affect the performance is preferable. In addition to this integrated representation h_{V+S} , AMMT forwards h_V and h_S as well because there are a number of cases where using only single modality is preferable (*e.g.*, sensor data containing undesirable noise).

4.2 Multi-modal Decoder

Sentence Generation. Given the aforementioned representations from the multi-modal encoder, our decoder generates a sentence $\hat{\mathcal{Y}} = \{\hat{y}_0, \hat{y}_1, \hat{y}_2, \dots, \hat{y}_N\}$, where $\hat{y}_i$ is a word with one-hot-vector encoding at step i . The objective function for sentence generation is

$$\theta^* = \arg \max_{\theta} \sum_{i=1}^N \log \Pr(\hat{y}_i | \hat{y}_{i-1}, \hat{y}_{i-2}, \dots, \hat{y}_0, z; \theta), \quad (2)$$

where θ denotes all parameters in the model, and z is a set of representations from the encoder. We compute the probability of a

⁵There were originally nine locations, but the data collected from two sensors were incomplete, and were therefore not included in the CMU-MMAC dataset.

word as follows:

$$\Pr(\hat{y}_i | \hat{y}_{i-1}, \hat{y}_{i-2}, \dots, \hat{y}_0, \mathbf{z}; \theta) \propto \hat{y}_i^T \text{softmax}(\mathbf{W}_h \mathbf{h}_i), \quad (3)$$

where $\mathbf{W}_h$ is a matrix to convert a hidden state to the same dimension as word vector $\hat{y}_i$, and $\mathbf{h}_i$ is an output hidden state of decoder at word step i by using gated recurrent unit (GRU). We use the special token <BOS> as the initial word $\hat{y}_0$ and <EOS> as the termination symbol, which enables variable-length sentences to be generated.

Dynamic Modal Attention (DMA). DMA takes representation vectors given by AMMT as input, and outputs the best-integrated features at each decoding step on the basis of attention on different types of representations.

Given representations $\mathbf{z} = (\mathbf{h}_V, \mathbf{h}_S, \mathbf{h}_{V+S})^T$ from AMMT, DMA takes the form of the weighted sum of each representation as follows.

$$g(\mathbf{z}) = \sum_k \zeta_{k,i} \mathbf{W}_k \mathbf{h}_k, \quad (4)$$

where $k \in \{V, S, V+S\}$, $\mathbf{W}_k$ are the linear transformations for adjusting the dimensions so that all representation vectors have the same length, and $\{\zeta_{k,i} | \zeta_k \in [0, 1], \sum_k \zeta_k = 1\}$ are the weights of each representation at word step i .

There are multiple choices for the design of $\zeta_{k,i}$ including simple softmax, Gumbel softmax [31], straight through (ST) Gumbel softmax [31], those with temperature parameters, and weighted variants of those designs to incorporate some preference on a certain type of representation. Here we define the general form of these variants as follows, and give the ablation study in Section 5.3.

$$\zeta_{k,i} = \frac{\exp((g_{k,i} + \log \pi_{k,i})/\tau)}{\sum_{k'} \exp((g_{k',i} + \log \pi_{k',i})/\tau)}, \quad (5)$$

where $g_{k,i} = 0$ in softmax, and $g_{k,i} = -\log(-\log u)$, $u \sim U(0, 1)$ in (ST) Gumbel softmax. $U(0, 1)$ is the uniform distribution. When we use ST Gumbel softmax, we can transform $\zeta_{k,i}$ into a one-hot vector in the forward pass by $\zeta'_{k,i} = \text{one-hot} \left(\arg \max_{k'} \zeta_{k',i} \right)$, and approximate $\zeta'_{k,i}$ by $\zeta_{k,i}$ in the backward pass to enable back-propagation. τ is the temperature to control the “hardness” of the attention. A smaller τ tends to give “harder” attention, which results in making an output vector close to a one-hot one. π_k are defined as follows.

$$\pi_{k,i} = \frac{c_k \rho_{k,i}}{\sum_{k'} c_{k'} \rho_{k',i}}, \quad (6)$$

where c_k are the weights that represent the preference for each representation. $\rho_{k,i}$ are defined as follows.

$$\rho_{k,i} = \frac{\eta_{k,i}}{\sum_{k'} \eta_{k',i}}, \quad (7)$$

where $\eta_{k,i} = \text{sigmoid}(\mathbf{W}'_{k,i} \mathbf{h}_k + \mathbf{b}_{k,i})$. Note that we calculate the attention $\zeta_{k,i}$ at each word step i , which means that DMA can dynamically and flexibly change the attention at each word step in a sentence depending on the context. See Section 5.3 for the ablation study on the different choices of $g_{k,i}$, τ , and c_k .

5 EXPERIMENT

5.1 Setup

We evaluated our approach on the MMAC Captions dataset. We split the dataset into training, validation, and test sets so that each

set includes all recipe types, resulting in 2,923, 838, and 1,241 data for the training, validation, and test sets, respectively.

Implementation Details. We set the maximum step for video frames (K for ϕ_V) and IMU signals (T for ϕ_S) to be 80 and 240, respectively. If the number of video frames did not match K , we down-sampled or duplicated the video frame. We also down-sampled IMU signals to match T .

We used the VGG16 model pre-trained with ImageNet for ϕ_V . The representations in the encoder had the following dimensions: $\mathbf{h}_V \in \mathbb{R}^{500}$, $\mathbf{h}_S \in \mathbb{R}^{120}$, and $\mathbf{h}_{V+S} \in \mathbb{R}^{620}$. Training was conducted by minimizing the cross-entropy loss with the Adam optimizer. We set the batch size to 100 sequences and the learning rates to 3×10^{-4} ($0 \leq \text{epoch} < 300$), 1×10^{-4} ($\text{epoch} < 400$), and 5×10^{-5} ($400 \leq \text{epoch}$). The maximum number of words in the generated sentence, *i.e.*, N in the decoder, were set to 15. We used Gumbel softmax for Eq. (5), and set $\tau = 0.05$ and $(c_V, c_S, c_{V+S}) = (1, 1, 1.5)$ unless otherwise stated. During training, we used scheduled sampling [5], which improves generalization performance. PyTorch was used for the implementation.

Metrics and Baselines. We used four metrics for the evaluation: BLEU, METEOR, CIDEr-D, and SPICE. BLEU is a precision metric of word n-grams between prediction and ground-truth sentence. METEOR is a more semantic metric that absorbs subtle differences in expression using a WordNet synonym. CIDEr-D measures the cosine similarity between a generated sentence and ground-truth sentence using term-frequency inverse-document-frequency and avoids the gaming effect with a stemming technique. SPICE is a metric that has shown correlations with human judgement. As in most previous studies, we used METEOR and CIDEr-D as the main metrics. They are computed using the MS-COCO caption evaluation tool [10]. We used average metrics of 6-times training to avoid the effect of random initialization.

We chose reasonable, versatile baselines to verify the performance of the novel multi-modal captioning task:

- (1) Egocentric-based method: ABiViNet [7]
- (2) RNN-based method: LSTM-YT [78], S2VT [77], and BA [4].
- (3) Transformer-based method: VTransformer [89].

5.2 Results

Quantitative Results. Table 1 lists the results on the MMAC Captions dataset. The proposed approach, which dynamically integrates vision and sensor data (Ours (V+S dynamic)) outperformed all the baselines. It achieved +1.4 gain in METEOR compared with the vision-only result (Ours (V)). As METEOR is close to a subjective evaluation, the results justify our claim that sensor-augmented multi-modal captioning is effective for generating detailed activity descriptions. Our approach when using only the sensors (Ours (S)) was inferior to the other baselines, but interestingly, it had comparable performance to LSTM-YT. This result also supports the effectiveness of sensor information. Ours (V+S fixed) is the result obtained by using $\mathbf{h}_{V+S}$ all the time without using AMMT and DMA. Although it outperformed both of the single-modality variants (Ours (V) and Ours (S)) as well as achieved comparable performance to the best baseline model (VTransformer), our proposed approach with AMMT and DMA (Ours (V+S dynamic)) acquired

Table 1: Captioning performance on MMAC Captions dataset. V indicates video data and S are sensor data. B-N represents BLEU score at N-gram. We computed the averaged value of each metric with six training runs.

Method	B-1	B-2	B-3	B-4	B-5	METEOR	CIDEr-D	SPICE
LSTM-YT [78]	52.9	45.8	39.8	34.6	30.7	24.8	292.9	0.423
S2VT [77]	64.5	58.0	52.7	48.5	45.6	33.8	428.9	0.557
ABiViRNet [7]	62.4	55.6	50.3	45.9	42.8	32.4	419.4	0.557
BA [4]	66.8	60.5	55.6	51.6	48.8	35.8	460.3	0.597
VTransformer [89]	67.9	61.8	56.7	52.5	49.6	36.5	471.4	0.625
Ours (V)	66.8	60.5	55.6	51.6	48.8	35.8	460.3	0.597
Ours (S)	49.5	41.2	35.3	30.9	27.8	23.0	278.7	0.384
Ours (V+S fixed)	67.6	61.5	56.6	52.6	49.8	36.5	476.9	0.613
Ours (V+S dynamic)	68.4	62.3	57.4	53.5	50.7	37.2	484.8	0.618

an additional gain of 0.7 in METEOR. This result indicates that it is important to dynamically change the attention to the different modalities, and DMA effectively worked from this perspective.

Qualitative Results. Qualitative results are shown in Figure 4. In each example, the top row is a snapshot of video frames. The middle row represents 21-dimensional gyro signals (3-axes by seven locations) arranged in chronological order. The bottom row shows the captioning results. We observed that the first words, (*i.e.*, verbs), are always generated by attending to a multi-modal representation $\mathbf{h}_{V+S}$, confirming that sensor-augmented representation is especially effective for generating verbs.

Figures 4 (a)–(b) indicate that sensor fusion helps infer the appropriate verbs and sometimes even nouns. While it is difficult for vision-only methods to distinguish visually similar actions such as *Opening* and *Taking* in these examples, the auxiliary sensor information was able to help generate the precise descriptions. Figures 4 (c)–(d) show the effectiveness of our DMA module. The colors of the generated words indicate that DMA appropriately shifts the attention from V+S to V and vice versa when necessary, resulting in better captioning results. As seen in these examples, it sometimes attends to V rather than V+S when generating nouns. We believe this is reasonable because visual information should be especially important for nouns and noisy sensor information sometimes harms the performance, although it is generally helpful for generating verbs. We will further study these points in the final paragraph of this section, referring to Figure 6. Figure 4 (e) represents a failure case where visual information alone is sufficient probably because the object is clear. In this case, sensor information caused some inaccurate inference (although it seems somewhat reasonable), but DMA failed to attend to V. Figure 4 (f) is a challenging case, where an activity took place in the out-of-camera-range. In this case, all the methods including the proposed method failed to accurately generate the description. Generating an accurate description even in such a case would require the incorporation of much broader contexts in the long term, which will be our future work, potentially by exploiting transformer-like architecture.

5.3 Further Analysis

Ablation Study. We conducted an ablation study to investigate the characteristics of the proposed sensor-fusion method. In

particular, we tested the following combinations: (i) *vision*: using only visual information, (ii) *vision+sensor*: multi-modal fusion using simple concatenation, (iii) *asymmetric fusion*: using AMMT for multi-modal fusion, (iv) *dynamic attention*: using DMA for sentence generation, and (v) *full*: using all the techniques.

As shown in Table 2, *vision+sensor* outperformed *vision*, which suggests the effectiveness of the sensor modality as auxiliary information. We found fusion with only DMA without AMMT (iv) were worse than fusion with simple concatenation (ii). We assume this is because it is important to adjust the different representations, *i.e.*, $\mathbf{h}_V$, $\mathbf{h}_S$, and $\mathbf{h}_{V+S}$, so that a possible imbalance among these representations can be calibrated before being input to DMA. Similarly, fusion with only AMMT (iii) resulted in a worse performance than fusion with simple concatenation (ii) probably because the transformation layer is unnecessary if DMA is not used, and the unnecessary redundant layer caused over-fitting. Finally, *full*, which uses both of AMMT and DMA, led to the best results.

Analysis of AMMT. We analyzed the design of AMMT in more detail. We conducted the ablation study by changing the symmetric and asymmetric patterns to see how the fusion module achieves the gain. The results presented in Table 3 confirm the effectiveness of the asymmetric transformation. We assume this is because the possible imbalance due to the difference between visual and sensor representation is mitigated by applying a transformation. Introducing symmetric modulation, or applying transformations to both of the representations, decreased the performance probably because of the redundant parameters. Incorporating redundant parameters, despite being sufficient to apply a transformation to either of the representations, may lead to over-fitting. In contrast, both the asymmetric modulations improve the performance, especially when the transformation is applied to the sensor representations. One possible hypothesis is that the sensor data sometimes contain undesirable noise, which can be surpassed by applying a transformation.

Analysis of DMA. As mentioned in Section 4.2, we explore the variants of $g_{k,i}$; softmax with temperature [27] and (ST) Gumbel softmax [31], namely, $g_{k,i} = 0$ in softmax, and $g_{k,i} = -\log(-\log u)$, $u \sim U(0, 1)$ in (ST) Gumbel softmax in Eq. (5).

Table 4 shows the results. We found that Gumbel softmax worked best. We assume this is because the stochastic sampling had a similar

Figure 4: Qualitative results. Notations: GT-Ground Truth, V-Vision, S-Sensor. Orange words represent those generated by attending multi-modal representation h_{V+S} , while blue ones represent those generated by h_V . ©Ekaterrina H. Spriggs.

Table 2: Results of ablation study for multi-modal fusion.

No	Fusion	AMMT	DMA	B-1	B-2	B-3	B-4	B-5	METEOR	CIDEr-D	SPICE
(i)				66.8	60.5	55.6	51.6	48.8	35.8	460.3	0.597
(ii)	✓			67.6	61.5	56.6	52.6	49.8	36.5	476.9	0.613
(iii)	✓	✓		67.5	61.3	56.4	52.5	49.7	36.4	474.0	0.607
(iv)	✓		✓	67.0	60.8	55.9	52.0	49.2	36.1	467.7	0.602
(v)	✓	✓	✓	68.4	62.3	57.4	53.5	50.7	37.2	484.8	0.618

effect as ensembling and led to a slightly better performance. ST Gumbel softmax was found to be worse than the vanilla Gumbel softmax, which indicates that it is better to use a soft attention rather than a hard one. However, ST Gumbel softmax would be a good choice if it is desirable to completely switch off certain modalities, e.g., for reducing sensor, computation, or memory cost.

Figure 5 (a) shows the sensitivity of hyper-parameter τ . We found that introducing a slightly hard attention by $\tau = 0.05$ gave the best

performance, whereas a too hard attention, namely $\tau = 0.001$, degraded the performance. This result agrees with the aforementioned observation on Gumbel and ST Gumbel softmax.

Figure 5 (b) shows the ablation for hyper-parameter c_k ; a modality preference. We set $c_V = c_S = 1.0$. The best performance was obtained on $c_{V+S} = 1.5$. The result suggests that it is better to slightly emphasize V+S over single modalities, which indicates the effectiveness of multi-modal information.

Table 3: Ablation study for AMMT. Linear (h_V/h_S) denotes the modality that linear transformation was applied in AMMT. The last row corresponds to Eq. (1). B: BLEU, M: METEOR, C: CIDEr-D, S: SPICE.

Linear(h_V)	Linear(h_S)	B-1	B-4	M	C	S
		67.0	52.0	36.1	467.7	0.602
✓	✓	67.0	51.7	35.9	465.0	0.602
✓		67.4	52.4	36.2	468.4	0.604
	✓	68.4	53.5	37.2	484.8	0.618

Table 4: A comparison on weighted representation for DMA.

Method	B-1	B-4	M	C	S
Softmax w/ temperature	67.9	52.8	36.5	474.8	0.612
ST Gumbel softmax	67.8	52.8	36.7	473.2	0.612
Gumbel softmax	68.4	53.5	37.2	484.8	0.618

Figure 5: Sensitivity analysis of hyper-parameters.

Detailed Analysis on Word-types. Finally, we analyze the relationship between the generated word-types and modal attention. Figure 6 (a) shows which modalities were attended to most, or $\text{argmax}_k \zeta_{k,i}$, by word-types. We can see that h_{V+S} was attended most in most cases. In particular, h_{V+S} was utilized for generating verbs in 99.6% of all the cases, suggesting the effectiveness of multi-modal information for describing motions. In addition, we confirmed that the visual representations h_V were emphasized more often when generating nouns, prepositions, and determiners, with rates ranging from 26.2% to 35.3%. This indicates that sensor data are not always useful and it is better to focus only on visual representation if sensor data do not contain useful information or contain harmful information such as noise or signals of irrelevant motions. Figure 6 (b) shows a scatter plot of the average attention rate of the modalities with respect to each noun words. The upper left region means that h_{V+S} was almost always attended to most

Figure 6: Further analysis on DMA: (a) Number of word-types for V and V+S, (b) Scatter plot on attending rate for nouns.

for generating the nouns in this region, and the lower right region means h_V was used more often to generate the nouns in this region. We found that easily identifiable objects with only vision data, e.g., *plate*, tended to be in the lower right region, while confusing objects, e.g., *ingredients*, tended to be in the upper left. The result indicates that DMA worked reasonably from a qualitative perspective as well.

6 CONCLUSION

This paper proposed a novel task of egocentric multi-modal captioning, which incorporates wearable sensors to generate detailed activity description. To address this task, we constructed a dataset called MMAC Captions. We proposed the model for sensor-augmented video captioning, which has two key modules; AMMT and DMA. Experimental results demonstrated that AMMT enabled to effectively integrate multi-modal representations, DMA gave more precise captioning by appropriately attending to the preferable modalities, and the combination of these two resulted in consistently superior performance compared with the baselines. We believe this study will facilitate the further research and open up an opportunity for developing new applications such as automatic generation of detailed industrial-operation-manuals.

ACKNOWLEDGMENTS

The CMU-MMAC data used in this paper was obtained from kitchen.cs.cmu.edu and the data collection was funded in part by the National Science Foundation under Grant No. EEE-0540865.

REFERENCES

- [1] Nayyer Afaaq, Naveed Akhtar, Wei Liu, Syed Zulqarnain Gilani, and Ajmal Mian. 2019. Spatio-temporal dynamics and semantic attribute enriched visual encoding for video captioning. In *Computer Vision and Pattern Recognition (CVPR)*.
- [2] Mohammad Al-Naser, Hiroki Ohashi, Sheraz Ahmed, Katsuyuki Nakamura, Takayuki Akiyama, Takuto Sato, Phong Xuan Nguyen, and Andreas Dengel. 2018. Hierarchical Model for Zero-shot Activity Recognition using Wearable Sensors. In *International Conference on Agents and Artificial Intelligence (ICAART)*.
- [3] Ling Bao and Stephen S. Intille. 2004. Activity recognition from user-annotated acceleration data. In *Pervasive Computing*.
- [4] Lorenzo Baraldi, Costantino Grana, and Rita Cucchiara. 2017. Hierarchical Boundary-Aware Neural Encoder for Video Captioning. In *Computer Vision and Pattern Recognition (CVPR)*.
- [5] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, Noam Shazeer, and Mountain View. 2015. Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks. In *Neural Information Processing Systems (NeurIPS)*.
- [6] Alejandro Betancourt, Pietro Morerio, Carlo S. Regazzoni, and Matthias Rauterberg. 2015. The evolution of first person vision methods: A survey. *IEEE Transactions on Circuits and Systems for Video Technology* 25, 5 (2015), 744–760.
- [7] Marc Bolaños, Álvaro Peris, Francisco Casacuberta, Sergi Soler, and Petia Radeva. 2018. Egocentric video description based on temporally-linked sequences. *Journal of Visual Communication and Image Representation* 50 (2018), 205–216.
- [8] Ming Chen, Yingming Li, Zhongfei Zhang, and Siyu Huang. 2018. TVT: Two-View Transformer Network for Video Captioning. In *Asian Conference on Machine Learning (ACML)*.
- [9] Shaoxiang Chen, Ting Yao, and Yu Gang Jiang. 2019. Deep learning for video captioning: A review. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- [10] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, and C. Lawrence Zitnick. 2015. Microsoft COCO Captions: Data Collection and Evaluation Server. (2015).
- [11] Xinpeng Chen, Lin Ma, Wenhao Jiang, Jian Yao, and Wei Liu. 2018. Regularizing RNNs for caption generation by reconstructing the past with the present. In *Computer Vision and Pattern Recognition (CVPR)*.
- [12] Xinlei Chen and C Lawrence Zitnick. 2015. Mind's Eye: A Recurrent Visual Representation for Image Caption Generation. In *Computer Vision and Pattern Recognition (CVPR)*.
- [13] Yangyu Chen, Shuhui Wang, Weigang Zhang, and Qingming Huang. 2018. Less Is More: Picking Informative Frames for Video Captioning. In *European Conference on Computer Vision (ECCV)*.
- [14] Dima Damen and Andrew Calway. 2014. You-do, I-learn : discovering task relevant objects and their modes of interaction from multi-user egocentric video. In *British Machine Vision Conference (BMVC)*.
- [15] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. 2018. Scaling Egocentric Vision: The EPIC-KITCHENS Dataset. In *European Conference on Computer Vision (ECCV)*.
- [16] Jacob Devlin, Ming Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- [17] Jeff Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *Computer Vision and Pattern Recognition (CVPR)*.
- [18] Chenyou Fan and David J. Crandall. 2016. DeepDiary: Automatic Caption Generation for Lifelogging Image Streams. In *ECCV Workshop*.
- [19] Alireza Fathi, Ali Farhadi, and James M. Rehg. 2011. Understanding egocentric activities. In *International Conference on Computer Vision (ICCV)*.
- [20] Alireza Fathi, Jessica K Hodgins, and James M Rehg. 2012. Social interactions: a first-person perspective. In *Computer Vision and Pattern Recognition (CVPR)*.
- [21] Alireza Fathi and James M Rehg. 2013. Modeling actions through state changes. In *Computer Vision and Pattern Recognition (CVPR)*.
- [22] Alireza Fathi, Xiaofeng Ren, and James M. Rehg. 2011. Learning to recognize objects in egocentric activities. In *Computer Vision and Pattern Recognition (CVPR)*.
- [23] Simon Ging, Mohammadreza Zolfaghari, Hamed Pirsiavash, and Thomas Brox. 2020. COOT: Cooperative hierarchical transformer for video-text representation learning. In *Neural Information Processing Systems (NeurIPS)*.
- [24] Sergio Guadarrama, Niveda Krishnamoorthy, Girish Malkarnenkar, Subhashini Venugopalan, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2013. Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition. In *International Conference on Computer Vision (ICCV)*.
- [25] Pirsiavash Hamed and Deva Ramanan. 2012. Detecting activities of daily living in first-person camera views. In *Computer Vision and Pattern Recognition (CVPR)*.
- [26] Lisa Anne Hendricks, Subhashini Venugopalan, Marcus Rohrbach, Raymond Mooney, Kate Saenko, and Trevor Darrell. 2016. Deep Compositional Captioning: Describing Novel Object Categories without Paired Training Data. In *Computer Vision and Pattern Recognition (CVPR)*.
- [27] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. In *NIPS Workshop*.
- [28] Chiori Hori, Takaaki Hori, Teng-Yok Lee, Kazuhiro Sumi, John R. Hershey, and Tim K. Marks. 2017. Attention-Based Multimodal Fusion for Video Description. In *International Conference on Computer Vision (ICCV)*.
- [29] Vladimir Iashin and Esa Rahtu. 2020. A Better Use of Audio-Visual Cues: Dense Video Captioning with Bi-modal Transformer. In *British Machine Vision Conference (BMVC)*.
- [30] Vladimir Iashin and Esa Rahtu. 2020. Multi-modal dense video captioning. In *CVPR Workshop on Multimodal Learning (CVPRW)*.
- [31] Eric Jang, Shixiang Gu, Ben Poole, Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2017. Categorical Reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations (ICLR)*.
- [32] Tao Jin, Siyu Huang, Ming Chen, Yingming Li, and Zhongfei Zhang. 2020. SBAT: Video captioning with sparse boundary-aware transformer. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- [33] Andrej Karpathy and G Toderici. 2014. Large-scale video classification with convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR)*.
- [34] Lei Ke, Wenjie Pei, Ruiyu Li, Xiaoyong Shen, and Yu Wing Tai. 2019. Reflective decoding network for image captioning. In *Proceedings of the IEEE International Conference on Computer Vision*.
- [35] Kris M. Kitani, Takahiro Okabe, Yoichi Sato, and Akihiro Sugimoto. 2011. Fast unsupervised ego-action learning for first-person sports videos. In *Computer Vision and Pattern Recognition (CVPR)*.
- [36] Quan Kong, Ziming Wu, Ziwei Deng, Martin Klunkigt, Bin Tong, and Tomokazu Murakami. 2019. MMAct: A Large-Scale Dataset for Cross Modal Human Action Understanding. In *International Conference on Computer Vision (ICCV)*.
- [37] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Nieves. 2017. Dense-Captioning Events in Videos. In *International Conference on Computer Vision (ICCV)*.
- [38] Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C. Berg, and Tamara L. Berg. 2013. Baby talk: Understanding and generating simple image descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 35, 12 (2013), 2891–2903.
- [39] Polina Kuznetsova, Vicente Ordonez, Alexander C Berg, Tamara L. Berg, Yejin Choi, and Stony Brook. 2012. Collective Generation of Natural Image Descriptions. In *Association for Computational Linguistics (ACL)*.
- [40] Rémi Lebre, Pedro O. Pinheiro, and Roman Collobert. 2015. Phrase-based Image Captioning. In *International Conference on Machine Learning (ICML)*.
- [41] Yong Jae Lee, Joydeep Ghosh, and Kristen Grauman. 2012. Discovering important people and objects for egocentric video summarization. In *Computer Vision and Pattern Recognition (CVPR)*.
- [42] Jie Lei, Liwei Wang, Yelong Shen, Dong Yu, Tamara L. Berg, and Mohit Bansal. 2020. MART: Memory-augmented recurrent transformer for coherent video paragraph captioning. In *Association for Computational Linguistics (ACL)*.
- [43] Yin Li, Alireza Fathi, and James M Rehg. 2013. Learning to predict gaze in egocentric video. In *Computer Vision and Pattern Recognition (CVPR)*.
- [44] Yin Li, Zhefan Ye, and James M. Rehg. 2015. Delving into egocentric actions. In *Computer Vision and Pattern Recognition (CVPR)*.
- [45] Chang Liu, Changhu Wang, Fuchun Sun, and Yong Rui. 2016. Image2Text: A Multimodal Caption Generator Chang. In *ACM international conference on Multimedia (ACM MM)*.
- [46] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. 2017. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Computer Vision and Pattern Recognition (CVPR)*.
- [47] Minghuang Ma, Haoqi Fan, and Kris M Kitani. 2016. Going deeper into first-person activity recognition. In *Computer Vision and Pattern Recognition (CVPR)*.
- [48] Takuya Maekawa, Yutaka Yanagisawa, Yasue Kishino, Katsuhiko Ishiguro, Koji Kamei, Yasushi Sakurai, and Takeshi Okadome. 2010. Object-based activity recognition with heterogeneous sensors on wrist. In *Pervasive*.
- [49] Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan Yuille. 2015. Deep Captioning with Multimodal Recurrent Neural Networks (m-RNN). In *International Conference on Learning Representations (ICLR)*.
- [50] Tushar Nagarajan, Yanghao Li, Christoph Feichtenhofer, and Kristen Grauman. 2020. EGO-TOPO: Environment Affordances from Egocentric Video. In *Computer Vision and Pattern Recognition (CVPR)*.
- [51] Katsuyuki Nakamura, Serena Yeung, Alexandre Alahi, and Li Fei-Fei. 2017. Jointly Learning Energy Expenditures and Activities using Egocentric Multimodal Signals. In *Computer Vision and Pattern Recognition (CVPR)*.
- [52] Motofumi Nakanishi, Shintaro Izumi, Sho Nagayoshi, Hironori Sato, Hiroshi Kawaguchi, Masahiko Yoshimoto, Takafumi Ando, Satoshi Nakae, Chiyoko Usui, Tomoko Aoyama, and Shigeo Tanaka. 2015. Physical activity group classification algorithm using triaxial acceleration and heart rate. In *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*.
- [53] Evonne Ng, Donglai Xiang, Hanbyul Joo, and Kristen Grauman. 2020. You2Me: Inferring Body Pose in Egocentric Video via First and Second Person Interactions. In *Computer Vision and Pattern Recognition (CVPR)*.

- [54] Thi Hoa Cuc Nguyen, Jean Christophe Nebel, and Francisco Florez-Revuelta. 2016. Recognition of activities of daily living with egocentric vision: A review. *Sensors* 16, 1 (2016).
- [55] Hiroki Ohashi, Mohammad Al-Naser, Sheraz Ahmed, Katsuyuki Nakamura, Takuto Sato, and Andreas Dengel. 2018. Attributes' importance for zero-shot pose-classification based on wearable sensors. *Sensors* 18, 8 (2018).
- [56] Hiroki Ohashi, M Al-Nasser, Sheraz Ahmed, Takayuki Akiyama, Takuto Sato, Phong Nguyen, Katsuyuki Nakamura, and Andreas Dengel. 2017. Augmenting wearable sensor data with physical constraint for DNN-based human-action recognition. In *ICML Time Series Workshop*.
- [57] Katsunori Ohnishi, Atsushi Kanehira, Asako Kanazaki, and Tatsuya Harada. 2015. Recognizing activities of daily living with a wrist-mounted camera. In *Computer Vision and Pattern Recognition (CVPR)*.
- [58] Vicente Ordonez, Girish Kulkarni, and Tamara L Berg. 2011. Im2Text: Describing Images Using 1 Million Captioned Photographs. In *Neural Information Processing Systems (NeurIPS)*.
- [59] Boxiao Pan, Haoye Cai, De-An Huang, Kuan-Hui Lee, Adrien Gaidon, Ehsan Adeli, and Juan Carlos Niebles. 2020. Spatio-Temporal Graph for Video Captioning with Knowledge Distillation. In *Computer Vision and Pattern Recognition (CVPR)*.
- [60] Yingwei Pan, Tao Mei, Ting Yao, Houqiang Li, and Yong Rui. 2016. Jointly Modeling Embedding and Translation to Bridge Video and Language. In *Computer Vision and Pattern Recognition (CVPR)*.
- [61] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. FiLM: Visual reasoning with a general conditioning layer. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [62] Yair Poleg, Chetan Arora, Tavi Halperin, Chetan Arora, Shmuel Peleg, Tavi Halperin, Chetan Arora, and Shmuel Peleg. 2015. EgoSampling: Fast-Forward and Stereo for Egocentric Videos. In *Computer Vision and Pattern Recognition (CVPR)*.
- [63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin. 2017. Attention Is All You Need. In *Neural Information Processing Systems (NeurIPS)*.
- [64] Tanzila Rahman, Bicheng Xu, and Leonid Sigal. 2019. Watch, listen and tell: Multimodal weakly supervised dense event captioning. In *International Conference on Computer Vision (ICCV)*.
- [65] Vasili Ramanishka, Abir Das, Dong Huk Park, Subhashini Venugopalan, Lisa Anne Hendricks, Marcus Rohrbach, and Kate Saenko. 2016. Multimodal Video Description. In *ACM international conference on Multimedia (ACM MM)*.
- [66] Marcus Rohrbach, Stefan Thater, Wei Qiu, Manfred Pinkal, Ivan Titov, and Max Planck. 2013. Translating Video Content to Natural Language Descriptions. In *International Conference on Computer Vision (ICCV)*.
- [67] M S Ryo, Brandon Rothrock, and Larry Matthies. 2015. Pooled motion features for first-person videos. In *Computer Vision and Pattern Recognition (CVPR)*.
- [68] Dandan Shan, Jiaqi Geng, Michelle Shu, and David F. Fouhey. 2020. Understanding Human Hands in Contact at Internet Scale. In *Computer Vision and Pattern Recognition (CVPR)*.
- [69] Suriya Singh, Chetan Arora, and C. V. Jawahar. 2017. Trajectory aligned features for first person action recognition. *Pattern Recognition* 62 (2017), 45–55.
- [70] Ekaterina H Spriggs, Fernando De La Torre, and Martial Hebert. 2009. Temporal segmentation and activity classification from first-person sensing. In *CVPR Workshop on Egocentric Vision (CVPRW)*.
- [71] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. VideoBERT: A joint model for video and language representation learning. In *International Conference on Computer Vision (ICCV)*.
- [72] George Toderici, Hrishikesh Aradhye, Marius Pasca, Luciano Sbaiz, and Jay Yagnik. 2010. Finding Meaning on YouTube: Tag Recommendation and Category Discovery. In *Computer Vision and Pattern Recognition (CVPR)*.
- [73] Bin Tong, Martin Klinkigt, Makoto Iwayama, Toshihiko Yanase, Yoshiyuki Kobayashi, Anshuman Sahu, and Ravigopal Vennelakanti. 2017. Learning to Generate Rock Descriptions from Multivariate Well Logs with Hierarchical Attention. In *International Conference on Knowledge Discovery and Data Mining (KDD)*.
- [74] Yoshitaka Ushiku, Tatsuya Harada, and Yasuo Kuniyoshi. 2011. Understanding Images with Natural Sentences. In *ACM international conference on Multimedia (ACM MM)*.
- [75] Yoshitaka Ushiku, Masataka Yamaguchi, Yusuke Mukuta, and Tatsuya Harada. 2015. Common subspace for model and similarity: Phrase learning for caption generation from images. In *International Conference on Computer Vision (ICCV)*.
- [76] Subhashini Venugopalan, Lisa Anne Hendricks, Raymond Mooney, and Kate Saenko. 2016. Improving LSTM-based Video Description with Linguistic Knowledge Mined from Text. In *Empirical Methods on Natural Language Processing (EMNLP)*.
- [77] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to sequence - Video to text. In *International Conference on Computer Vision (ICCV)*.
- [78] Subhashini Venugopalan, Huijuan Xu, Jeff Donahue, Marcus Rohrbach, Raymond Mooney, Kate Saenko, U C Berkeley, Marcus Rohrbach, U C Berkeley, Raymond Mooney, and Kate Saenko. 2015. Translating Videos to Natural Language Using Deep Recurrent Neural Networks. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- [79] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Computer Vision and Pattern Recognition (CVPR)*.
- [80] Xiaohan Wang, Yu Wu, Linchao Zhu, and Yi Yang. 2020. Symbiotic Attention with Privileged Information for Egocentric Action Recognition. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [81] Zuxuan Wu, Yanwei Fu, Yu-Gang Jiang, and Leonid Sigal. 2016. Harnessing Object and Scene Semantics for Large-Scale Video Understanding. In *Computer Vision and Pattern Recognition (CVPR)*.
- [82] Jun Xu, Ting Yao, Yongdong Zhang, and Tao Mei. 2017. Learning Multimodal Attention LSTM Networks for Video Captioning. In *ACM Multimedia (ACM MM)*.
- [83] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard Zemel, and Yoshua Bengio. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *International Conference on Machine Learning (ICML)*.
- [84] Li Yao, Atousa Torabi, Kyunghyun Cho, Nicolas Ballas, Christopher Pal, Hugo Larochelle, and Aaron Courville. 2015. Describing videos by exploiting temporal structure. In *International Conference on Computer Vision (ICCV)*.
- [85] Ryo Yonetani and Kris M Kitani. 2016. Recognizing micro-actions and reactions from paired egocentric videos. In *Computer Vision and Pattern Recognition (CVPR)*.
- [86] Ye Yuan and Kris Kitani. 2019. Ego-Pose Estimation and Forecasting as Real-Time PD Control. In *International Conference on Computer Vision (ICCV)*.
- [87] Ziqi Zhang, Yaya Shi, Chunfeng Yuan, Bing Li, Peijin Wang, Weiming Hu, and Zhengjun Zha. 2020. Object Relational Graph with Teacher-Recommended Learning for Video Captioning. In *Computer Vision and Pattern Recognition (CVPR)*.
- [88] Luowei Zhou, Yannis Kalantidis, Xinlei Chen, Jason J. Corso, and Marcus Rohrbach. 2019. Grounded video description. In *Computer Vision and Pattern Recognition (CVPR)*.
- [89] Luowei Zhou, Yingbo Zhou, Jason J. Corso, Richard Socher, and Caiming Xiong. 2018. End-to-End Dense Video Captioning with Masked Transformer. In *Computer Vision and Pattern Recognition (CVPR)*.