

A Space Information-Enhanced Dense Video Caption for Indoor Human Action Recognition

Bin CHEN
ISEE, Kyushu University
Kyushu University
Fukuoka, Japan
bin.367@s.kyushu-u.ac.jp
*Corresponding author

Yugo NAKAMURA
ISEE, Kyushu University
Kyushu University
Fukuoka, Japan
y-nakamura@ait.kyushu-u.ac.jp

Shogo FUKUSHIMA
ISEE, Kyushu University
Kyushu University
Fukuoka, Japan
shogo@ait.kyushu-u.ac.jp

Yutaka ARAKAWA
ISEE, Kyushu University
Kyushu University
Fukuoka, Japan
arakawa@ait.kyushu-u.ac.jp

Abstract—Dense video captioning tasks are used to detect interesting events and provide descriptive text for these events from untrimmed videos. This technology has the potential to be used in security surveillance and human care applications. However, current methods often overlook the relationships between objects in the video, which limits their applicability and makes it challenging to adapt them to specific domains, such as video summarization for indoor human activities. In these scenarios, human activities are closely intertwined with the objects in the scene. In this paper, we propose a plug-and-play module designed to enhance existing dense video captioning methods with spatial information. Specifically, we extract spatial information about the interesting objects using Red-Green-Blue-depth (RGB-D) images and the results of image segmentation. We then integrate this information into the captions generated by the Dense Video Captioning (DVC) method using a fine-tuned Large Language Model (LLM).

We evaluate the performance of our model on a custom dataset and demonstrate that our system provides a convenient and effective approach for obtaining space-enhanced captions.

Index Terms—Dense video caption, Image segmentation, Depth camera, Human action recognition

I. INTRODUCTION

Video is an important information medium in security monitoring and smart home applications. The large volume of video content demands the need for automated methods to summarize and compactly represent the essential content [1]. One promising approach to creating content summaries is to use dense video captions, a technique that generates descriptive text for each frame of video [3].

However, existing dense video captioning methods and related datasets have often overlooked spatial information and object relationships within the visual contents of videos. This limitation hinders the adaptability of these methods to other application scenarios, such as generating activity log files for indoor individuals or indoor security monitoring, where the same activity can have varying interpretations with different indoor positions. Intuitively, to clearly and unambiguously describe indoor human behavior, captions should include information about the action class, location, and associated objects.

In this paper, we introduce a plug-and-play module to enhance spatial information in captions, building upon the current popular approach. As illustrated in Fig. 1, we use

Red-Green-Blue-Depth (RGB-D) data as input and directly extract accurate spatial information from the RGB-D images. Subsequently, we combine this spatial information with object relationships using a fine-tuned sequence-to-sequence (Seq2Seq) text generation model to generate the unambiguous and detailed captions. In our scheme, the Region of Interest (RoI) to which we pay particular attention include: 1. The human’s position in relation to the main object. 2. The position of the object interacting with the human. 3. The ambiguous object in relation to other objects of the same class. We collected a dense video caption dataset enhanced with spatial information and object relationships of interesting objects and then evaluated our system on the dataset to show its effectiveness. Our results demonstrate that the system is capable of effectively generating dense video captions with space information.

Our contribution is here:

- 1) We collected a dataset for dense video captioning, enriched with spatial information and the relationships among interesting objects.
- 2) We designed a plug-and-play module for existing methods to generate space-enhanced unambiguous captions.

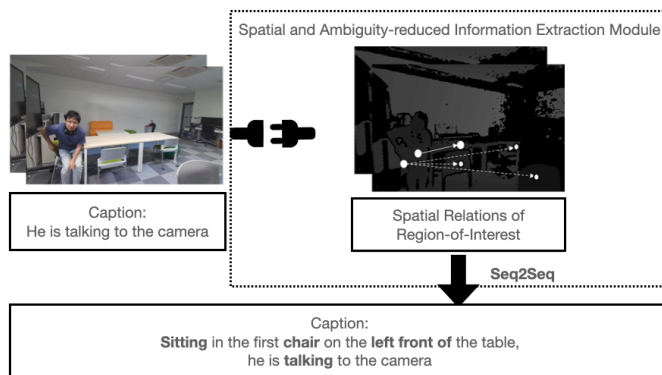


Fig. 1. Main scheme of the system

II. RELATED WORK

A. Dense video captioning and Dataset

Dense video captioning is an emerging sub-field in video understanding, with a focus on event localization and event captioning within long, untrimmed videos that often contain a mix of events and irrelevant background content. Compared to traditional video captioning methods, dense video captioning aims to generate descriptions for all events within a video, making it suitable for application scenarios that require detailed textual descriptions.

Based on the input, there are two main categories of methods in this field: 1. Multi-modality-based methods and 2. Vision-modality-based methods. For example, Iashin et al. utilize audio, speech, and visual information as input and merge the extracted features with a multi-modal generator to output a distribution over the vocabulary [1]. Yang et al. propose vid2seq [2], which uses video and transcribed speech as input to achieve state-of-the-art results in the dense video captioning task. However, in the real world, audio or speech inputs often come with significant background noise, which can potentially impact the results of an application. Fortunately, there are also models based solely on vision that have achieved state-of-the-art results.

For instance, Wang et al. introduced a framework for end-to-end dense video captioning with parallel decoding (PDVC) using only visual input. They formulated dense caption generation as a set prediction task. In their pipeline, they employed CNN and Transformer encoders to extract features. In the parallel decoding phase, they implement a Transformer decoder with three heads (event counter, localization header, and caption header) to generate the localization and caption independently.

However, current vision-based end-to-end methods often neglect the spatial information and fail to address ambiguity. Furthermore, the existing datasets used in dense video captioning do not adequately incorporate this information, despite the presence of some spatial details within the dataset's ground truth sentences. The current popular datasets used in dense video captioning include YouCook2 and ActivityNet. YouCook2 is one of the largest task-oriented instructional video datasets, containing 2000 long untrimmed videos from 89 cooking recipes. On average, each distinct recipe has 22 videos [5]. The ActivityNet Captions dataset is based on ActivityNet v1.3 and includes 20,000 YouTube untrimmed videos with 100,000 caption annotations. These videos have an average length of 120 seconds [6]. We observed that the ground truth captions in ActivityNet often contain sentences that describe the relationships among interesting objects for accurate captions. These datasets are typically used to train end-to-end models and require substantial human effort for collection. Hence, there is a need for a more accurate and efficient method to incorporate spatial information and expand the range of application scenarios.

B. Scene graph generation and image segmentation

Scene Graph Generation (SGG) is an important technology for vision-based scene understanding. Its primary goal is to analyze and recognize visual content and represent their relationship as a graph structure. In previous methods, objects are detected using bounding boxes, followed by the prediction of their pairwise relationships. In 2022, J. Yang et al. introduced Panoptic Scene Graph Generation (PSG), a novel task that challenges models to create a more comprehensive scene graph representation based on panoptic segmentation, as opposed to rigid bounding boxes. In contrast to SGG, the PSG task consistently involves three sub-tasks: semantic segmentation, instance segmentation, and relationship prediction. In their paper, they present a transformer-based method and define 56 predicate relationships. Following their design, Zijian Zhou et al. introduced a new method, Hilo [9], for the PSG task. In their paper, they utilize a DETR-based panoptic segmentation method, Mask2Former [8], as the preprocessing step to obtain accurate panoptic segmentation results. They subsequently generate more precise outcomes.

As a result, the existing methods are typically required to predict all the relationships they detect. However, in application scenarios where only spatial information about interesting objects is necessary, existing methods encounter 3 issues: 1. An excessive number of relations. 2. Predictions are not always accurate. 3. Less emphasis on spatial information.

III. PROPOSED METHOD

A. Problem Setting

Dense video captioning aims to identify a set of temporally-localized captions of suitable length. In our paper, we model this task with the following equation 1, where t_j^s , t_j^e , S_j represent the starting time of the event, ending time of an event, and the caption of an event, respectively. f_e , f_o , f_l represent the video caption module, spatial and ambiguity-reduced information extraction module, and sequence-to-sequence text generation module, respectively. In our evaluation section, we will use t_j^s and t_j^e to measure localization performance, while S_j will be used to measure dense video caption performance. We model the constitute of S_j as equation 2, where S_j^s , S_j^a , S_j^o represent space information, action information, and object information.

$$f(t_j^s, t_j^e; S_j) g_{j=1}^{N_{set}} = f_l(f_e(RGB); f_o(GB D)) \quad (1)$$

$$S_j = f(S_j^s; S_j^a; S_j^o) \quad (2)$$

B. Framework

The core insight of our system is to design a system that can incorporate spatial information to generate detailed and unambiguous captions. The entire pipeline is illustrated in Fig. 2. The system comprises 3 main components: video caption generation, spatial and ambiguity-reduced information extraction, and sequence-to-sequence(seq2seq) text generation.

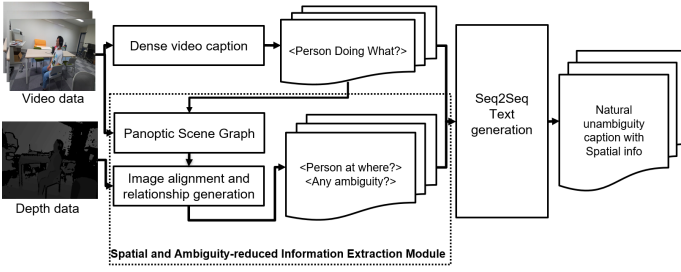


Fig. 2. Architecture of the System

Feature Extractor. Humam Alwassel et al. [14], highlight that traditional feature extractors are often trained for trimmed action classification tasks. Consequently, to obtain features more suitable for untrimmed videos, we have implemented the TSP feature extractor to extract the corresponding features, replacing the use of the C3D [10] features.

Video Captioning Generation. The Parallel Decoding architecture for dense video captioning is a novel design introduced by Teng Wang et al. in 2022 [3]. In contrast to the traditional two-stage pipeline, they employ a transformer encoder and decoder architecture to simultaneously predict the location and caption, followed by merging the location and caption with the predicted scores.

The localization head is designed to perform box prediction and binary classification tasks using a multi-layer perceptron. The results will be saved as sets of tuples $f; t_j^s; t_j^e; c_i^{loc}; g_{j=1}^N$, where c_i^{loc} stand for the confidence of the event. The caption head is responsible for generating sentences that describe the event. As this head cannot determine the boundaries of the event, the deformable soft attention (DSA) and LSTM are used to generate a sentence $S_j = \{w_{j1}; \dots; w_{jM_j}\}$, where the w_{jt} means the word in the j proposal and M_j is the sentence length.

Spatial and Ambiguity-reduced Information Extraction. In order to extract spatial and unambiguous information from RGB-D video sequences and generate concise sentences S_j^s and S_j^o for descriptions, we utilize the previous output $(t_j^s; t_j^e)_{j=1}^{N_{set}}$ as the input and consider the RGB and depth images with a time stamp computed as $(t_j^s + t_j^e)/2$ as the key for the j^{th} sentence.

We implement a Mask2Former to obtain a more accurate image segmentation result and based on that, we implement a fine-tuned OpenPSG [7] module to generate all the relationships of the visual contents.

To extract spatial and unambiguous information from keyframes, we adhere to the following design principles: 1. The main object is chosen from the three largest objects and centered. 2. Additional clarification is introduced only when ambiguity arises with other objects. We reuse the ‘human action’ class from OpenPSG, employ panoptic segmentation from Mask2Former, and align RGB-D images to directly generate information about: [human position related to the main object], [human interact with what] and [which object he interact with] based on the depth of the interest objects.




Fig. 3. Spatial and Ambiguity-reduced Information Extraction

Seq2Seq Text Generation. To generate natural sentences, we utilize a Large Language Model (LLM) to combine and summarize the short sentences. ChatGPT is a well-known model in the field of natural language processing and has produced excellent results for a variety of tasks. However, ChatGPT is not freely available. To make the system applicable to a wide range of applications, we implemented a fine-tuned ‘falcon-7b’¹ model that integrates all information into longer sentences with a more natural structure and human comprehension.

IV. EXPERIMENT

A. Experiment setting

Dataset. There is currently no benchmark dataset that comprises RGB-D images of indoor scenes and annotations in the form of natural sentences describing human behavior. To address this gap, we have gathered a dataset consisting of 30 untrimmed, long videos for training and testing. Within the dataset, we have designed 3 different scenes, 4 position classes, and 2 subjects. Additionally, each video includes both unambiguous and ambiguous action classes.

Metrics. The method is evaluated in two aspects: 1. For location performance, the average precision and average recall across IOU (Intersection over Union) at [0.3, 0.5, 0.7, 0.9] are used. 2. For dense captioning performance, we calculate the BLEU4 [10], METEOR [11], and CIDEr [12] of the matched pairs between generated captions and ground with IOU thresholds of [0.3, 0.5, 0.7, 0.9].

B. Implementation details

We collect the dataset with one Microsoft Azure Kinect. The video is 30 FPS with 720P. We implemented a TSP feature-updated PDVC and pre-trained it on ActivityNet dataset as the backbone. In the spatial and ambiguity-reduced information extraction module, we implement a Mask2Former model which is trained on the coco panoptic dataset for detection and segmentation of daily objects. We design 4 position relations (left, right, front, back) and 4 human-object interaction classes (walking on, standing on, lying on, sitting on) to generate based on OpenPSG.

¹Falcom LLM <https://falconnllm.tii.ae>

