

Document downloaded from:

<http://hdl.handle.net/10251/141941>

This paper must be cited as:

Bolaños, M.; Peris-Abril, Á.; Casacuberta Nolla, F.; Soler, S.; Radeva, P. (01-2). Egocentric video description based on temporally-linked sequences. *Journal of Visual Communication and Image Representation*. 50:205-216. <https://doi.org/10.1016/j.jvcir.2017.11.022>



The final publication is available at

<https://doi.org/10.1016/j.jvcir.2017.11.022>

Copyright Elsevier

Additional Information

Accepted Manuscript


Egocentric Video Description based on Temporally-Linked Sequences

Marc Bolaños, Álvaro Peris, Francisco Casacuberta, Sergi Soler, Petia Radeva

PII: S1047-3203(17)30231-6
DOI: <https://doi.org/10.1016/j.jvcir.2017.11.022>
Reference: YJVICI 2097

To appear in: *J. Vis. Commun. Image R.*

Revised Date: 9 November 2017
Accepted Date: 28 November 2017



Please cite this article as: M. Bolaños, A. Peris, F. Casacuberta, S. Soler, P. Radeva, Egocentric Video Description based on Temporally-Linked Sequences, *J. Vis. Commun. Image R.* (2017), doi: <https://doi.org/10.1016/j.jvcir.2017.11.022>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Egocentric Video Description based on Temporally-Linked Sequences

Marc Bolaños^{a,b}, Álvaro Peris^c, Francisco Casacuberta^c, Sergi Soler^a, Petia Radeva^{a,b}^aUniversitat de Barcelona, Barcelona, Spain^bComputer Vision Center, Bellaterra, Spain^cPRHLT Research Center, Universitat Politècnica de València, València, Spain**Abstract**

Egocentric vision consists in acquiring images along the day from a first person point-of-view using wearable cameras. The automatic analysis of this information allows to discover daily patterns for improving the quality of life of the user. A natural topic that arises in egocentric vision is storytelling, that is, how to understand and tell the story relying behind the pictures.

In this paper, we tackle storytelling as an egocentric sequences description problem. We propose a novel methodology that exploits information from temporally neighboring events, matching precisely the nature of egocentric sequences. Furthermore, we present a new method for multimodal data fusion consisting on a multi-input attention recurrent network. We also release the *EDUB-SegDesc* dataset. This is the first dataset for egocentric image sequences description, consisting of 1,339 events with 3,991 descriptions, from 55 days acquired by 11 people. Finally, we prove that our proposal outperforms classical attentional encoder-decoder methods for video description.

Keywords: egocentric vision; video description; deep learning; multi-modal learning

1. Introduction

Egocentric vision (Doherty et al., 2013; Betancourt et al., 2015; Bolaños et al., 2017) is a recent topic in the computer vision field, with the goal of analyzing the visual information provided by wearable cameras, which have the capability to acquire images or videos from a first person point-of-view. The analysis of these images can provide useful information about the behavior of the user for several complementary topics like social interactions, scene understanding, time-space-based localization, action recognition, nutritional habits, among others. Thus, enabling to understand the whole story and user's behavior behind the pictures (i.e. automatic storytelling) followed by inferring his/her actions and habits, could lead to a better quality of life for the user. Considering the sheer amount of data wearable cameras provide, there is a need to create automatic algorithms to analyze and summarize them. In this paper, we focus on the specific topic of creating automatic diaries of the life of the user by means of textual descriptions. One of the possible health-related applications for the automatic

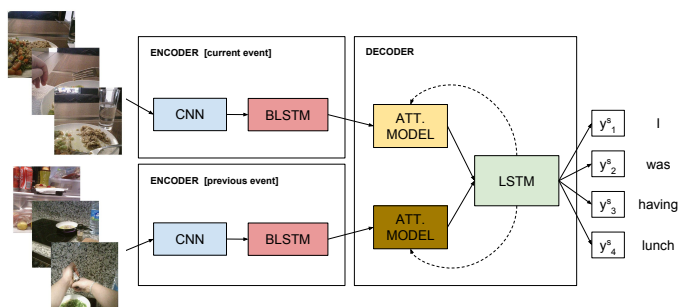


Figure 1: General outline of the proposed temporally-linked multi-input attention model (TMA).

Email addresses: marc.bolanos@ub.edu (Marc Bolaños), lvapeab@prhlt.upv.es (Álvaro Peris), fcn@prhlt.upv.es (Francisco Casacuberta), ssolerso8@alumnos.ub.edu (Sergi Soler), petia.ivanova@ub.edu (Petia Radeva)

diary construction could be the treatment of patients with dementia. As proven by [Spector et al. \(2003\)](#); [Sellen et al. \(2007\)](#), the daily review of egocentric pictures taken by this kind of patients can help them recover partially their cognitive capabilities. The incorporation of additional automatically extracted textual information and comparing it to the user’s one could give novel tools to complete cognitive frameworks for memory enhancement.

The egocentric captioning problem can be seen as an instantiation of a video description task ([Venugopalan et al., 2015](#)): the system receives as input a sequence of images and must produce a sentence as a sequence of words describing it. The problem is challenging due to two main reasons: on the one hand, most wearable cameras have a small field of view the other hand, lifelogging cameras have low temporal resolution (2-3 fpm) that is events are not videos, but collection of temporally ordered images. Hence, the problem is much more difficult than conventional videos and we need powerful techniques and algorithms from computer vision and natural language understanding to address the problem of video description. The recent development of deep learning techniques has allowed a breakthrough in the computer vision and natural language processing fields. The ability of training complex models has pushed forward many research areas and nowadays, Convolutional Neural Networks (CNNs) constitute one of the most powerful tools in the computer vision field. As the problem at hand involves sequence learning or prediction, a natural choice are Recurrent Neural Networks (RNNs), which are powerful sequence modelers. Specifically, the use of gated units, such as Long Short-Term Memory (LSTM) ([Hochreiter and Schmidhuber, 1997](#)) or gated recurrent units ([Cho et al., 2014](#)) allows to properly model sequences with long and complex relationships.

The combination of RNNs together with CNNs is therefore widely employed for tackling multi-modal learning tasks, involving the combination of vision and language. Some examples of multi-modal learning problems are image or video description, dense captioning, visual question answering, multi-modal interaction to mention a few.

In this work, we develop a fully-neural end-to-end system for egocentric captioning that we call Temporally-linked Multi-input Attention model (TMA). We hypothesize that, within the egocentric sequences captioning problem, given a day, some of the events that compose it can follow a temporally logical relation. That is, previous actions occurred during a day can influence the following ones. Therefore, we need a model able to capture and learn this relation. Our proposal is able to embed previous information coming from either image or language. In [Fig. 1](#), we show a simplified outline of the proposed method. In this example, we illustrate how previous event’s frames can be employed in order to help to the process of video description of the analyzed event. Note that in the example, the user is cooking and next he proceeds to take the food. This temporal relationship between actions aids the model when predicting the current description.

Furthermore, we publish the first egocentric dataset for event captioning, composed of 55 complete days containing nearly 55,000 images in total, acquired by 9 different people.

The main contributions of this work are the following:

- We present a novel captioning model, which incorporates the information from previous events into the current decoding state.
- We present a new LSTM model capable of combining information from multiple inputs and modalities, as well as applying a separate attention mechanism to each of them.
- We conduct experiments on the new dataset and compare our model with other classical captioning architectures. Results show that using information from previous events provides better generalization.
- We present the first dataset for egocentric sequences captioning, named EDUB-SegDesc, based on describing the events that take place along a day.
- We make public both dataset and model, in order to make results reproducible and foster the research in this topic.

The rest of the paper is structured as follows: the related work is reviewed in [Section 2](#). Next, we describe our model and its main components in [Section 3](#). We present the egocentric captioning dataset in [Section 4](#). We set up our experimental framework and show the obtained results in [Section 5](#). Such results are analyzed in [Section 6](#). Finally, conclusions and future research lines are drawn in [Section 7](#).

2. Related work

In recent years, deep learning techniques have provided tremendous advances in the computer vision field. More precisely, CNNs (Krizhevsky et al., 2012) have proved to excel on the task of learning rich image representations and consequently have served as feature extractors, providing record-breaking results in most tasks. On the other hand, RNNs have proved to be powerful sequence modelers. They have been recently used in many sequences learning tasks, including machine translation (Sutskever et al., 2014; Bahdanau et al., 2015), image (Xu et al., 2015) and video (Yao et al., 2015; Peris et al., 2016) captioning, or visual question answering (Fukui et al., 2016; Bolaños et al., 2016). Most of these problems involve tackling multi-modal data (i.e. text and images) (Toselli et al., 2011), which means that, in all of them, both CNNs and RNNs are commonly employed.

Such systems are built under the encoder-decoder framework: an encoder processes the input and computes a meaningful representation of it. Next, the decoder takes this representation and produces the desired output, typically a sentence in natural language. Given their power as feature extractors, CNNs are usually employed as encoders (Xu et al., 2015; Yao et al., 2015; Fukui et al., 2016). Nevertheless, if the input signal has a sequential component (as in machine translation or video captioning), RNNs, solely or together with CNNs, can also act as encoders (Sutskever et al., 2014; Peris et al., 2016).

The vanilla encoder-decoder architecture was enhanced with attention mechanisms (Bahdanau et al., 2015; Xu et al., 2015; Yao et al., 2015), allowing the decoder to selectively focus on parts of the input during the text generation process. Therefore, the system is able to properly deal with long and complex inputs.

Regarding the video description problem, which consists in generating a natural language description given an input video or sequence of images, several proposals have been published in recent years. The general framework used also consists of applying one or several CNNs to the input images, followed by a generative LSTM for sentence generation. Several improvements and variations have been proposed to the main model, some examples being the use of additional CNN representations like optical flow (Venugopalan et al., 2015; Yao et al., 2015), Bidirectional LSTMs (BLSTM) in the encoder (Peris et al., 2016), attention mechanisms in the decoder (Yao et al., 2015), hierarchical information (Pan et al., 2016), external linguistic knowledge (Venugopalan et al., 2016) or multi-modal attention mechanisms (Hori et al., 2017), among others.

All these works relating video captioning assume a sample-wise independence. This is that a sample is meant to be unrelated to the next one. This may represent a limitation in tasks which aim to model continuous events, arbitrarily split. The task addressed in this work belongs to this class of problems: we aim to describe the whole day of a user. For tackling it, we divide a day into events and describe each one of these events. Due to this division, our samples are conditioned between them. For example, if the user is at the office, it is likely that in the following event he/she uses a computer. Therefore, the classical sample-wise independence assumption becomes excessively severe, as critical information may be potentially lost. We propose a novel model that takes into account information from past events, being suitable for this kind of tasks.

Only one work has recently proposed to take into account mid-term temporal information in conventional videos (Krishna et al., 2017). Their model extracts C3D features on variable-sized conventional videos and computes several temporal segmentation proposals in short actions. Later, it generates textual descriptions for each action incorporating contextual information from past and future actions belonging to the same video. In comparison to our problem of egocentric lifelogging data, apart from the different perspective of the sequences, their videos are much shorter. Moreover, instead of consisting of several long-term events, theirs contain mid-term short actions belonging to the same event. Even though there are important differences to both problems, the authors prove that incorporating past and future information helps predicting the output of current events or actions.

In the field of egocentric vision (Doherty et al., 2013; Betancourt et al., 2015; Bolaños et al., 2017), some authors have worked on related problems like activity recognition (Iwashita et al., 2014; Cartas et al., 2017) or event classification (Castro et al., 2015). The main handicap of these methods is that they provide simple labels from a user-oriented perspective, which provide more limited and predefined information compared to generating free textual descriptions. Regarding any of these methods, the first step needed in order to provide a coherent description of the actions and events happening in egocentric images is the application

of an automatic segmentation (Dimiccoli et al., 2017; Poleg et al., 2016; Lu and Grauman, 2013) of the complete day of the user. After the application of the day segmentation and having as output the set of events of the day, we can acquire the self-contained units of information needed in order to apply either event captioning, activity recognition or event classification models.

Considering the specific problem at hand, which is the generation of textual descriptions of egocentric events, few works have been proposed in the state-of-the-art. Only one of them tackled it as a video description problem and from an end-to-end perspective. Fan and Crandall (2016) explored the problem of creating image diaries in order to apply image retrieval. As a step of its process, the authors proposed an image captioning method that processes one image at a time by applying a CNN for features extraction and a RNN for sentence generation. Later, they proposed grouping the images and applying a sentences fusion technique for providing the final captions for each event along the day. Goel and Naik (2016) also focused on the task of image retrieval for both conventional and egocentric videos. It applied a simple method of video description composed of a CNN for feature extraction, a Bidirectional LSTM for image sequence combination and a LSTM in the decoder for applying the final sentence generation for 5-seconds-long clips of video. The purpose of this method in the work was to provide semantic information of the available data for the posterior retrieval.

In this work, we focus on event textual description. We assume that the events are previously extracted, manually or automatically (Dimiccoli et al., 2017). To this aim, we propose an end-to-end model specifically trained for egocentric day sequences captioning. As argued before, it is important to take into account relevant information from previous samples, for generating descriptions of the current event. Thus, our model treats the image sequences from different events as temporally-linked units. It jointly models and exploits intra-eventual information (flowing through the frames of a single event) together with inter-eventual information (linking temporal sequences of neighboring events).

3. Methodology

The problem, we face in this work, involves a multi-modal learning task. On the one hand, we have a huge amount of information coming from the sequence of images captured by the wearable camera. On the other hand, each sequence of images has the correspondent caption associated, which describes the event occurring in the life of the user from an egocentric point of view.

Although the sequences of images may be segmented in independent semantic units to some extent (i.e. events), given the nature of the egocentric images, it is obvious that there exists a dependency between temporally neighboring events (i.e. most given events experienced in the day of a person have a relationship to the previous one he/she lived). We aim to exploit this information by incorporating into our model, at a given temporal point s , the information extracted from the previous event $s - 1$. This information can be either textual (previous caption), visual (previous sequence of images) or both. Therefore, we develop a model that takes into account the information from both sources: the current sequence of frames together with the information coming from previous events.

3.1. Egocentric captioning

Similarly to the video description problem, in the egocentric task, we have as input a sequence of frames and we want to output a sentence that describes the input. This problem has already been tackled in the literature on conventional videos with multiple variations (Pan et al., 2016; Venugopalan et al., 2015, 2016; Yao et al., 2015; Peris et al., 2016). In the latter work, the frames were encoded by a CNN and a BLSTM network. The representation obtained was fed to an LSTM decoder equipped with an attention mechanism, which generated the corresponding caption. We use this architecture as starting point for developing our system.

The main difference that characterizes the egocentric captioning problem is that, unlike in the classical video description approach, we do have temporally-linked events, which share a relationship (e.g. if in an event the user enters into an office and sits on his/her table, it is highly likely that in the following event he/she will be using a computer).

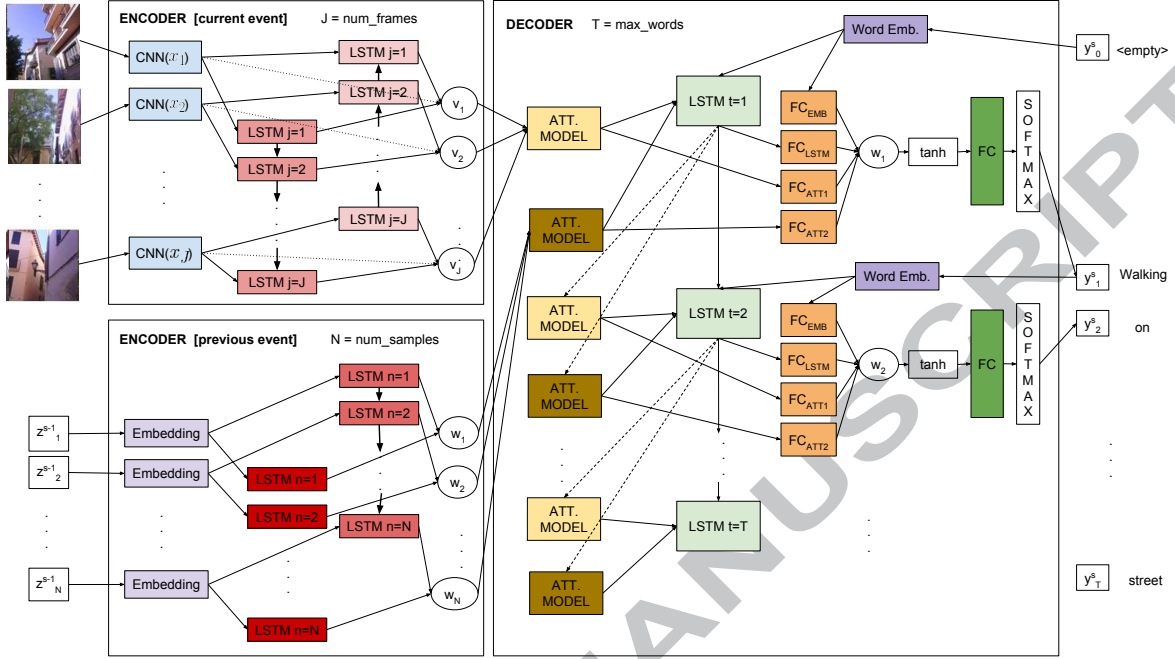


Figure 2: Architecture of the proposed TMA model. Unlike the traditional approaches, our architecture consists of at least two encoder stages; one for the current sequence and another one (or more) for the previous sequence; and a decoder stage that combines the information of all the previous stages using a multi-input attention LSTM.

Thus, we propose a system able to take advantage from both the current sequence of egocentric images and the action happened in the previous event. In Fig. 2 we detail the architecture of our model. The input to the system is a sequence of frames $X^s = x_1, \dots, x_J$ and the sequence of information providing from the previous event, $Z^{s-1} = z_1^{s-1}, \dots, z_N^{s-1}$. This latter sequence can either be the previous textual description, Y^{s-1} or the previous sequence of frames, X^{s-1} , or both. The current sequence of images is processed in the encoder of the current event, which is composed of a CNN (blue) and of a BLSTM (light red), producing the sequence of visual features v_1, \dots, v_J . Meanwhile, the data Z^{s-1} is processed in the encoder of the previous event, which is composed of an embedding mechanism (light purple) and of another BLSTM (dark red), which computes the final sequence of representations w_1, \dots, w_N . The embedding mechanism used in this encoder can either be a CNN (blue) when dealing with images or a word embedding matrix (dark purple) when dealing with a textual representation. Note that this word embedding matrix is the same as the one in the decoder. These sequences of feature vectors are given as input to the decoder LSTM by means of two independent attention mechanisms (yellow and brown). In addition, at each time-step t , the decoder takes as input the word embedding of the previously generated word ($\mathbf{E}(y_{t-1}^s)$). Finally, a series of skip connections is combined in an element-wise summation after embedding the different modalities by using fully-connected layers (orange). We include a skip connection for each attention mechanism, one for the LSTM hidden state and another one for the previously generated word. The sequence of time-steps for the current caption is indexed by t , while the sequence of events is indexed by s .

We extend the classical encoder–decoder approach by adding different encoders: one for the current sequence of images and another for the information from previous events. Such encoders are built using CNNs and BLSTMs networks, which we detail in Section 3.2. In Section 3.3, we proceed to explain the decoder, which is also extended by adding a multi-input attention LSTM for dealing with input sequences of multiple modalities. Finally, in Section 3.4, we summarize the whole pipeline applied in the proposed TMA model.

3.2. Encoder

3.2.1. Convolutional Neural Networks

CNNs have proved to be great modelers of image representations. Several models have been proposed in the state-of-the-art (e.g. AlexNet (Krizhevsky et al., 2012), VGG (Simonyan and Zisserman, 2014), GoogLeNet (Szegedy et al., 2015), ResNet (He et al., 2016), etc.), most of them originally for the problem of object recognition or detection (Russakovsky et al., 2015). Furthermore, they have been proven to serve as excellent feature extractors for other related tasks in the computer vision field.

Without loss of generality, in our proposal, we make use of the well-known GoogLeNet architecture (Szegedy et al., 2015) pre-trained on the ILSVRC challenge data as an extractor for training our complete TMA model. Although it varies depending on the problem and data, GoogLeNet has proven to be one of the best models for feature extraction, offering a trade-off between performance, number of parameters and computational time. We have to note that it is possible to use any alternative state-of-the-art CNN model for dealing with the images representation in our TMA model.

3.2.2. Long Short-Term Memory and Bidirectional networks

Since the input of our system is a sequence (of images), we can effectively model it and its relationships by using a RNN. More precisely, in order to learn a representation of the input sequences and its relationships, we use a LSTM network (Hochreiter and Schmidhuber, 1997; Gers et al., 2000). Unlike simple units, LSTM cells feature an additional state, the so-called memory state. The network controls how information flows through the unit by means of three gates, namely input, output and forget gates. Such gates modulate the amount of information that the network will incorporate from the input, from the output at the current time-step, or the amount that it will store for future time-steps. Therefore, LSTM networks are able to model long and complex sequences of information, reducing the vanishing gradient problem (Bengio et al., 1994).

Hence, at a given time-step j , the hidden state \mathbf{v}_j is the memory state \mathbf{c}_j controlled by the output gate \mathbf{o}_j . The memory state depends on the previous memory state (modulated by the forget gate \mathbf{f}_j) and an updated memory state (modulated by the input gate \mathbf{i}_j). The updated state is obtained from the input \mathbf{x}_j and the previous hidden state \mathbf{v}_{j-1} . All three gates are also computed from \mathbf{x}_j and \mathbf{v}_{j-1} (Gers et al., 2000).

With LSTM units, we manage how information flows through the network, and we are able to model long-term relationships. Nevertheless, this information only flows in a time direction. If we have the full day to analyze, we can take profit from relationships from the previous to the next events, but also from information flowing from next to previous events. Bidirectional networks (Schuster and Paliwal, 1997) cope with both sequence directions by maintaining two independent recurrent layers: the forward layer processes the sequence from the past to the future and the backward layer processes the sequence reversed in time. Therefore, we can extract relationships flowing in both time directions of our input signal. Once the full sequence has been processed in both directions, forward and backward layers are combined, typically by concatenating their hidden states (Bahdanau et al., 2015; Peris et al., 2016).

3.3. Decoder: Multi-input Attention LSTM

Since we are in a multi-modal scenario dealing with sequences of text and images, we propose to use an LSTM network that accepts multiple inputs—from multiple sources—and combines them after applying




Figure 3: Architecture of a multi-input LSTM cell at time-step, t . In this case, the inputs to the cell are the word embedding of the previous word ($\mathbf{E}(y_{t-1}^s)$), the previous LSTM hidden state, (\mathbf{h}_{t-1}) and the context vectors from two different attention mechanisms, ($\mathbf{z}_t, \mathbf{z}'_t$).

an attention mechanism for each input. The combination, as well as the attention mechanisms, are thus learned together with the rest of the model.

Attention mechanisms (Bahdanau et al., 2015; Xu et al., 2015) help the decoder to selectively focus on parts of the input sequence, depending on the decoding step. More formally, given a sequence of J input vectors $\mathbf{v}_1, \dots, \mathbf{v}_J$ previously calculated by the encoder, at each decoding time-step t , the attention mechanism weights them into a single context vector \mathbf{z}_t :

$$\mathbf{z}_t = \sum_{j=1}^J \alpha_{jt} \mathbf{v}_j, \quad (1)$$

being α_{jt} the weight given to the j -th feature vector at time, t . At each time-step, the weights are calculated according to the scores obtained by a soft-alignment model, which determines how much a feature vector does influence the outcome of the current word. The aligner is implemented by a single-layer perceptron:

$$e_{jt} = \mathbf{w}^\top \tanh(\mathbf{W}_a \mathbf{h}_{t-1} + \mathbf{U}_a \mathbf{v}_j), \quad (2)$$

where \mathbf{w} , \mathbf{W}_a and \mathbf{U}_a are the perceptron parameters and \mathbf{h}_{t-1} is the decoder hidden state from the previous time-step. For clarity, we omit the bias term in this equation, as well as in the rest of the paper. The aligner is then followed by a softmax function in order to obtain the final normalized weights, α_{jt} :

$$\alpha_{jt} = \frac{\exp(e_{jt})}{\sum_k \exp(e_{kt})}. \quad (3)$$

Our multi-input decoder is a natural extension of classical LSTM networks, which is able to process N different inputs. For simplicity, in this section we assume that we have inputs from two different modalities. Such inputs are processed by the respective attention models. Hence, at a given time-step, t , we have two inputs from the attention models (\mathbf{z}_t and \mathbf{z}'_t). Furthermore, in order to boost the decoder capabilities, we introduce the word embedding of the previously generated word as an additional input. Fig. 3 illustrates the LSTM cell.

As in regular LSTMs, in our multi-input network, the hidden state depends on the memory state and the output gate:

$$\mathbf{h}_t = \mathbf{o}_t \odot \mathbf{c}_t, \quad (4)$$

where \odot denotes the element-wise multiplication. \mathbf{c}_t is the memory state, defined as:

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t, \quad (5)$$

where \mathbf{c}_{t-1} is the previous time-step memory state and $\tilde{\mathbf{c}}_t$ is the updated memory state. \mathbf{f}_t and \mathbf{i}_t are the forget and input gates, which modulate them. $\tilde{\mathbf{c}}_t$ is computed taking into account the attended representations of the inputs (\mathbf{z}_t and \mathbf{z}'_t), the last word generated by the decoder (y_{t-1}) and the previous hidden state (\mathbf{h}_{t-1}):

$$\tilde{\mathbf{c}}_t = \tanh(\mathbf{W}_c \mathbf{E}(y_{t-1}) + \mathbf{U}_c \mathbf{h}_{t-1} + \mathbf{A}_c \mathbf{z}_t + \mathbf{B}_c \mathbf{z}'_t), \quad (6)$$

where \mathbf{E} is a word embedding matrix and $\mathbf{E}(y_{t-1})$ denotes the word embedding of the previously generated word. \mathbf{W}_c , \mathbf{U}_c , \mathbf{A}_c and \mathbf{B}_c are the weight matrices for the word embedding of y_{t-1} , the previous hidden state and both context vectors from the alignment attention models, respectively.

The forget, input and output gates also depend on the context vectors, the previous word and the previous hidden state:

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{E}(y_{t-1}) + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{A}_f \mathbf{z}_t + \mathbf{B}_f \mathbf{z}'_t), \quad (7)$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{E}(y_{t-1}) + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{A}_i \mathbf{z}_t + \mathbf{B}_i \mathbf{z}'_t), \quad (8)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{E}(y_{t-1}) + \mathbf{U}_o \mathbf{h}_{t-1} + \mathbf{A}_o \mathbf{z}_t + \mathbf{B}_o \mathbf{z}'_t). \quad (9)$$

3.4. Temporary-linked Multi-input Attention model

Considering our TMA model as a whole, first we process the sequence of images captured by the camera for the current event s , $X^s = x_1^s, \dots, x_J^s$. As image encoder, we use a CNN that extracts features from X^s . Next, we apply a BLSTM network, in order to capture the temporal relationships existing in the sequence. Finally, for each frame x_j^s , we compute a feature vector \mathbf{v}_j by concatenating the representation extracted by the CNN together with the forward and backward hidden states from the BLSTM.

Simultaneously, we process the previous event (Z^{s-1}). The information of the previous event can be either the previous sequence of egocentric images (X^{s-1}), the previous caption (Y^{s-1}), or both. In case of processing the previous sequence of images, the method is the same as the one applied on the current event: a CNN combined with a BLSTM network. If the information providing from the previous event is the caption $Y^{s-1} = y_1^{s-1}, \dots, y_N^{s-1}$, words are projected to the continuous space by means of the embedding matrix, shared with the decoder. Next, an additional BLSTM processes this sequence of word embeddings. We concatenate the forward and backward BLSTM hidden states for obtaining a contextual representation for each input word. In case, we are using both modalities, the model is modified for having three different encoders: 1) current event frames X^s , 2) previous event frames X^{s-1} and 3) previous event output caption Y^{s-1} , as well as three attention mechanisms in the decoder, one for each input. For the sake of clarity and without loss of generality, let us assume from now on that we only have one modality as additional input from the previous event.

Therefore, our encoder produces two sequences of features: $\mathbf{v}_1, \dots, \mathbf{v}_J$, referring to the current video and $\mathbf{w}_1, \dots, \mathbf{w}_N$, related to the previous event. We integrate this into the decoder by means of two independent attention mechanisms. Each attention mechanism weights the elements of its respective sequence and computes a joint representation of it, \mathbf{z}_t and \mathbf{z}'_t respectively, taking into account the previous decoding state \mathbf{h}_{t-1} (see Section 3.3). Such representations, together with the previously generated word, are the inputs of the multi-input LSTM (described in Section 3.3).

Finally, at each time-step t , and in order to predict the following word in the output sequence, we define a probability distribution over the task vocabulary as follows:

$$\mathbf{p}_t = \text{softmax}(\mathbf{U}_p \tanh(\mathbf{W}_a \mathbf{h}_t + \mathbf{W}_b \mathbf{z}_t + \mathbf{W}_c \mathbf{z}'_t + \mathbf{W}_d \mathbf{E}(y_{t-1}^s))), \quad (10)$$

where \mathbf{h}_t is the hidden state from the decoder, \mathbf{z}_t and \mathbf{z}'_t are the contextual representations of the current and previous inputs computed by their respective attention models and $\mathbf{E}(y_{t-1}^s)$ is the word embedding of the preceding word in the current sequence. $\mathbf{W}_a, \mathbf{W}_b, \mathbf{W}_c, \mathbf{W}_d$ are the skip-connections weight matrices (orange in Fig. 2). \mathbf{U}_p is the last layer matrix (dark green in Fig. 2).

The produced distribution, \mathbf{p}_t represents the probability of a word given the input video X^s , the previous event Z^{s-1} , and the words generated so far in the current sequence y_1^s, \dots, y_{t-1}^s :

$$\mathbf{p}_t = p(y_t^s | X^s, Z^{s-1}, y_1^s, \dots, y_{t-1}^s). \quad (11)$$

We approximate the most likely caption by using a beam search method (Sutskever et al., 2014). The complete model parameters (θ) are jointly estimated over a dataset \mathcal{S} , which consists of S image sequences and caption pairs. The training objective is to maximize the log-likelihood of \mathcal{S} with respect to θ :

$$\hat{\theta} = \arg \max_{\theta} \sum_{s=1}^S \sum_{t=1}^{T_s} \log(p(y_t^s | X^s, Z^{s-1}, y_1^s, \dots, y_{t-1}^s; \theta)), \quad (12)$$

where T_s is the length of the s -th caption. If the previous event, Z^{s-1} of a given image sequence X^s is undefined, we introduce an artificial empty event as Z^{s-1} .

4. EDUB-SegDesc dataset

EDUB-SegDesc¹ is a dataset that can be used either for egocentric events segmentation (Dimiccoli et al., 2017) or for egocentric sequences description. It was acquired by the wearable camera Narrative², taking a picture every 30 seconds (2 fpm). It consists of 55 days acquired by 9 people. Each day was manually segmented in events or sequences following the same criteria as in Dimiccoli et al. (2017): “An event is a semantically perceptual unit that can be inferred by visual image features, without any prior knowledge of what the camera wearer is actually doing.” In Fig. 4 a histogram with the duration (in minutes) of the resulting segments is shown. We can observe a wide variability in duration. Most of the segments (around 65%) have relatively short durations of 15 minutes or less, but there also exist several long events (around 5%) with a duration longer than 100 minutes.




Figure 4: Histogram depicting the duration of the segments in the EDUB-SegDesc dataset. All segments with a duration equal or greater than 200 minutes appear grouped in the last bin (+200).

The dataset contains a total of 48,717 images, divided in 1,339 events (or image sequences) and 3,991 captions, and has an average of 3 captions per event. It was divided in training, validation and test splits making sure that all the sequences from the same day should belong to the same data split. The division results in the figures depicted in Table 1.

Table 1: Figures of the EDUB-SegDesc dataset, according to each partition, training, validation and test.

EDUB-SegDesc	Training	Validation	Test	Total
#days	39	7	9	55
#images	32,664	7,301	8,752	48,717
#segments	889	204	246	1,339
#descriptions	2,652	598	741	3,991

In Fig. 5 we show the number of co-occurrences in consecutive segments from some manually chosen keywords. This highlights the natural relationships found in consecutive events. Some notable examples of concepts appearing in consecutive events include: ‘people’ in past events followed by ‘talked’ in current events (social events); ‘laptop’ followed by ‘work’ (work-related events); ‘street’ followed by ‘entered’ (going from an outdoors to an indoors environment); ‘station’ followed by ‘train’ (transport-related events) or ‘phone’ followed by ‘street’ (events related to using the mobile phone on the street).

¹<http://www.ub.edu/cvub/edub-segdesc>

²www.getnarrative.com

Fig. 6 shows several word-related statistics: Occurrences of the most common words and bigrams and histogram sentence lengths. Note that the number of appearances of the word 'I' is very high both in the single word and the bigram counts given the egocentric nature of the dataset. Other commonly occurring words and bigrams are the ones related to events where the user is 'walking' and/or on the 'street'. Curiously, there is a certain bias in the number of words contained in the annotations, where a considerable number of the sentences are composed of 5 words. Some examples of 5-word sentences are 'I went to my office', 'I worked with my laptop', or 'I walked on the street'.

Finally, we show some examples of the events and sentences contained in the dataset (Fig. 7). The low temporal resolution of the camera used (2 frames per minute) becomes clear in dynamic events, where the user moves and this causes to have highly variable environments. This fact, together with the limited information present in some of the images highlights the difficulty of the problem.




Figure 6: Word-related statistics of the dataset. Occurrences of the most common words (left), occurrences of the most common bigrams (center), and histogram of number of words in all the sentences (right).

The EDUB-SegDesc dataset was specifically acquired and labeled for the purpose of developing a model for describing and understanding all the events appearing along the day of a person. Thus, the main application of using egocentric sequences for textual descriptions generation is providing a memory aid for MCI patients.

5. Experiments and results

In this section we set up the experimental environment. We also define the metrics employed to evaluate it. Finally, we show the results obtained.

5.1. Metrics and evaluation

We evaluated our proposal using standard image and video captioning metrics. We used the COCO-Caption evaluation package (Chen et al., 2015) and computed three metrics:

BLEU (Papineni et al., 2002): compares the ratio of n-gram structures that are shared between the system hypotheses and the reference sentences. We report BLEU-4 in our results.

METEOR (Lavie and Denkowski, 2009): this metric was introduced to solve the lack of the recall component when computing BLEU. It computes the F1 score of precision and recall between hypotheses and references. In addition, it considers exact, stemmed, synonyms and paraphrase matches.

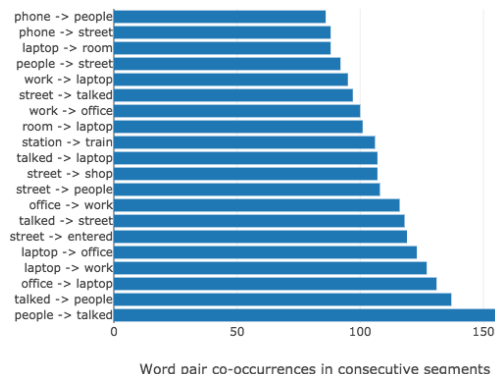


Figure 5: Number of co-occurrences in consecutive segments of several manually chosen keywords. Word pairs are shown in the following format: `word1 -> word2`, where `word1` appears in the past event and `word2` appears in the current event.

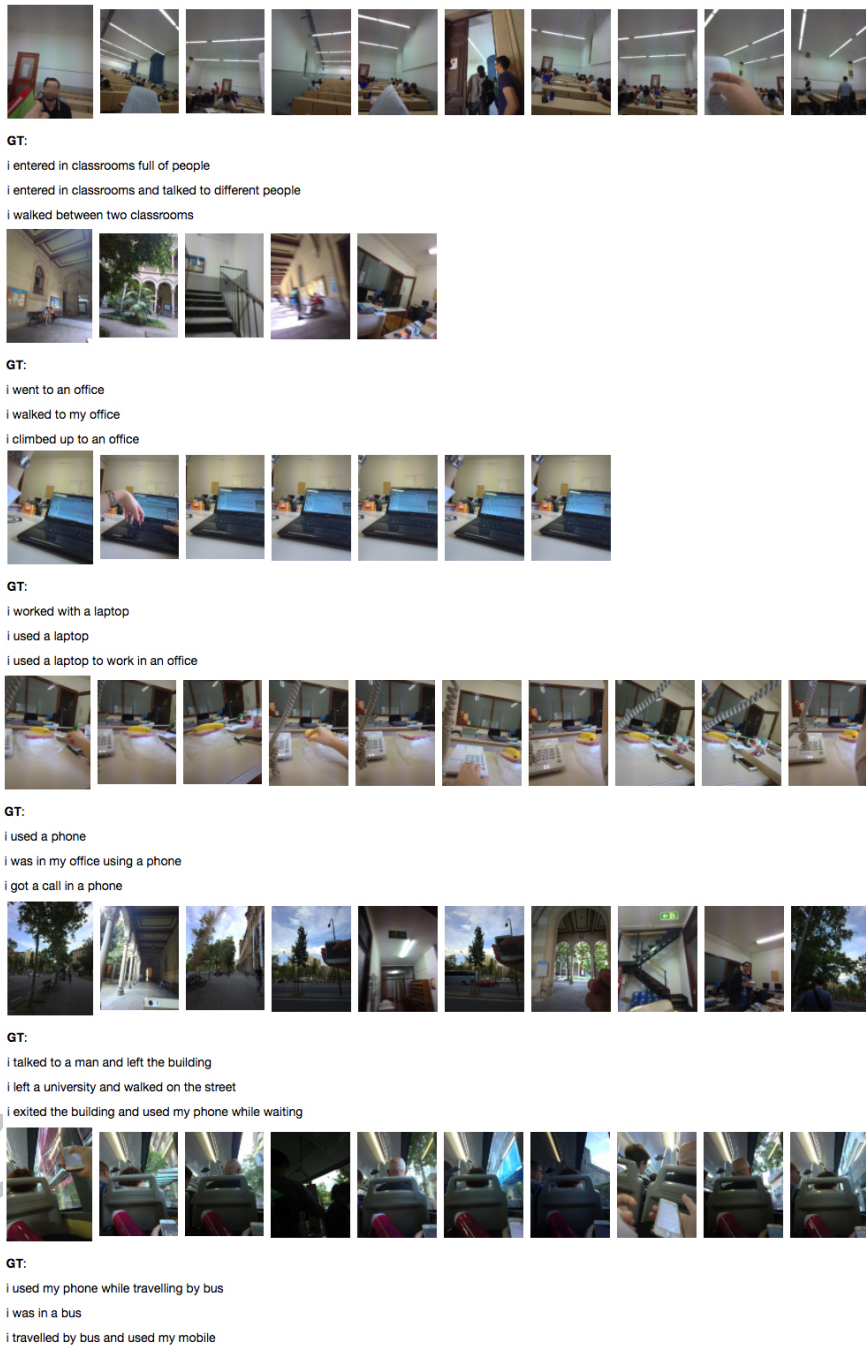


Figure 7: Subset of a day from the dataset EDUB-SegDesc. We show some consecutive events and their respective GT sentences. The difficulty of the problem is highlighted by the dynamism and instability of the scene, when the user is moving. Particularly difficult examples can be seen in the second event: the GT sentences explain that the user is heading to his/her office, but this aspect is only manifested in the last image of the sequence.

CIDEr (Vedantam et al., 2015): similarly to BLEU, it computes the number of matching n-grams, but penalizing n-grams frequently found in the training set. The CIDEr metric ranges from 0 (minimum quality) to 10 (maximum quality).

5.2. Experimental setup

All the neural models that we compare were built with the Keras³ and Theano (Theano Development Team, 2016) libraries. We release the source code of our implementation for future comparisons⁴. The full model was jointly trained end-to-end on the EDUB-SegDesc dataset, except for the CNN, which was already pre-trained for object detection on ImageNet (Russakovsky et al., 2014) and remained static during the training of our model.

The main hyper-parameters of the model were selected according to the analysis reported in the video description task (Peris et al., 2016). Moreover, since we conducted experiments with pre-trained models, we must keep fixed some hyper-parameters. Therefore, we used word embeddings of size 301, the encoder BLSTM had 717 units in the forward and 717 more units in the backward layers and the decoder LSTM had 484 units. The initial state of the decoder LSTM was initialized with the hyperbolic tangent of the mean of the video features obtained by the encoder (Xu et al., 2015). We took at most 26 frames evenly distributed from each complete sequence of our dataset.

We trained our model by stochastic gradient descent (SGD). After doing experiments either using Adam or Adadelata optimizers, we observed that for some model variations, the optimal performance was reached either using Adadelata (Zeiler, 2012) with a learning rate of 1.0 and without learning rate decay; or using Adam (Kingma and Ba, 2014) with an initial learning rate of 0.001 and a decay of 0.995 at the end of every epoch. Thus, we report the best results on each case. During training, the norm of the gradients was clipped to 10 (Pascanu et al., 2013).

In order to prevent over-fitting, we used batch normalization (Ioffe and Szegedy, 2015). Contrarily to other works, we observed that the use of dropout (Srivastava et al., 2014) combined with batch normalization, produced better performance, ($p = 0.5$). We also applied weight decay (10^{-4}) and Gaussian noise ($\sigma = 10^{-2}$) to the non-recurrent weights.

We set our batch size to 64 and used an early stop criterion on the validation set based on BLEU-4, setting our patience to 20 and checking the performance each 50 updates. The size of the beam during the search was 10.

In order to minimize the influence of randomness in our results, mostly due to the weights random initialization, each experiment was run 5 times, and reported the median value of such runs. The model was trained either from scratch, exclusively using the data from the EDUB-SegDesc dataset or reusing the pre-trained weights from certain layers that were learned in different model variations. We studied the inclusion of word embeddings, obtained using the skip-gram model from Mikolov et al. (2013) and trained on part of Google News dataset. We also tested training the decoder as a language model on the 1 Billion words dataset (Chelba et al., 2013), but results in both cases were not better. Finally, we also tested reusing the pre-trained weights from the video captioning model from Peris et al. (2016) (trained on the MSVD dataset published by Chen and Dolan (2011)), which proved to improve the results under certain model configurations. The Microsoft Research Video Description Corpus (MSVD) dataset (Chen and Dolan, 2011) contains 1,970 short clips from YouTube annotated by different users, accounting for more than 80,000 training samples. In terms of coverage, approximately the 98% of the words from EDUB-SegDesc are present in the MSVD dataset.

5.3. Results

In this section, we study the influence on the final performance of different architectural proposals. First, we show the performance of a classical video caption system and study the influence of several pre-training methods. Next, we study some of the architectural choices that led to the best TMA system. Finally, we compare the state of the art with our best proposals proposal based on the influence that information from previous events has on the system.

Table 2 shows the results obtained when using different pre-training techniques for the language model. In all the experiments we tackle the problem as classical video captioning without incorporating information from previous events. The captioning model is the same as in Peris et al. (2016).

³www.keras.io

⁴<https://github.com/MarcBS/TMA>

As shown in [Table 2](#), the inclusion of word2vec vectors worsens the performance of the system in terms of BLEU and METEOR. We hypothesize that this is due to the different domains on which the word embeddings are trained. The word2vec vectors were trained on a more general domain and, therefore, their capabilities cannot be exploited to the full in our problem.

If we use the parameters learned from MSVD data, the BLEU score is also lowered, although in a lower extent than with word2vec. Nevertheless, the performance of the MSVD model in terms of METEOR and CIDEr is increased.

Table 2: EDUB-SegDesc validation and test set results for different pre-training techniques on the language model. ABiViRNet refers to the video captioning system from [Peris et al. \(2016\)](#). BLEU and METEOR metrics are given in percentage. We compare the basic model trained from scratch with the same model with certain pre-trained components. #params denotes the number of parameters to estimate, given in millions.

	Validation			Test			#params
	BLEU-4	METEOR	CIDEr	BLEU-4	METEOR	CIDEr	
ABiViRNet	31.2	21.3	0.97	29.6	20.3	0.79	27.3M
ABiViRNet + word2vec	30.1	21.0	1.04	26.0	20.1	0.90	27.3M
ABiViRNet + MSVD	32.5	22.0	1.11	28.5	21.2	0.89	35.1M

In [Table 3](#) we report a further set of comparisons that show several additional tests that did not provide good results. In the first part of the table we compare several models either using or not dropout. It appears that, combining the use of dropout, batch normalization and Gaussian noise in the same model clearly helps obtaining better results in all the cases. The sole use of dropout as regularization strategy produced even worse results. In the second part of the table we compare either using all the images available in the dataset or removing all non-informative images ([Lidon et al., 2015](#)) (i.e. images which are either dark, blurry or that point to the sky or ground without showing any object).

Table 3: EDUB-SegDesc validation and test set results for several additional configurations tested that did not provide good enough results. For each bad result we report its counter example configuration with better results. BLEU and METEOR metrics are given in percentage. #params denotes the number of parameters to estimate, given in millions. The best results for each measure are shown in boldface. Results with the symbol * were obtained with the Adam optimizer instead of Adadelta.

	Validation			Test			#params
	BLEU-4	METEOR	CIDEr	BLEU-4	METEOR	CIDEr	
ABiViRNet no dropout	29.3	21.6	1.09	24.8	18.6	0.80	27.3M
ABiViRNet	31.2	21.3	0.97	29.6	20.3	0.79	27.3M
ABiViRNet + MSVD no dropout	30.2	21.0	1.05	27.9	19.4	0.89	35.1M
ABiViRNet + MSVD	32.5	22.0	1.11	28.5	21.2	0.89	35.1M
TMA previous-caption no dropout	31.9	21.6	1.02	28.1	20.1	0.94	39.1M
TMA previous-caption	32.6	21.7	1.04	30.6	20.4	0.90	39.1M
TMA previous-video + MSVD - NonInfo*	31.7	21.7	1.13	30.3	21.5	0.99	51.0M
TMA previous-video + MSVD*	35.4	23.5	1.18	31.9	22.1	1.07	51.0M

The effect of including information from previous events is shown in [Table 4](#). For comparison, we also include the only similar approach in the egocentric video captioning literature, *DeepSeek* ([Goel and Naik, 2016](#)). This model consists of a non-attentional BLSTM encoder with two layers. The encoder feeds a similar decoder to the one from [Yao et al. \(2015\)](#). We performed additional tests with several state-of-the-art models for video captioning: *Enc-Dec Global* ([Yao et al., 2015](#)), *hLSTMAt* ([Song et al., 2017](#)) and *ABiViRNet* ([Peris et al., 2016](#)). Note that none of them considers long-term temporal information from different events.

As detailed in [Section 3.3](#), we tested our TMA model introducing the previous caption, the previous sequence of images or both previous caption and images. As before, we distinguish between training from scratch with the EDUB-SegDesc dataset or start from the weights learned with MSVD pre-trained model. Given the results observed in [Table 2](#), we drop the use of word2vec word embeddings.

Table 4: EDUB-SegDesc validation and test set results for variations of our TMA model compared to the state-of-the-art, which do not consider information from previous events. We use either the previous caption, the previous sequence of images (video) or both for making the current prediction. BLEU and METEOR metrics are given in percentage. #params denotes the number of parameters to estimate, given in millions. The best results for each measure are shown in boldface. Results with the symbol * were obtained with the Adam optimizer instead of Adadelta.

	Validation			Test			#params
	BLEU-4	METEOR	CIDEr	BLEU-4	METEOR	CIDEr	
Enc-Dec Global (Yao et al., 2015)	30.1	21.9	1.02	28.1	20.8	0.88	44.4M
hLSTMat (Song et al., 2017)	31.6	23.3	1.28	25.6	20.8	0.88	—
ABiViRNet (Peris et al., 2016)	31.2	21.3	0.97	29.6	20.3	0.79	27.3M
DeepSeek (Goel and Naik, 2016)	33.6	24.1	1.26	27.9	21.4	0.99	27.2M
TMA previous-caption	32.6	21.7	1.04	30.6	20.4	0.90	39.1M
TMA previous-caption + MSVD	33.9	23.5	1.21	28.3	21.3	0.94	46.9M
TMA previous-video*	34.4	23.7	1.21	31.0	21.4	1.01	43.3M
TMA previous-video + MSVD*	35.4	23.5	1.18	31.9	22.1	1.07	51.0M
TMA previous-video-caption	36.4	24.6	1.29	30.4	21.8	1.00	55.1M
TMA previous-video-caption + MSVD*	34.0	23.0	1.19	29.7	22.1	0.93	62.8M

According to Table 4, if we compare the results obtained by the state of the art methods (lines 1-4), which only consider the current sequence of images, to the different configurations of our method (lines 5-10), we can see that, in most of the cases, our method outperforms the rest. This behaviour can be explained by the inclusion of information from previous events in the TMA model. Since it considers a broader context, it is able to better understand the given event, and therefore, generally succeeds at increasing the performance of the system. Furthermore, considering the characteristics of egocentric lifelogging photostreams, the data provided by a single event often lacks enough information to easily understand what is happening. Thus, providing context from previous events usually improves the captioning results. On the other hand, some state of the art methods outperform certain configurations of our model (see line 3 BLEU-4, and line 4 METEOR and CIDEr on the test set). We understand that this phenomenon can occur, because although providing context from previous events can usually be useful, in some cases noise could be introduced for two reasons: 1) the error of the predictions of previous events can be propagated, and 2) some consecutive events could lack any semantical relationship or have a very low number of samples in the training set.

Considering the differences between different configurations of our TMA model, the inclusion of the previous video event as input yields the best generalization results during test evaluation. The previous caption also enhances the system, but in a lower extent. On the other hand, the inclusion of both previous frames and caption produced better results, but only in the validation set. This phenomenon could imply that although the model has a greater potential, it is also more complex considering the number of parameters, which produces a quicker over-fitting on the training data and the consequent impact produced by the model selection of the validation set (Reunanen, 2003).

The results obtained with the TMA model show that taking into account the previous video event is more effective than considering the previous caption. This effect is partially produced because, at test time, we use as input an output previously generated by the model. Obviously, this output may be erroneous. Therefore, it may lead to the introduction of noise to the system and to error propagation.

From Table 2 and Table 4, it can be concluded that, when fine-tuning from the MSVD model, the METEOR scores are always better than when starting from scratch, but BLEU-4 scores are lowered. This is probably due to vocabulary differences from both tasks. The model pre-trained tends to produce captions with the structure that it learned from the MSVD dataset, while the model trained from scratch on the EDUB-SegDesc generates ad-hoc captions for the task at hand. Since BLEU-4 is based on n-gram counts, the latter model generally obtains larger values than the first one, because its captions are more literal. Since the METEOR metric stems and employs synonyms, it is less rigid than BLEU-4. Therefore, pre-trained models are able to score better than those trained from scratch.

6. Discussion

In this section, we review some illustrative examples, in order to understand the weaknesses and strengths of our model, as well as the major difficulties appeared in the task at hand.

Fig. 8 shows some examples of the predictions produced by our model on consecutive events in the test set. We can observe the influence that the previous event had in the captioning of the second sample. In sample #2, the *TMA previous-video* model is aware that the user was previously in a restaurant. This conditions the model, which generates the caption "I went to the bathroom", which is a likely action when a person is in a restaurant. The *ABiViRNet* model is unable to capture this information.

The influence of previous actions can also be observed in events #5 and #6. At event #5, the user is shopping and in the next event, he/she continues shopping. The *TMA previous-video* model can distinguish the shopping action occurred in event #6 due to incorporation of the visual information from the previous one. On the other hand, the *ABiViRNet* model is not able to do this inference and deduces that the user is in the street. Similarly, the *TMA previous-caption* model is more prone to error propagation: when fails to understand a certain event, it is more likely to be also wrong in the following one. For example, in sample #4 the model generates "I was in a supermarket", which leads the model to fail also on the consecutive event #5.

Fig. 9 depicts some examples of initial events of different days. We can see that, in the case of TMA models, the fact that information from previous events can not be used in this particular cases does not influence the model and obtains better or comparable results than non-TMA models.

Fig. 10 shows additional examples of successes/failures of the system. In the success case #1, we can see that the model is able to understand that the user went to the bathroom, although it is not specifically described in the GT and it is hard to distinguish just looking at the images (note the third image where the user is drying his hands). In addition, this exemplifies one of the major challenges occurred within the event captioning problem: during an event, many situations may occur. The system may focus only on part of the event, resulting in captions that, although correct, are not found in the GT.

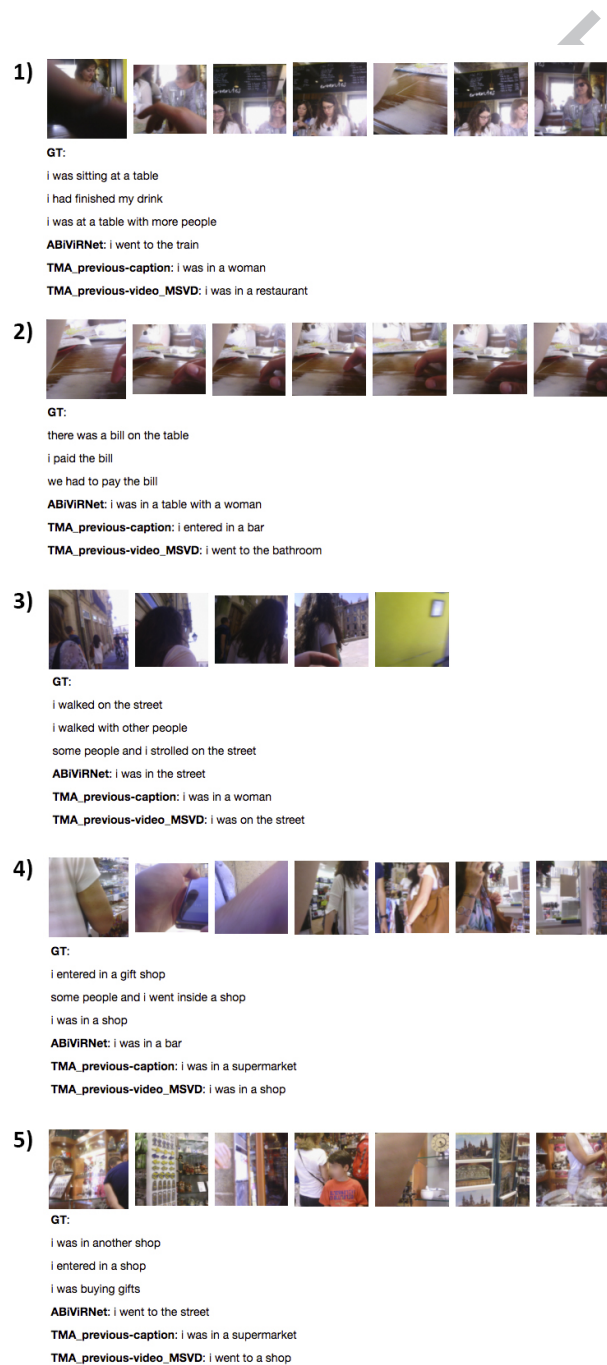


Figure 8: Comparison of the results obtained by the baseline method ABiViRNet and two of our proposals TMA previous-caption and TMA previous-video + MSVD. A set of consecutive events from the test split is shown.

In the success case #2, we can see that both *ABiViRNet* and *TMA previous-video* model correctly describe the first event. But as we move to the next one, the first model is unable to detect that the user is having lunch, while the *TMA* model is capable to infer that the event has changed. Therefore, it incorporates this information and correctly guesses that the user is having lunch.

The right examples show failure cases caused under certain conditions in the input images. In the first failure case, the model infers that the user is in a supermarket due to the vending machine in the first image. In the second failure case, the model interprets the images as a parade due to the multiple persons and colored lights appearing.

7. Conclusions and future work

In this work, we addressed the challenging problem of egocentric captioning as a video description task considering a natural characteristic of this problem: since the egocentric sequences were captured consecutively along a day, there exists a relationship between a given situation and the previous one. We aimed to include such dependency in an automatic captioning system. For doing this, we developed a natural extension of LSTM networks, able to deal with multiple inputs, even when coming from different modalities.

For assessing our proposal, we also constructed a dataset for egocentric captioning. Image sequences were obtained using a wearable camera and were manually segmented and annotated. Both source code and dataset are made publicly available. We carried out an automatic evaluation of the system, with clear results: the inclusion of previous information effectively enhances the performance of the system.

If we have available the sequence of events of a full day, we could introduce not only previous events to the system, but also the following ones. Since the *TMA* model defined in this work supports an arbitrary number of inputs, the inclusion of the context coming from following events becomes natural. Therefore, the captioning results could be refined, not only by the previous, but also by incorporating the following information. Furthermore, we could aim to look further than the immediately previous event, incorporating longer-term memory. Very recently, [Kaiser et al. \(2017\)](#) proposed a memory-augmented network, able to learn very long-term relationships. We plan to test whether such architectures can deal with our problem in an effective way.

Additionally, in a hypothetical real application, the inclusion of an interactive correction process could be considered. Interactive neural systems provide encouraging results ([Knowles and Koehn, 2016](#); [Peris et al., 2017](#)) in the machine translation field. In such systems, the user corrects the errors committed by the system, the system takes into account these corrections and changes its outputs, to produce a hopefully better caption. Finally, closely related with interactivity, another interesting extension of this work is the application of online learning techniques to the egocentric captioning pipeline. This would allow the development of better user-tailored and adaptive captioning systems, definitely more useful to the final user.

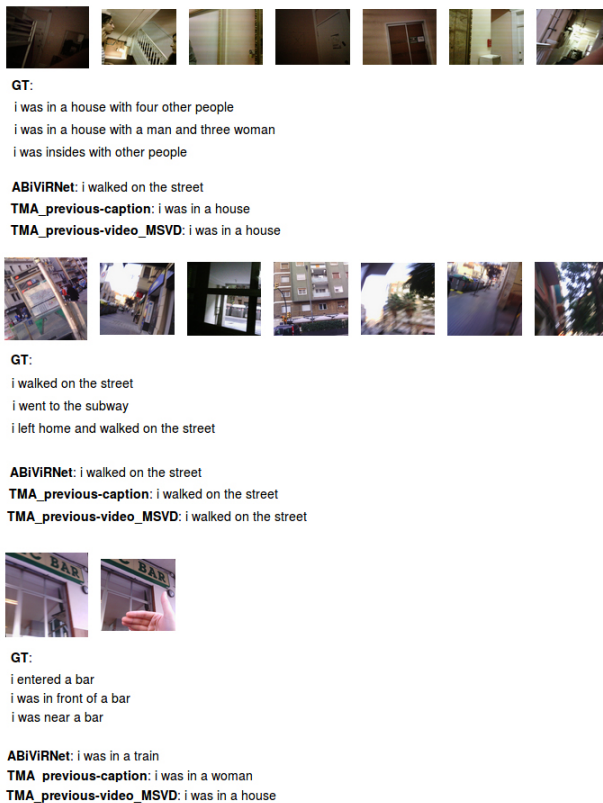


Figure 9: Comparison of the results obtained by the baseline method *ABiViRNet* and two of our proposals *TMA previous-caption* and *TMA previous-video + MSVD*. Events corresponding to the start of the day are shown.

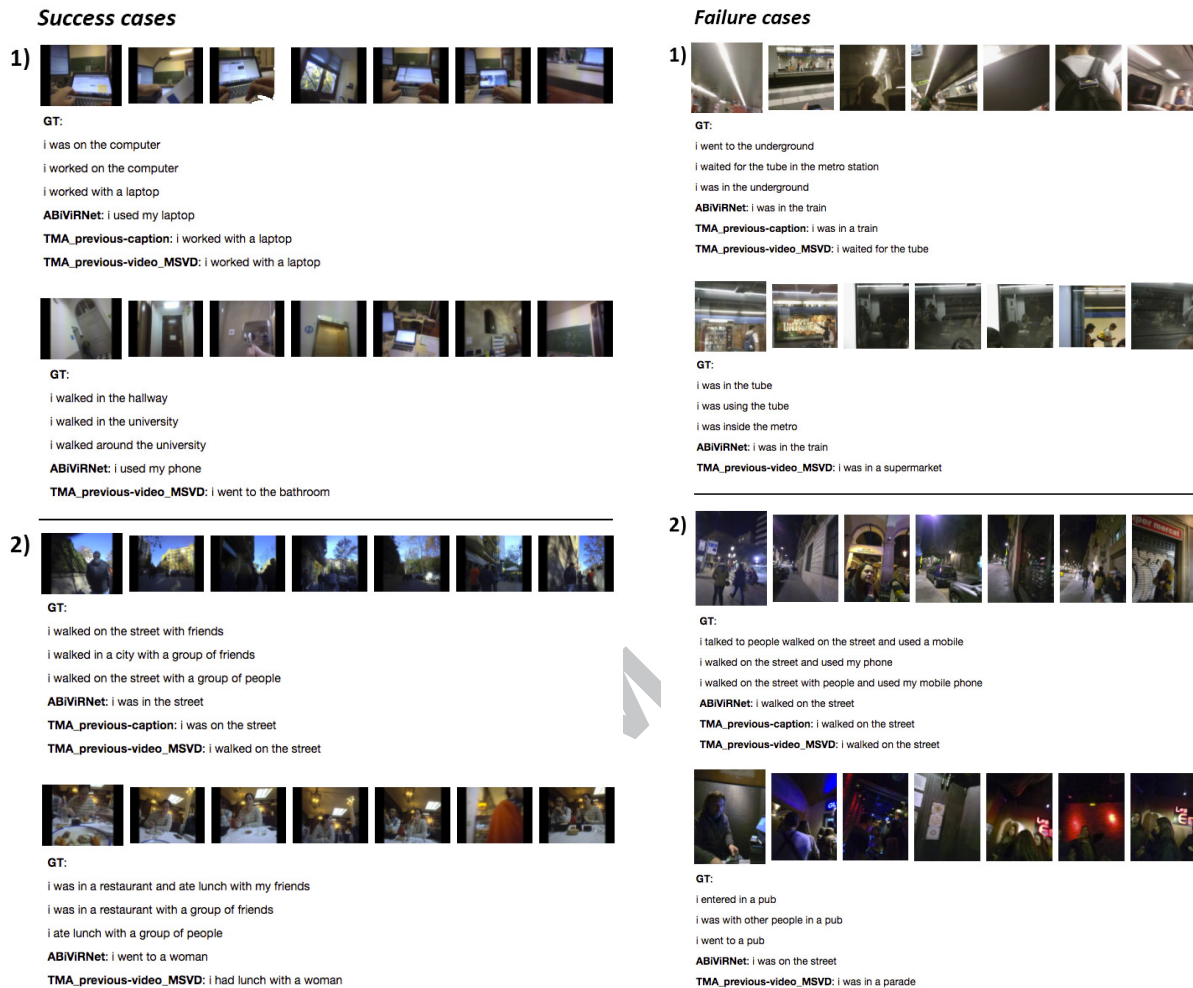


Figure 10: Two success cases (left) and two failure cases (right) of our model *TMA previous-video + MSVD*. All the examples belong to the test set. For each case, the top sequence is the previous event and the bottom sequence is the event that we are predicting. In success case #1, we can see that our method correctly recognizes that the user went to the bathroom even though it is not specified in the GT sentences and it is not straightforward to see in the images (note the third image using a hand drier). The two samples at the right exemplify cases in which our model fails due to certain images appearing in the sequences. For instance, in the failure case #1 our model considers that the user is in a supermarket due to the first image, where a vending machine appears. In failure case #2, our model considers that the user is seeing a parade due to the multiple people and colored lights appearing in the images.

Acknowledgments

This work was partially funded by TIN2015-66951-C2, SGR 1219, CERCA, Grant 20141510 (Marató TV3), PrometeoII/2014/030 and R-MIPRCV network (TIN2014-54728-REDC). Petia Radeva is partially funded by ICREA Academia'2014. Marc Bolaños is partially funded by an FPU fellowship. We gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan X GPU used for this research. The funders had no role in the study design, data collection, analysis, and preparation of the manuscript.

References

- Bahdanau, D., Cho, K., Bengio, Y., 2015. Neural machine translation by jointly learning to align and translate. arXiv:1409.0473.
- Bengio, Y., Simard, P., Frasconi, P., 1994. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks* 5 (2), 157–166.
- Betancourt, A., Morerio, P., Regazzoni, C. S., Rauterberg, M., 2015. The evolution of first person vision methods: A survey. *IEEE Transactions on Circuits and Systems for Video Technology* 25 (5), 744–760.
- Bolaños, M., Dimiccoli, M., Radeva, P., 2017. Toward storytelling from visual lifelogging: An overview. *IEEE Transactions on Human-Machine Systems* 47 (1), 77–90.
- Bolaños, M., Peris, Á., Casacuberta, F., Radeva, P., 2016. VIBIKNet: Visual bidirectional kernelized network for visual question answering. arXiv:1612.03628.
- Cartas, A., Dimiccoli, M., Radeva, P., 2017. Batch-based activity recognition from egocentric photo-streams. arXiv preprint arXiv:1708.07889.
- Castro, D., Hickson, S., Bettadapura, V., Thomaz, E., Abowd, G., Christensen, H., Essa, I., 2015. Predicting daily activities from egocentric images using deep learning. In: *proceedings of the 2015 ACM International symposium on Wearable Computers*. ACM, pp. 75–82.
- Chelba, C., Mikolov, T., Schuster, M., Ge, Q., Brants, T., Koehn, P., Robinson, T., 2013. One billion word benchmark for measuring progress in statistical language modeling. arXiv:1312.3005.
- Chen, D. L., Dolan, W. B., 2011. Collecting highly parallel data for paraphrase evaluation. In: *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. pp. 190–200.
- Chen, X., Fang, H., Lin, T.-Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C. L., 2015. Microsoft COCO captions: Data collection and evaluation server. arXiv:1504.00325.
- Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y., 2014. On the properties of neural machine translation: Encoder-decoder approaches. arXiv:1409.1259.
- Dimiccoli, M., Bolaños, M., Talavera, E., Aghaei, M., Nikolov, S. G., Radeva, P., 2017. SR-clustering: Semantic regularized clustering for egocentric photo streams segmentation. *Computer Vision and Image Understanding* 155, 55–69.
- Doherty, A. R., Hodges, S. E., King, A. C., Smeaton, A. F., Berry, E., Moulin, C. J., Lindley, S., Kelly, P., Foster, C., 2013. Wearable cameras in health. *American journal of preventive medicine* 44 (3), 320–323.
- Fan, C., Crandall, D. J., 2016. Deepdiary: Automatically captioning lifelogging image streams. In: *Proceedings of European Conference on Computer Vision*. pp. 459–473.
- Fukui, A., Park, D. H., Yang, D., Rohrbach, A., Darrell, T., Rohrbach, M., 2016. Multimodal compact bilinear pooling for visual question answering and visual grounding. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pp. 457–468.
- Gers, F. A., Schmidhuber, J., Cummins, F., 2000. Learning to forget: Continual prediction with LSTM. *Neural computation* 12 (10), 2451–2471.
- Goel, K., Naik, J., 2016. Deepseek: A video captioning tool for making videos searchable.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural computation* 9 (8), 1735–1780.
- Hori, C., Hori, T., Lee, T.-Y., Sumi, K., Hershey, J. R., Marks, T. K., 2017. Attention-based multimodal fusion for video description. arXiv:1701.03126.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Proceedings of the International Conference on Machine Learning* 32, 448–456.
- Iwashita, Y., Takamine, A., Kurazume, R., Ryoo, M. S., 2014. First-person animal activity recognition from egocentric videos. In: *Pattern Recognition (ICPR), 2014 22nd International Conference on*. IEEE, pp. 4310–4315.
- Kaiser, L., Nachum, O., Roy, A., Bengio, S., 2017. Learning to remember rare events. In: arXiv:1703.03129.
- Kingma, D., Ba, J., 2014. Adam: A method for stochastic optimization. arXiv:1412.6980.
- Knowles, R., Koehn, P., 2016. Neural interactive translation prediction. In: *Proceedings of the Association for Machine Translation in the Americas*. pp. 107–120.
- Krishna, R., Hata, K., Ren, F., Fei-Fei, L., Niebles, J. C., 2017. Dense-captioning events in videos. arXiv preprint arXiv:1705.00754.
- Krizhevsky, A., Sutskever, I., Hinton, G. E., 2012. Imagenet classification with deep convolutional neural networks. In: *Proceedings of the Neural Information Processing Systems Conference*. pp. 1097–1105.
- Lavie, A., Denkowski, M. J., 2009. The METEOR metric for automatic evaluation of machine translation. *Machine translation* 23 (2-3), 105–115.
- Lidon, A., Bolaños, M., Dimiccoli, M., Radeva, P., Garolera, M., Giró-i Nieto, X., 2015. Semantic summarization of egocentric photo stream events. arXiv:1511.00438.
- Lu, Z., Grauman, K., 2013. Story-driven summarization for egocentric video. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2714–2721.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., Dean, J., 2013. Distributed representations of words and phrases and their compositionality. In: *Proceedings of the Neural Information Processing Systems Conference*. pp. 3111–3119.
- Pan, P., Xu, Z., Yang, Y., Wu, F., Zhuang, Y., 2016. Hierarchical recurrent neural encoder for video representation with application to captioning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1029–1038.

- Papineni, K., Roukos, S., Ward, T., Zhu, W.-J., 2002. BLEU: a method for automatic evaluation of machine translation. In: Proceedings of the Annual Meeting of the Association for Computational Linguistics. pp. 311–318.
- Pascanu, R., Mikolov, T., Bengio, Y., 2013. On the difficulty of training recurrent neural networks. Proceedings of the International Conference on Machine Learning 28, 1310–1318.
- Peris, Á., Bolaños, M., Radeva, P., Casacuberta, F., 2016. Video description using bidirectional recurrent neural networks. In: Proceedings of the International Conference on Artificial Neural Networks. pp. 3–11.
- Peris, Á., Domingo, M., Casacuberta, F., 2017. Interactive neural machine translation. Computer Speech & Language 45, 201–220.
- Poleg, Y., Ephrat, A., Peleg, S., Arora, C., 2016. Compact cnn for indexing egocentric videos. In: IEEE Winter Conference on Applications of Computer Vision. IEEE, pp. 1–9.
- Reunanen, J., 2003. Overfitting in making comparisons between variable selection methods. Journal of Machine Learning Research 3 (Mar), 1371–1382.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al., 2014. Imagenet large scale visual recognition challenge, 2014. arXiv:1409.0575.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al., 2015. Imagenet large scale visual recognition challenge. International Journal of Computer Vision 115 (3), 211–252.
- Schuster, M., Paliwal, K. K., 1997. Bidirectional recurrent neural networks. IEEE Transactions on Signal Processing 45 (11), 2673–2681.
- Sellen, A. J., Fogg, A., Aitken, M., Hodges, S., Rother, C., Wood, K., 2007. Do life-logging technologies support memory for the past?: an experimental study using sensecam. In: Proceedings of the SIGCHI conference on Human factors in computing systems. ACM, pp. 81–90.
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. arXiv:1409.1556.
- Song, J., Guo, Z., Gao, L., Liu, W., Zhang, D., Shen, H. T., 2017. Hierarchical lstm with adjusted temporal attention for video captioning. arXiv preprint arXiv:1706.01231.
- Spector, A., Thorgrimsen, L., Woods, B., Royan, L., Davies, S., Butterworth, M., Orrell, M., 2003. Efficacy of an evidence-based cognitive stimulation therapy programme for people with dementia. The British Journal of Psychiatry 183 (3), 248–254.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. Journal of Machine Learning Research 15 (1), 1929–1958.
- Sutskever, I., Vinyals, O., Le, Q. V., 2014. Sequence to sequence learning with neural networks. In: Proceedings of the Neural Information Processing Systems Conference. Vol. 27. pp. 3104–3112.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1–9.
- Theano Development Team, 2016. Theano: A Python framework for fast computation of mathematical expressions.
- Toselli, A. H., Vidal, E., Casacuberta, F., 2011. Multimodal interactive pattern recognition and applications. Springer Science & Business Media.
- Vedantam, R., Lawrence Zitnick, C., Parikh, D., 2015. Cider: Consensus-based image description evaluation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4566–4575.
- Venugopalan, S., Hendricks, L. A., Mooney, R., Saenko, K., 2016. Improving LSTM-based video description with linguistic knowledge mined from text. arXiv:1604.01729.
- Venugopalan, S., Rohrbach, M., Donahue, J., Mooney, R., Darrell, T., Saenko, K., 2015. Sequence to sequence-video to text. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 4534–4542.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R., Bengio, Y., 2015. Show, attend and tell: Neural image caption generation with visual attention. In: Proceedings of the International Conference on Machine Learning. pp. 2048–2057.
- Yao, L., Torabi, A., Cho, K., Ballas, N., Pal, C., Larochelle, H., Courville, A., 2015. Describing videos by exploiting temporal structure. In: Proceedings of the International Conference on Computer Vision. pp. 4507–4515.
- Zeiler, M. D., 2012. Adadelta: an adaptive learning rate method. arXiv:1212.5701.

- We present a novel captioning model, which incorporates the information from previous events into the current decoding state.
- We present a new LSTM model capable of combining information from multiple inputs and modalities, as well as applying a separate attention mechanism to each of them.
- We conduct experiments on the new dataset and compare our model with other classical captioning architectures. Results show that using information from previous events provides better generalization.
- We present the first dataset for egocentric sequences captioning, named EDUB-SegDesc, based on describing the events that take place along a day.
- We make public both dataset and model, in order to make results reproducible and foster the research in this topic.