

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# Guiding Attention using Partial-Order Relationships for Image Captioning

Murad Popattia<sup>1</sup>, Muhammad Rafi<sup>1</sup>, Rizwan Qureshi<sup>1,2</sup>, Shah Nawaz<sup>3†</sup> <sup>1</sup>National University of Computer and Emerging Sciences, Karachi, Pakistan, <sup>2</sup>Hamad Bin Khalifa University, Doha, Qatar

<sup>3</sup>Pattern Analysis & Computer Vision (PAVIS) - Istituto Italiano di Tecnologia (IIT)

muradmansoor189@gmail.com, muhammad.rafi@nu.edu.pk, riahmed@hbku.edu.qa,

shah.nawaz@iit.it

#### Abstract

The use of attention models for automated image captioning has enabled many systems to produce accurate and meaningful descriptions for images. Over the years, many novel approaches have been proposed to enhance the attention process using different feature representations. In this paper, we extend this approach by creating a guided attention network mechanism, that exploits the relationship between the visual scene and text-descriptions using spatial features from the image, high-level information from the topics, and temporal context from caption generation, which are embedded together in an ordered embedding space. A pairwise ranking objective is used for training this embedding space which allows similar images, topics and captions in the shared semantic space to maintain a partial order in the visual-semantic hierarchy and hence, helps the model to produce more visually accurate captions. The experimental results based on MSCOCO dataset shows the competitiveness of our approach, with many state-of-the-art models on various evaluation metrics.

# 1. Introduction

Recent success of deep neural networks in computer vision, speech, and natural language processing have prompted academics to think beyond these fields as separate entities, instead solving challenges at their intersections [2, 6, 9, 17, 19, 33]. Generating descriptive and meaningful captions for images, and to capture its semantic meaning, is one such multimodal inference problem [3, 10]. Despite its complexity, it has various applications, including visually- impaired assistance, intelligent chat-bots, medical report generation, self- driving cars, and many more [23]. In general, an image captioning model should be able to



GT: A double decker bus sitting on a green field ATT: A bus that is sitting in the grass T-OE-ATT: A double decker bus is parked in a grassy field



GT: An image of a herd of giraffes in the water ATT: A group of giraffes standing on a dirt road T-OE-ATT: Three giraffes are drinking

water from a watering hole



find objects, their positions, map the relationship, as well as express this relationships in a human understandable language.

A typical image caption system consists of a convolutional neural network (CNN) and a recurrent neural network (RNN), with CNN as the image encoder and RNN as the sentence decoder [21, 26]. However, in order to capture the spatial context from the image in an efficient manner, other approaches such as [8, 31, 32, 34] incorporate highlevel information from topics or detected objects as semantic features to the decoder model. Another line of research was to make use of cross-modal associations between image and text features in a joint-embedding space. Earlier research work [13,15] treated images and caption as a symmetric relationship by using Euclidean or cosine distances to gauge similarities between these two modalities. On the other hand, in [25] treated these associations as asymmetric by enforcing a hierarchical order within the embedding space, and has shown to perform better than symmetric relationships.

Further improvement in this framework, is the introduc-

<sup>†</sup> Current Affiliation: Deutsches Elektronen-Synchrotron (DESY) Email: shah.nawaz@desy.de

tion of attention mechanism [22], which allows the decoder to focus on a sub-region of the image, when predicting the next word in the caption [28]. Despite of focusing only on spatial attention, Lu et al. [16] presented a novel adaptive mechanism for helping the attention module to learn, when to shift between spatial and temporal context during word prediction. In addition, Anderson et al. [1] improves the attention process by first detecting a set of salient image regions (bottom-up) and then attending to these fixated regions (top-down). [30] builds upon this concept by exploiting the semantic relationships between the detected spatial regions using GCN (Graph Convolution Networks). [11] also make use of a similar approach but instead modify the attention module by adding self-attention module on top of the conventional attention mechanism, which helps the decoder to draw relations between various attended vectors.

On the other hand, Jiang et al. [12], focused on increasing the semantic information fed to the decoder by using a fusion of multiple encoders, each focusing on a different view point, to build better representations for the decoder. Likewise, Wang et al. [27] also worked in a similar direction that guides attention using a hierarchy of semantic features. However, lack of inter-feature correlations between these encoders makes it difficult for the decoder to leverage the association from the resulting joint representations. Lastly, despite relying on spatial cues from encoded features, Ke et al. [14] worked on improving the temporal coherence of words during descriptions by applying attention on both visual and textual domains.

Alongside the same line of work of incorporating semantic associations between different spatial regions using GCNs [30], our idea is to make use of multi-modal representations such as ordered embeddings [25] as our semantic feature vectors to guide the attention module. Similar to the late-fusion of features as done in [30], we instead use a weighted summation as our fusion mechanism to fuse these embeddings.

Overall the main contributions of our work are three-fold:

- We make use of ordered embedding features for topics and images to guide the attention module instead of feeding them as low-level features. This step has been shown to improved metrics, see ablation study (*Section 3.3.1*).
- We incorporate a global weighted sum for fusing "visual" and "temporal" states instead of feeding them at each time-step separately which helps the model to learn the best estimation of the attention required for each image.
- Lastly, we present an ablation study of each contribution and how it effects the overall performance of the model on the MSCOCO dataset.

## 2. Methodology

#### 2.1. Overall Framework

Our approach follows the traditional encoder-decoder framework, where the encoder is responsible to pass on features used by the decoder to output the most likely word during captioning. Figure 2 illustrates the overall framework.

Similar to recent approaches of sending objects or topics during encoding [29, 34], we used topics instead of objects to capture both the "actors" as well the "activities" binding them. The encoder consists of three components: 1) topic classifier 2) feature extractor and the 3) retrieval model. We use a pre-trained deep CNN model as a feature extractor to extract visual features from the image and train a multi-label topic classifier to predict topics for given images. After that, we train a retrieval model which embeds captions, image and topics into a shared semantic space, in which a similarity score can be calculated between them. Interestingly, using embeddings helps to better learn the latent relationships between image and topic features, lost during feature extraction. This helps the attention module in describing and discriminating spatial regions more effectively. (Details in Section 2.3)

Inspired from the simple yet effective architecture defined in [1], we used two LSTM branches in the decoder i.e. the *guiding-lstm* and the *core-lstm*. Here, we define a weighted sum of the semantic embeddings of both the images and topics, as input to the *guiding-lstm* at the first timestep, which gives the model a better understanding of the alignment of visual features and the topics. We then utilize its hidden state  $h_{t-1}$  for guiding the language LSTM and the context vector  $z_t$  used for attention. Similar to using a visual sentinel [16], we used a weighted summation for fusing the attention weights instead of a sentinel gate to shift between spatial and temporal attentions. This allows for a more simpler architecture in terms of learning parameters involved, whilst maintaining the accuracy during word prediction.

#### 2.2. Topic Classifier

For extracting topics *T*, the ground-truth captions are concatenated to form documents *D*, where each document *d* corresponds to captions *C*, for a given image and contains a set of words *W*. After that, we train a Latent Dirichlet Allocation (LDA) model [4], which is a probabilistic model to learn the topic representations from documents. The trained topic model outputs a set of topic probabilities  $T = \{T_1, T_2, T_3, ..., T_n\}$ .

For training the classifier, the topics are sampled and converted to one-hot representations using the following function:



Figure 2. The overall framework of the proposed model, where  $\Sigma_w$  represents a weighted-summation and  $\oplus$  denotes matrix addition. The model consists of a feature extractor and a topic classifier to extract spatial features and topics given in an image. These semantic attributes are then fed into a retrieval model which arranges an image, topic and caption triplet in a partial-order hierarchy. The resultant embeddings are then late-fused using weighted summation and then fed into a 'guiding LSTM'. The 'core-lstm' then makes use of this hidden state for temporal attention. Consequently, two separate attention blocks are used, each attending to different aspects of decoding and the resulting attention vectors are late-fused again in a weighted fashion to produce captions.

$$f_{t_i \subseteq T_i}(x) = \begin{cases} 1 & \text{if } P(x) \ge 0.1\\ 0 & \text{else} \end{cases}$$
(1)

where  $t_i$  represents a single topic from a set of topics T, for image i from a set of images I and P(x) represents the topic-confidence from LDA. We formulate this as a multilabel classification problem, since an image can have multiple topics. A pre-trained CNN model is used to extract image features which are then fed into a feed-forward neural network with a sigmoid activation for the prediction layer. This layer outputs an  $(N_i \times N_t)$  vector where  $N_i$  corresponds to the number of images and  $T_t$  are the number of topics. We report the evaluation for the topic classifier in *Section 4.1* of the paper.

#### 2.3. Retrieval Model

The architecture of the retrieval model is inspired by the approaches in [25, 32]. It follows the idea of [13] to align caption and images in the same space, but with a partialorder relation rather than a symmetric relation. This is a more intuitive approach as images have captions with different levels of details, and because the captions are so dissimilar, it is impossible to map both their embeddings close to the same image embedding using a symmetric distance measure like cosine similarity. Nevertheless, maintaining order is robust to such affect, as dissimilar caption can have embeddings placed very far away from the image, while remaining above it in the partial order. The partial order relation can be defined as:

 $x \leq y \iff \forall x \forall y (x \geq y)$ . This imposes for all values of the vector x to be greater than all values of the vector y in the embedding space to maintain order.

We start with three entities i.e. images *I*, topics *T* and captions *C*. As per [25], we utilized domain-specific encoders to extract features for training the embeddings. For images and topics, we utilized the fully-connected features from the feature-extractor and the topic features from the topic classifier respectively. While for captions, we used a Gated Relu Unit (GRU) as the RNN based text-encoder instead of an LSTM, because of its computational efficiency. These feature vectors are then weighted with  $W_I$ ,  $W_T$  and  $W_C$  before being projected in the embedding space:

$$O_{\rm i} = \|W_{\rm I} \cdot f_{\rm FE}(I)\|^2 \tag{2}$$

$$O_{\rm t} = \|W_{\rm T} \cdot f_{\rm TC}(T)\|^2 \tag{3}$$

$$O_{\rm c} = \|W_{\rm C} \cdot GRU(C)\|^2 \tag{4}$$

 $O_i$ ,  $O_c$  represents the order embeddings of image, topics, and captions respectively.  $f_{FE}(I)$  represents the image features from the feature-extractor, while  $f_{TC}(T)$  represents the features from the topics classifier. We use L2-Norm during encoding instead of an absolute value function to mitigate overfitting [25].

Similarity Function The general notion of similarity between two vectors x and y in the embedding space can hence be quantified as the degree to which a pair of points violates the partial order  $x \leq y$  [25]:

$$S(x,y) = -(\|max(0,O_{y} - O_{x})\|^{2})$$
(5)

where  $O_x$  and  $O_y$  represents the encoded feature vector in the embedding space. The negative sign constitutes to the fact that a positive difference between  $O_y$  and  $O_x$  denotes violation of the order penalty.

Loss Function As previous works which learn embedding in cross-modal retrieval tasks [13, 15], we re-use the pair-wise ranking loss objective to increase the similarity for the matching pairs and vice-versa for the contrastive terms by a margin  $\alpha$ :

$$L(x,y) = \sum_{(x,y)} (\sum_{x'} \max\{0, \alpha - S(x,y) + S(x',y)\} + \sum_{y'} \max\{0, \alpha - S(x,y) + S(x,y')\})$$
(6)

where (x, y) is the ground-truth pair while (x', y) and (x, y') are constrastive terms. Our hierarchy has image at the top of the partial order, followed by captions which are then bounded by the topics. Hence, the total loss can be defined as the summation of losses over all three partial orders:

$$L = L(I, C) + L(I, T) + L(C, T)$$
(7)

#### 2.4. Caption Generation

We now describe the decoding phase of the model. The trained encoding functions  $O_i$  and  $O_t$  are used to produce relevant embeddings for image and topics during feature extraction. We then used a weighted-summation ( $\Sigma_w$ ), of these embeddings:

$$\Sigma_{w(OE)} = \lambda \cdot O_i + (1 - \lambda) \cdot O_t \tag{8}$$

where  $\lambda$  is a learnable parameter. The reason for a weighted-sum is to allow the model to learn the relative importance of each embedding during training. Different from the approach of [32], we focused on guiding the decoder in a 3-way manner i.e. using the embedding information, visual features and reliance on past information from the hidden states.

**Dual-LSTM branch** We used an auxiliary *guiding-lstm*, to process the information from the learned embeddings and

feeding the hidden state information to both the *attention* vector  $z_t$  and the *core-lstm* at initial timestep t = -1:

$$h_{t-1} = LSTM_g(\Sigma_{w(OE)}) \tag{9}$$

$$z_t = W_g h_{t-1} + W_c h_t (10)$$

$$h_t = LSTM_c(x_t, h_{t-1}) \tag{11}$$

where  $h_{t-1}$  and  $h_t$  represents the hidden states at relevant timesteps,  $W_g$  and  $W_c$  are learnable parameters in the context vector  $z_t$ .  $LSTM_g$  and  $LSTM_c$  represent the guiding and core LSTMs respectively. The initial hidden state for  $LSTM_g$  is essentially zeroes and hence not shown in the formulation.

**Spatial Attention Block** This block is responsible to generate the attention distribution vector over the important visual regions of the image. Similar to the idea of *soft-attention* [16], we utilize the context-vector  $z_t$  from equation 10 instead of just the hidden state information done in [16], in order to guide attention pertaining to the partial-order relation between the image and topic:

$$\alpha_t = softmax(W_{\alpha}[W_f F_L + W_z z_t]) \tag{12}$$

$$\rho_s = \sum_{i=1}^{N} \alpha_{ti} f_i \tag{13}$$

where  $F_L = \{f_1, f_2, ..., f_N\}$  represent the local image features from the convolution layer just before the FC layer of the feature extractor,  $\alpha_t$  denotes the attention weights over the features in  $F_L$ ,  $\alpha_{ti}$  denotes the weight over the i<sup>th</sup> part of  $F_L$  and  $\rho_s$  denotes the *spatial-context* vector.

ŀ

**Temporal Attention Block** The temporal block guides the attention module whether the information is required at all, or the next word can be predicted using the past information stored within the decoder [16]. Likewise, we utilize the information from the LSTM's memory cell along with the context vector  $z_t$  which contains the residual embedding information from the previous timestep. It helps the temporal block decide whether the current timestep requires attending to visual features or not. This is illustrated below:

$$\rho_t = \tanh(c_t) \bigodot \sigma(W_x x_t + W_{z'} z_t)$$
(14)

where  $c_t$  is the memory cell of the *core-lstm*,  $x_t$  is the word vector at timestep t,  $z_t$  denotes the context vector,  $\bigcirc$  refers to an element-wise product and  $\rho_t$  denotes the *temporal-context* vector.

**Word Prediction** Instead of keeping track of the temporal information for each word, we let the model generalize the ratio between these attentions using a weighted-summation ( $\Sigma_w$ ). This is because ideally it is a more

simpler approach to rely less on the attention gate at each timestep and generalize from the embedding context obtained from  $z_t$ .

$$\Sigma_{w(ATT)} = \mu \cdot \rho_s + (1 - \mu) \cdot \rho_t \tag{15}$$

We then calculate the word probability over a vocabulary of possible words at time t:

$$p_t = softmax(f_{MLP}(z_t + \Sigma_{w(ATT)}))$$
(16)

where  $f_{MLP}$  denotes a dense layer with ReLU activation.

#### **3. Experiments**

#### **3.1. Implementation Details**

As our model is divided into sub-components, we train each part separately instead of training them end-to-end.

Feature Extractor We use a ResNet-152 [7] model trained on ImageNet dataset. The FC features are taken from the last layer of the CNN which have a dimension of  $2048 \times 1$ . We use  $F_L = \{f_1, f_2, ..., f_N\}, f_i \in \mathbb{R}^{512}$  to represent the spatial CNN features at each of the N grid locations where N = 49.

**Topic Classifier** For the training the topic model, we limit our vocabulary to top 5000 and train the LDA on these features for 100 iterations. We empirically set the number of topics to be 80 for our case. Increasing the topics made the topic vectors more sparse and decreased the recall for the topic classifier. For the topic classifier, we used the image features  $R^{2048\times1}$  to be fed into a 5-layer feed-forward NN, with the prediction layer  $R^{80}$  having a sigmoid activation. The classifier was optimized using SGD with a learning-rate of 0.1 and momentum 0.9. The learning-rate was changed in case of plateauing with a patience of 0.4 and a factor of 0.2.

**Retrieval Model** For the retrieval model, we reused the FC image features  $R^{2048\times1}$  from the feature extractor and the  $R^{80}$  topic features from the topic classifier in *Section* 2.2. The dimensions of the embedding space and the GRU hidden state in equation (4) were set to 1024, and the margin  $\alpha$  is set to 0.05 as per [25].

**Caption Model** For the decoder, our model used LSTMs. The *guiding* and *core* LSTMs both have a dimension of 512. For the captions, we use a word embedding size of 256. During training, we see that downsizing and concatenating FC image features with this embedding improved results. The initial value for  $\lambda$  and  $\mu$  in equation (8) is set to 0.5 for both, and learned during training. Furthermore, the number of units for  $f_{MLP}$  was set to 1024. Lastly, for sampling the captions, we use beam size of 1. The whole model was optimized using Adam optimizer with a minibatch size of 128 and learning rate of 0.001. The model trained for 10 epochs on a Tesla T4 GPU and the training finished in 10 hours to produce the results.

#### **3.2.** Datasets

We conducted experiments on the popular benchmark: Microsoft COCO dataset <sup>1</sup> as this has been widely used for benchmarking in the related literature. Also, we adopt the 'Karpathy' splits setting [14], which includes 118,287 training images, and 5K testing images for evaluation. Some images had more than 5 corresponding captions, the excess of which are discarded for consistency. We directly use the publicly available code <sup>2</sup> provided by Microsoft for result evaluation, which includes BLEU, METEOR, ROUGE-L and CIDEr.

#### 3.3. Evaluation

#### 3.3.1 Ablation Study

To study the effects of guiding the attention module, we design an ablation experiment to assess the effect of 1) using an embedding space 2) using a different feature extractor and 3) using embedding along with attention as shown in **Table 2**. We see that the initial approach of feeding topics as low-level features performs poorly. A dramatic improvement was seen when using an embedding space in the process. This confirms the hypothesis that embeddings serve as a better auxiliary guidance for attention. We term this as (**T-OE**). Moreover, we assess the model's performance on a less accurate feature extractor such as VGG-19 [20] which only incurred as small change in the metrics signifying that trained embeddings are robust to changes in the feature extractor. Lastly, we incorporate attention in the process (T-OE-ATT) and guide them using the trained embeddings which shows an improved score in all metrics, signifying the importance of the embeddings to guide attention.

#### **3.3.2** Quantitative Evaluation

In **Table 1**, we compare our proposed architecture with recent state-of-the-art models on the MSCOCO dataset that make use of LSTMs in their decoder architecture. For fair comparison, we report the scores for single model for each approach that use the same CNN backbone as ours (ResNet [7]), without using ensembling and CIDEr optimizations.

Our approach is able to outperform RF-Net [12] and HAN [27] signifying that using partial order is more suitable for building joint multi-modal representations as compared to using domain-specific encoders alone. Moreover, incorporating attention with T-OE, as shown in Table 2, also helps us outperform RDN [14] over notable metrics such as **METEOR, ROUGE-L** and **CIDEr** which show that orthogonal improvements to encoder or decoder alone are less susceptible to improvement as compared to jointly improving both the feature representations and the caption genera-

<sup>&</sup>lt;sup>1</sup>https://cocodataset.org/

<sup>&</sup>lt;sup>2</sup>https://github.com/tylin/coco-caption

Approaches	BLEU-I	<b>BLEU-II</b>	<b>BLEU-III</b>	<b>BLEU-IV</b>	METEOR	<b>ROUGE-L</b>	CIDEr
Adaptive ATT [16]	74.2	58.0	43.9	33.2	26.6	-	108.5
LSTM-A [31]	75.4	-	-	35.2	26.9	55.8	108.8
<b>RF-Net</b> [12]	76.4	60.4	46.6	35.8	27.4	56.5	112.5
Up-Down ATT [1]	77.2	-	-	36.2	27.0	56.4	113.5
HAN [27]	77.2	61.2	47.7	36.2	27.5	56.6	114.8
<b>RDN</b> [14]	77.5	61.8	47.9	36.8	27.2	56.8	115.3
GCN-LSTM [30]	77.4	-	-	37.1	28.1	57.2	117.1
AoA-Net [11]	77.4	-	-	37.2	28.4	57.5	119.8
Ours (T-OE-ATT)	77.0	61.2	47.1	35.9	28.4	57.3	115.9

Table 1. Performance comparison on MSCOCO 'Karpathy' test split trained on a single-model using cross-entropy loss without CIDEr optimization. (-) indicates metrics not provided. All values are provided in percentages (%) with the highest bold-faced.

Approach	B-IV	METEOR	ROUGE-L	CIDEr
Topic	25.5	22.9	50.1	80.2
T-OE <sub>(VGG)</sub>	34.4	27.8	56.5	112.7
T-OE(Resnet)	35.4	28.2	57.0	114.4
T-OE-ATT	35.9	28.4	57.3	115.9

Table 2. Ablation study on MSCOCO 'Karpathy' test split.

tion process. It is worth noting that compared to our architecture, RDN [14] and RF-Net [12] have a greater number of learning parameters (1.15B parameters [14] for RDN), whilst our decoder contains comprises of only 29M parameters and yet is able to produce competitive results. Both GCN-LSTM [30] and AoA-Net [11] use Faster-RCNN as their feature encoder which is able to feed in region-level information while our model uses only the fully connected features from the ResNet backbone and is still competitive over METEOR and ROUGE-L scores. It should also be noted that AoA-Net [11] leverage the use of self-attention mechanisms which have been used alongside transformers and are able to produce state-of-the-art results. On the contrary, our work can be extended to incorporate region level information alongside topics or to use a different attention mechanism to improve results and has not been explored in this study.

As our model uses LSTMs for caption generation, hence this comparison does not take into account transformerbased architectures [5, 24]. Transformers are a different class of architecture as compared to LSTMs as they do not follow the auto-regressive nature of LSTMs and process the inputs in a parallel fashion [24] so incorporating partialorder embeddings alongside this class of architecture could also be a favourable research direction.

#### 3.3.3 Qualitative Evaluation

We assess our model qualitatively as illustrated in Figure 1. The baseline model is based on model's output based on *topic and image features*, while the guided attention model is based on *topic and image embeddings*. Without embeddings, we see that attention lack descriptiveness of the context associated with the visual features such as **doubledecker**, grassy, drinking etc.



GT: A group of people crossing the curved street
T-OE: A bus driving down a street next to a truck
T-OE-ATT: A truck driving down a street next to a traffic light



**GT:** A beautiful woman cooking food in the kitchen **ATT:** A woman is frying in the wok

with a knife

T-OE-ATT: A person in a kitchen preparing food with a knife

Figure 3. Examples of inaccurate captions from the model.

We also see an influence when comparing ground truth captions where the model was able to capture semantic context like *parked* instead of *sitting* and *drinking water* instead of *in the water*. It is because the model is able to draw associations between objects and actions due to partial-order information from the underlying topics of the captions fed into the decoder module denoting how attention was guided.

However, as denoted in Figure 3, the attention module can pick up on noise from these embedded features such as confusing between *a bus* and *a truck*. This is evident from **T-OE**, where the caption contains **truck** even though it is absent from the image. An explanation can be *bus* and *truck* being semantically closer in the embedding space. Moreover, relying on spatial attention can also lead to misclassifying objects in the image from *spatula* to *knife*. This can be seen from the caption generated from the model without **T-OE** where the object is misidentified as a *knife*.

# 4. Discussion

# 4.1. Evaluation of topic classifier and retrieval model

As the topic classifier and the embedding sub-space act as intermediaries to the final model, hence, we evaluate their performance on relative metrics as well. The output of the topic classifier is a set of sigmoid probabilities which are converted into one-hot encodings. Using precision solely for evaluating one-hot encodings is not enough as we can see that a higher precision does not mean our model has a good recall. Hence, we use F1-score with a  $\beta$  more inclined towards recall. The highest F1-score was achieved in the COCO dataset which may be due to a larger amount of data being used to train the model. We summarize these results in **Table 3**.

Dataset	Precision	Recall	F1-Score
Flickr30k	60.08	42.33	43.56
MSCOCO	77.54	60.48	61.52

Table 3. Performance results of the topic classifier on validation sets of Flick30k and MSCOCO dataset

For the order embedding model, we assess the quality of the model by treating it as a Caption Retrieval task. The metric used in this experiment was Recall@K which refers to the percentage of recall achieved in top-k items. We summarize these results in **Table 4**.

Dataset	R@1	R@5	R@10
Flickr30k	35.2	61.9	73.4
MSCOCO	49.5	79.6	89.3

Table 4. Performance results of the retrieval model on validation sets of Flick30k and MSCOCO dataset

Nevertheless, the scores for both the topic classifier and the retrieval model were not state-of-the-art but were enough to extract suitable features for the training images. Respective improvements to the models in terms of finetuning or using a different architecture, might positively impact the overall accuracy during captioning but is beyond the scope of this paper.

## 4.2. Visualizing the embedding space

In this section, we present a high-level visualization of the partial-order structure between images, topics and captions in the embedding space, as shown in Figure 4.

The embedding space consist of three modalities, with images being at the highest order, captions being at the center and topics being at the lowest order of the hierarchy posing a lower bound for the captions. This hierarchical arrangement also conforms with the cognitive arrangement of these modalities. Images are generally abstract points from



Figure 4. Representation of order in the embedding space

which we derive meaning about its context while separate words such as topics can be used to complement images but do not contribute to any meaning on their own. Captions on the other hand, describe a story which the spatial cues of the image support.

We can then visualize these captions as a collection of words each of which can constitute to a topic. Treating the problem as a caption retrieval task, where given an image the model outputs the set of all possible captions, setting a lower bound with topics helps constraint this search space and helps reduce noise from overlapping caption regions. [32]

# 4.3. Analysis of the weighted summation for attention

Contrasting to the approach followed in [16], where the model is trained to shift attention at each word prediction step, we constraint the model in determining an overall ratio of the spatial or temporal attention needed for word prediction and keep this as a static value for all succeeding predictions. However, despite setting the values randomly, we allow the decoder to generalize from a set of caption on the amount of attention needed for each caption. For testing, we set the temporal context  $\mu$  to be 0.3 for spatial and consequently, 0.7 for temporal attention. The reason for a higher ratio for temporal context is because it complements the RNNs capability to work with sequences. For the model we use the **ATT** approach where the image features are fed directly as spatial cues to the decoder. **Figure 5** shows the learned ratios after several iteration of training.

It can be seen that the model gradually learns to increase the gradient flow from the spatial block of the attention module, signifying the need of visual attention. However, we do notice some peaks for the flow of temporal informa-



Figure 5. Weight distributions of spatial and temporal attention for several iterations on MSCOCO dataset.

tion. A plausible reason is because while visual information is necessary, it may not always be inline with temporal coherence when describing images. Hence, we sample captions with different values for  $\mu$  as shown in Figure 6.

For a lower value of  $\mu$  in the formulation 15 in Section 2.4, the model allows the flow of temporal information in the decoder and hence we see a time-relative aspect in sentences with phrases such as "about to hit" and "is laying on". On the contrary, if we shift the value of  $\mu$ higher, it boosts the gradient flow from the spatial block filling in visual details from the image such as "white shirt", "white shorts", "laptop computer", "top of desk". However, we see that despite being rich in scene-specific details, the model misses out the global context of the image imposing the need for a good balance between both the attention modules. This is the reason we allow the model to learn these weights during training.

# 5. Limitations

In this section, we discuss the architectural limitations to our work and also explore future extensions to this approach. Firstly, the performance of the decoder phase is dependant on the output from the topic classifier and bottlenecks the overall improvement from training. Moreover, most recent works such as GCN-LSTM [30] and AoA-Net [11] make use of Faster-RCNN to feed in region-level information and hence incorporating these object-level associations alongside topics in the multi-modal embedding space are susceptible to increase in efficacy of the approach used. Another limitation of our work is the use of traditional attention mechanisms. Our study make use of softattention mechanism which involves the averaging of feature maps. Comparing our approach with HAN [27] which also makes use of soft-attention mechanism, we gain a relative improvement as discussed in Section 3.3.2. However, our approach struggles against AoA-Net [11] which uses





GT: a young girl lying in a hospital hooked up to machines

(µ = 0.3) : a woman is laying on hospital bed

(µ = 0.6) : a laptop computer sitting on top of a desk

GT : a man holding a tennis racket getting ready to hit a tennis ball

 $(\mu = 0.3)$ : a tennis player is **about to hit** the ball

(µ = 0.6) : a man in white shirt and white shorts playing tennis

Figure 6. Sampled captions on varying  $\mu$  values on COCO dataset. A higher value of  $\mu$  denotes more weight being given to the spatial flow of information within the decoder and viceversa.

a more robust attention mechanism. Moroever, the use of self-attention has been shown to improve performance over traditional attention mechanisms such as [11], more notably in transformers [5, 24] and hence can be incorporated with the use of these multi-modal embeddings to improve performance. Lastly, using recent reinformcement learning based techniques such as CIDEr optimizations [18] have yielded state-of-the-art results for image captioning, incorporating them with our study may further boost the performance over the metrics used.

# 6. Conclusion

In this work, we proposed a new approach to guide the attention model by exploiting partial-order relationships between image, captions and topics. Arranging the image and textual modalities in an asymmetric fashion results in more effective learning of the latent space. Hence, we make use of a multi-modal embedding space that is able to arrange the visual and textual modalities in an asymmetrical hierarchy where the caption embeddings are bounded between image and topic features. We then make use of these joint representations to guide the attention module. An extensive ablation study was also performed to indicate that using ordered embeddings, the attention model was able to draw accurate links between semantically important regions of the image when attending to them, which helped improve the overall interpretability, syntax and descriptiveness of the captions. The proposed architecture was not only simpler in terms of complexity, but also competitive with many recent LSTMbased architectures. For next steps, a promising direction can be to incorporate the highlighted approach with transformers or leveraging the model architecture to be trained in an end-to-end manner.

# References

- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. 2, 6
- [2] Omer Arshad, Ignazio Gallo, Shah Nawaz, and Alessandro Calefati. Aiding intra-text representations with visual context for multimodal named entity recognition. In 2019 International Conference on Document Analysis and Recognition (ICDAR), pages 337–342. IEEE, 2019. 1
- [3] Shuang Bai and Shan An. A survey on automatic image caption generation. *Neurocomputing*, 311:291–304, 2018. 1
- [4] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003. 2
- [5] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning, 2020. 6, 8
- [6] Ignazio Gallo, Alessandro Calefati, and Shah Nawaz. Multimodal classification fusion in real-world scenarios. In 2017 14th IAPR international conference on document analysis and recognition (ICDAR), volume 5, pages 36–41. IEEE, 2017. 1
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [8] Simao Herdade, Armin Kappeler, Kofi Boakye, and Joao Soares. Image captioning: Transforming objects into words, 2020. 1
- [9] Chiori Hori, Takaaki Hori, Teng-Yok Lee, Ziming Zhang, Bret Harsham, John R Hershey, Tim K Marks, and Kazuhiko Sumi. Attention-based multimodal fusion for video description. In *Proceedings of the IEEE international conference on computer vision*, pages 4193–4202, 2017. 1
- [10] MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, and Hamid Laga. A comprehensive survey of deep learning for image captioning. ACM Computing Surveys (CsUR), 51(6):1–36, 2019. 1
- [11] Lun Huang, Wenmin Wang, Jie Chen, and Xiaoyong Wei. Attention on attention for image captioning. *CoRR*, abs/1908.06954, 2019. 2, 6, 8
- [12] Wenhao Jiang, Lin Ma, Yu-Gang Jiang, Wei Liu, and Tong Zhang. Recurrent fusion network for image captioning, 2018. 2, 5, 6
- [13] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015. 1, 3, 4
- [14] Lei Ke, Wenjie Pei, Ruiyu Li, Xiaoyong Shen, and Yu-Wing Tai. Reflective decoding network for image captioning, 2019. 2, 5, 6
- [15] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neu-

ral language models. *arXiv preprint arXiv:1411.2539*, 2014. 1, 4

- [16] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 375–383, 2017. 2, 4, 6, 7
- [17] Shah Nawaz, Muhammad Kamran Janjua, Ignazio Gallo, Arif Mahmood, Alessandro Calefati, and Faisal Shafait. Do cross modal systems leverage semantic relationships? In Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops, pages 0–0, 2019. 1
- [18] Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. *CoRR*, abs/1612.00563, 2016.
- [19] Muhammad Saad Saeed, Muhammad Haris Khan, Shah Nawaz, Muhammad Haroon Yousaf, and Alessio Del Bue. Fusion and orthogonal projection for improved face-voice association. arXiv preprint arXiv:2112.10483, 2021. 1
- [20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 5
- [21] Moses Soh. Learning cnn-lstm architectures for image caption generation. Dept. Comput. Sci., Stanford Univ., Stanford, CA, USA, Tech. Rep, 2016. 1
- [22] Mike W Spratling and Mark H Johnson. A feedback model of visual attention. *Journal of cognitive neuroscience*, 16(2):219–237, 2004. 2
- [23] Gargi Srivastava and Rajeev Srivastava. A survey on automatic image captioning. In *International Conference on Mathematics and Computing*, pages 74–83. Springer, 2018.
- [24] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. arXiv preprint arXiv:1706.03762, 2017. 6, 8
- [25] Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun.
   Order-embeddings of images and language, 2016. 1, 2, 3, 4, 5
- [26] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 3156–3164, 2015. 1
- [27] Weixuan Wang, Zhihong Chen, and Haifeng Hu. Hierarchical attention network for image captioning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):8957– 8964, Jul. 2019. 2, 5, 6, 8
- [28] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015. 2
- [29] Zhongliang Yang, Yu-Jin Zhang, Sadaqat ur Rehman, and Yongfeng Huang. Image captioning with object detection and localization, 2017. 2
- [30] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *Proceedings*

of the European Conference on Computer Vision (ECCV), September 2018. 2, 6, 8

- [31] Ting Yao, Yingwei Pan, Yehao Li, Zhaofan Qiu, and Tao Mei. Boosting image captioning with attributes, 2016. 1, 6
- [32] Niange Yu, Xiaolin Hu, Binheng Song, Jian Yang, and Jianwei Zhang. Topic-oriented image captioning based on order-embedding. *IEEE Transactions on Image Processing*, 28(6):2743–2754, 2018. 1, 3, 4, 7
- [33] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. Unified vision-language pretraining for image captioning and vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13041–13049, 2020. 1
- [34] Zhihao Zhu, Zhan Xue, and Zejian Yuan. Topic-guided attention for image captioning. In 2018 25th IEEE International Conference on Image Processing (ICIP), pages 2615– 2619, 2018. 1, 2