

MHSAN: Multi-Head Self-Attention Network for Visual Semantic Embedding

Geondo Park

Chihye Han

Daeshik Kim

Wonjun Yoon

Korea Advanced Institute of Science and Technology (KAIST)

Lunit Inc.

{geondopark, ckelseyhan, daeshik}@kaist.ac.kr

wonjun@lunit.io

Abstract

Visual-semantic embedding enables various tasks such as image-text retrieval, image captioning, and visual question answering. The key to successful visual-semantic embedding is to express visual and textual data properly by accounting for their intricate relationship. While previous studies have achieved much advance by encoding the visual and textual data into a joint space where similar concepts are closely located, they often represent data by a single vector ignoring the presence of multiple important components in an image or text. Thus, in addition to the joint embedding space, we propose a novel multi-head self-attention network to capture various components of visual and textual data by attending to important parts in data. Our approach achieves the new state-of-the-art results in image-text retrieval tasks on MS-COCO and Flickr30K datasets. Through the visualization of the attention maps that capture distinct semantic components at multiple positions in the image and the text, we demonstrate that our method achieves an effective and interpretable visual-semantic joint space.

1. Introduction

Various computer vision applications such as image captioning, visual question answering (VQA), and image-text retrieval rely on multi-modal information from images and natural languages. Visual-semantic embedding is to map high-dimensionally represented visual and textual features in the common space where their complex relationship is captured, and it is generally assessed by image-text retrieval. The key to visual-semantic embedding is how accurately input data from each modality are expressed. One popular approach for learning the visual-semantic embedding is to employ a two-path network, which represents visual and textual data separately, using the mono-modal encoders [12, 4, 3].

Recent works on two-path networks collapse the visual features into a single vector by fully connected layer [12, 4] or selective spatial pooling [3] and the textual features using

the last hidden state of the recurrent neural network. However, since an image or a sentence consists of a set of sub-components, there can be various descriptions for a single image depending on where the focus is placed within the image, and it is important to propagate information about each component separately when embedding the respective semantics. Previous methods can easily identify the decisive features, but rarely consider the relative importance of each feature and complex semantics.

Motivated by this observation, we propose a multi-head self-attention network for visual-semantic embedding that captures various aspects and components in images and texts. Departing from previous methods, our proposed model generates multiple attention weights focusing on different visuospatial features of the image and semantic features of the sentence in each encoder. Features with a stronger focus are considered more important and have more influence toward the final representation. By explicitly learning ‘on what point in the image’ and ‘on which word in the sentence’ to focus from various perspectives, our model provides more abundant information about visual and textual data for joint embedding. In other words, our model encodes important regions and words separately and weighs the relative importance between them, which is simply collapsed by selective spatial pooling or even single attention.

We evaluate visual-semantic embedding generated by our proposed model on image-text retrieval tasks using the benchmark MS-COCO [16] and Flickr30K [31] datasets. The results show that our multi-head attention network establishes the new state-of-the-art record on the two benchmark datasets. Most notably, on Flickr30K [31], our embedding network achieves the relative improvement compared to the current state-of-the-art method [10] by 25.4% in sentence retrieval, and 27.5% in image retrieval. On MS-COCO [16], our model also outperforms the current state-of-the-art methods by a large margin. In addition, we perform ablation studies and visualization of the resulting attention maps to demonstrate that the proposed methods of self-attention, multi-head, and diversity regularization each meaningfully contribute to the improved performance.

Contribution. Overall, our main contributions are as follows:

1. We propose a new visual-semantic embedding network based on the multi-head self-attention module that extracts multiple aspects of semantics in visual and textual data.
2. We demonstrate that our model successfully learns semantically meaningful joint embedding by setting new state-of-the-art records in the image-text retrieval on standard benchmark datasets: MS-COCO and Flickr30K.
3. From the advantage of the interpretability of our model, we validate the effectiveness of multi-head self-attention and diversity regularization.
4. While the method of generating multiple attention is used in the natural language process [17] [23] [25], to the best of our knowledge, this is the first work that generates multiple self-attention maps for image encoding as well as visual-semantic embedding.

2. Related Work

Visual-Semantic Embedding. Learning visual-semantic embedding aims to map natural images and texts in a common space which captures the complex relations between them and is evaluated by performing image-text retrieval tasks. Recent deep learning based approaches [12, 4, 3, 28] proposed two-path architectures for learning visual-semantic embedding, which generates image and text representations separately through each mono-modal encoder and connects the two-path with fully connected layers. Kiros *et al.* [12] first attempted to create a multi-modal joint embedding space learned using a two-path network with triplet ranking loss. Faghri *et al.* [4] suggested advanced triplet ranking loss incorporating hard negative mining and showed a significant improvement in image-text retrieval tasks while keeping the two-path architecture. The aforementioned works used the fully connected layer of the convolutional network trained on the classification task [22] as the image encoder and recurrent neural network [9] as the text encoder. For further improvements, Engilberge *et al.* [3] extracted one-dimensional image representation by applying selective spatial pooling to image features to enhance the visual feature encoder. Huang *et al.* [10] proposed learning semantic concepts and order method to enhance the image representation, which is the current state-of-the-art performance in the image-text retrieval tasks. Wehrmann *et al.* [28] proposed a character-level convolution text encoder using a modified Inception Module [24] to improve the textual encoder. Wu *et al.* [30] improved robustness of visual semantic embedding through structured semantic representations of textual data.

With a promising improvement of performance gained by integrating the attention network on many tasks, several

works [20, 14] have tried to perform image-text retrieval tasks with the attention mechanism. DAN [20] applies a cross-attention network, which focuses on the relevant parts of the features between images and texts through multiple step process. Lee *et al.* [14] proposed the cross-attention network considering relative importance between localized regions of an image and the words that requires the bounding box ground truth which are expensive to pre-train object proposals as used in Fast-RCNN [6]. In this study, while maintaining two-path, we utilize the self-attention mechanism to focus on important regions and words without object proposals for visual-semantic embedding.

Self-Attention. Self-attention networks [21], also called intra-attention networks, are trainable end-to-end, that is, simultaneously train feature extraction and attention generation. There have been several approaches utilizing self-attention to focus on the important parts in order to compute a representation of image and text data [21, 29, 1, 19]. From the motivation to focus on various aspects of the data, several previous works generated attention weights of the same input. Vaswani *et al.* [25] improved machine translation by attending to multiple positions of encoded features. Sukhbaatar *et al.* [23] represented the question sentence into a continuous vector by generating attention sequentially for question and answering task. Lin *et al.* [17] have proposed a structured self-attention sentence embedding that represents the sentences with multiple points of attention. MCB [5] achieved a core improvement for visual question answering by their own fusion method and attention mechanism that computes one attention weight from multiple image features and the text feature. In contrast with previous studies, we propose a multi-head self-attention network which generates multiple attention weights that encode the varying localized regions from various perspectives for an image. In our proposed model, self-attention apprehends the various semantics in the image and text for learning visual-semantic embedding.

3. Approach

In order to reflect various semantics of visual and textual data in the joint embedding space, we propose multi-head self-attention model generating multiple attention weights, where each weight focuses on the different regions in the image and on the different words in the sentence. Basically, the proposed model follows the two-path network that consists of separate image and text encoders, shown in Figure 1. We describe the proposed model in four components: (1) visual feature extractor, (2) textual feature extractor, (3) multi-head self-attention, and (4) contrastive triplet ranking loss with diversity regularization.

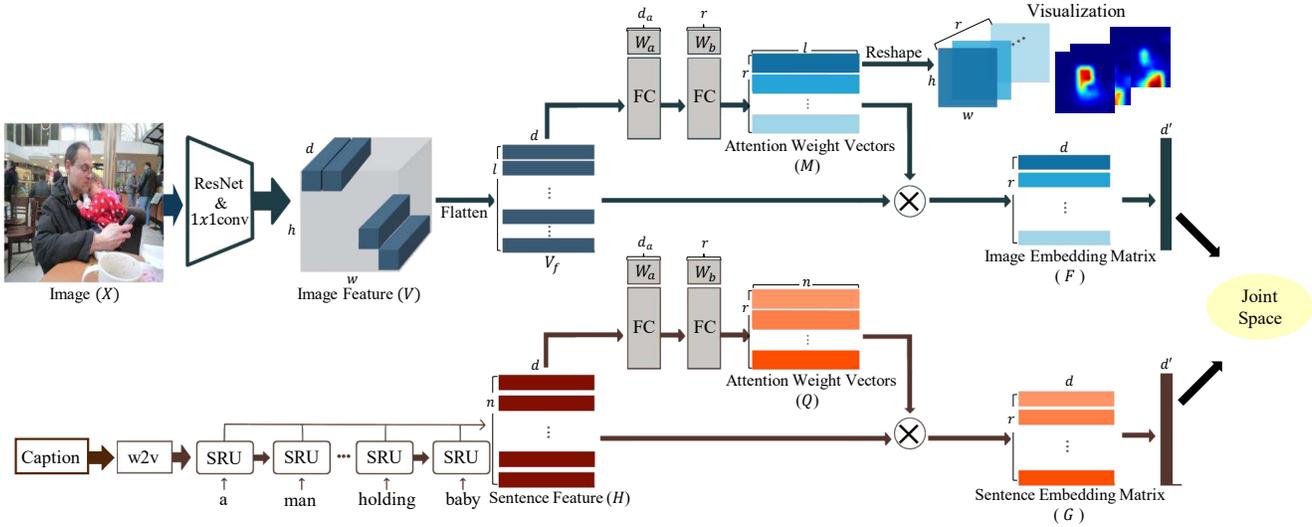


Figure 1: Details of the proposed architecture for learning semantic-visual embedding with multi-head self-attention. For each encoder, features of image(V) and sentence(H) are extracted by ResNet and SRU. On top of the feature extractor, multi-head self-attention network generates attention weight vectors. The final embedding of the image and text is a 2-dimensional matrix.

3.1. Visual Feature Extractor

Similar to [3], we use a ResNet-152 [8] pre-trained on the ImageNet classification task to extract visual features in the images. In accordance with our purpose, we eliminated all of the fully-connected layers and global average pooling and only used the convolutional layers to extract 3-dimensional feature maps preserving the spatial information of the images. We then superinduce an additional 1×1 convolutional layer to perform linear adaptation and to make features in the desired dimension of d . Given an image $X \in (0, 255)^{W \times H \times 3}$ where W, H are the width and height of the image, the transformed ResNet-152 outputs the 3-dimensional visual feature of size $V \in \mathbb{R}^{w \times h \times d}$ where $(w, h) = (\frac{W}{32}, \frac{H}{32})$. Then, we flatten the 3-D feature map into a 2-D feature map as $V_f = (v_1, v_2, \dots, v_l)$, where $l = w \times h$ and $v_{1,2,\dots,l} \in \mathbb{R}^d$ are features containing spatial information in the images.

3.2. Textual Feature Extractor

Before extracting feature of the sentences, each sentence is divided into a word unit through the tokenizer and each word is expressed as a vector representation. Given a sentence S with length of n tokens, $S = (w_1, w_2, \dots, w_n)$, $S \in \mathbb{R}^{k \times n}$. Here, w_i is a k -dimensional word vector of a i -th word in the sentence and S is a set of word vectors that are independent of each other. In this study, we employ $k = 620$ word embedding used in [13] for the word representation. To extract sentence features containing the dependency between each word in a sentence, we use the 4-layer simple recurrent units (SRU) [15] with a dropout:

$h_i = \overrightarrow{SRU}(w_i, h_{i-1})$ where h_i is a i -th hidden state of SRU. In contrast to other works which use only the last hidden state h_n from recurrent neural networks as a sentence feature, we use all the hidden states as a sentence feature where each hidden state represents the context-aware word in the sentence. Thus, we define the final sentence feature as $H = (h_1, h_2, \dots, h_n)$, $h_{1,2,\dots,n} \in \mathbb{R}^d$ where d is the dimension of the hidden states.

3.3. Multi-Head Self-Attention

We apply the self-attention module on the image and text encoders separately to focus on important spatial regions for image representation and words for sentence representation. Moreover, due to the complex relationship between images and texts, it is difficult to contain the semantic meanings of the image and text within a single vector calculated by a single attention weight. To solve this problem, we propose the multi-head self-attention mechanism on the image and text encoder to capture the diverse aspects of semantic meaning in the visual and textual data. Another advantage of the multi-head attention architecture is that the representation vectors of an image and sentence retain the important information that could be missed out on single attention which captures globally important information in the whole image scene or sentence. On top of the feature extractor of each modality, the multi-head self-attention module is applied. In case of the image encoder, our goal is to learn a set of attention weight vectors used for refining important spatial features, where each attention weight vector focuses on different subsets of spatial feature. To infer the multiple at-

tention weight vectors, the multi-head self-attention module takes the flattened image features as an input and outputs 2-dimensional attention weight matrix (M), as illustrated in Figure 1. In particular, the flattened features pass through a multi-layer perceptron (MLP) with one hidden state and one activation function. Then, we apply the softmax along the second dimension of output from MLP to ensure the computed each attention weights sum up to 1. Therefore, each row of the attention weight matrix attends to a different part in the image. In short, the multi-head attention weight matrix is computed as follows:

$$M = \text{softmax}(W_a \text{ReLU}(W_b V_f^T)) \quad (1)$$

where $V_f^T \in \mathbb{R}^{d \times l}$, $W_b \in \mathbb{R}^{d_a \times d}$, $W_a \in \mathbb{R}^{r \times d_a}$, $M \in (0, 1)^{r \times l}$, and d_a is a hidden state dimension and r is the number of attention weight. In the case of $r = 1$, we denote it as single-head self-attention where the model calculates only one attention weight for each image spatial feature.

With a set of attention weight vectors, the model generates 2-dimensional image embedding matrix through a linear combination of the flattened image feature,

$$F = M V_f \quad (2)$$

where $F \in \mathbb{R}^{r \times d}$ is an image embedding matrix which consists of r -vectors. More specifically, given the $F = (f_1, f_2, \dots, f_r)$, each row vector of F is calculated as $f_r = \sum_{l=1}^l M_{r,l} * v_l$ where $f_r \in \mathbb{R}^d$ and $l = w \times h$.

The same attention module is applied in the text encoder replacing image features with sentence features. As a result of the text encoder, the sentence embedding matrix (G) is generated with a set of attention weight vectors (Q).

To map image and text embedding matrix into a joint embedding space, we concatenate r -vectors and project into a single vector using additional projection layer with normalization.

Further, to visualize the attention maps in Section 4.3, we reshape each attention weight vector into a feature size of (w, h) and resize into the same size as the input image using bi-linear interpolation. For multi-head self-attention, r -attention heatmaps where each heatmap focuses on a different local region in an image are provided.

3.4. Contrastive Triplet Ranking Loss with Diversity Regularization

We basically use the triplet ranking loss with hard negative ($Loss_T$), which has shown a significant improvement on image-text retrieval tasks in the previous studies [4, 3, 14]. Since the multi-head attention mechanism can suffer from the redundancy problem where each attention vector focuses on the very similar region, we use the diversity regularization loss ($Loss_D$) with a coefficient(λ)

employed in [17]:

$$Loss_{total} = Loss_T + \lambda Loss_D \quad (3)$$

Given a sampled training set of image and text pairs in a batch: $B = \{(X_n, Y_n)\}_{n=1}^N$ where X_n is the n -th image and Y_n is the n -th sentence in a batch, the image vector x_n denote the image X_n and sentence vector y_n denote the sentence Y_n in the joint embedding space. Then, triplet ranking loss with hard negative is defined as:

$$Loss_T = \frac{1}{N} \sum_{n=1}^N \left(\max_{m \in N, m \neq n} [\alpha - \cos(x_n, y_n) + \cos(x_n, y_m^-)]_+ + \max_{m \in N, m \neq n} [\alpha - \cos(y_n, x_n) + \cos(y_n, x_m^-)]_+ \right) \quad (4)$$

where y_m^- , x_m^- are negative instances, (x_n, y_n) is a positive pair in a batch, α is the margin and $[t]_+ = \max(0, t)$. The contrastive triplet ranking loss encourages the matched embedding pairs to be closer at least margin α apart than mismatched embedding pairs in the joint embedding space. Following the hard negatives mining strategy, the loss function focuses only on a single contrastive example that has the highest similarity score for image and text in the batch while not similar in meaning.

To ensure the diversity of multiple attention weight vectors in each encoder, we add the diversity regularization loss defined as:

$$Loss_D = \|(MM^T - I)\|_F^2 + \|(QQ^T - I)\|_F^2 \quad (5)$$

where M and Q is the R -multiple attention weight matrix from the each image and text encoder, $I \in \mathbb{R}^{r \times r}$ is the identity matrix and $\|\cdot\|_F$ denotes the squared frobenius norm. Minimizing the regularization term with triplet ranking loss, the matrices M and Q become sparse so that each of the attention vector focuses on the different region and word in each encoder. We show the effect of the diversity regularization in Section 4.4.

3.5. Implementation Details

Image Encoder First, we take a random rectangular crop from an image and resize the image into a fixed size of (256×256) for training images. Then, we use ResNet-152 and 1×1 convolution where the output dimension (d) is set to 2400. We set a hidden state of MLP, d_a , to 350 and the joint embedding dimension, d' to 2400. We use 0.5 probability of dropout on the MLP of self-attention mechanism and the projection layer for the joint embedding.

Text Encoder We use tanh and 0.25 dropout in the SRU text encoder. The self-attention mechanism and the projection layer is set the same as the image encoder. Unlike the self-attention mechanism in the image encoder, we use tanh instead of relu between MLP.

Method	Sentence Retrieval			Image Retrieval		
	R@1	R@5	R@10	R@1	R@5	R@10
	1K Test Images					
Order-embeddings [26]	46.7	-	88.9	37.9	-	85.9
Embedding network [27]	50.4	79.3	89.4	39.8	75.3	86.6
2-Way Net [2]	55.8	75.2	-	39.7	63.3	-
CHAIN-VSE [28]	61.2	89.5	95.8	46.6	81.9	90.92
NAN [10]	61.3	87.9	95.4	47.0	80.8	90.1
UniVSE [30]	64.3	89.2	94.8	48.3	81.7	91.2
VSE++ [4]	64.6	90.0	95.7	52.0	84.3	92.0
DPC [32]	65.6	89.8	95.5	47.1	79.9	90.0
GXN [7]	68.5	-	97.9	56.6	-	94.5
Engilberg <i>et al.</i> [3]	69.8	91.9	96.6	55.9	86.9	94.0
SCO [10]	69.9	92.9	97.5	56.7	87.5	94.8
Ours (1-head)	72.4	93.7	97.5	57.5	88.0	94.4
Ours (5-head)	72.6	93.7	97.6	58.5	88.0	94.5
Ours (10-head)	73.5	94.2	98.1	59.1	88.5	94.8
	5K Test Images					
Order-embeddings [26]	23.3	-	84.7	31.7	-	74.6
VSE++ [4]	41.3	71.1	81.2	30.3	59.4	72.4
GXN [7]	42.0	-	84.7	31.7	-	74.6
SCO [10]	42.8	72.3	83.0	33.1	62.9	75.5
Ours (10-head)	49.3	78.3	87.1	36.2	67.2	78.7

Table 1: Results of experiments on MS-COCO, the best in bold.

Training Process We set the margin of triplet ranking loss, α , to 0.2 and the coefficient of diversity regularization to 0.1. We use the Adam optimizer for the whole training. We first update only the last projection layer of the image encoder for the 4 epochs with 0.005 initial learning rate. Then, we update the SRU for the 15 epochs with 0.1 decayed learning rate. Finally we train the rest of the parameters in the image encoder together for the 40 epochs with 0.1 decayed learning rate.

4. Experiment

We first evaluate the quantitative result of our cross-modal retrieval model on the standard image-text aligned datasets, MS-COCO and Flickr30K, and compare our models with recent other methods. Then, we perform additional experiments to demonstrate that the self-attention, multi-head, diversity regularization on each encoder contribute to respective performance improvement. Additionally, through visualization of the attention maps, we qualitatively show how our model works and what our model has propagated into the joint embedding space. Our pytorch implementation is available as open source.¹

Dataset and Evaluation Metric. MS-COCO and Flickr30K datasets consist of 123,287, 31,783 images respectively and 5 captions are annotated for each image. We follow training/validation/test splits used in [11]: 1)MS-COCO, 113,287 training, 5,000 validation, 5,000 test images. 2)Flickr30K, 29,783 training, 1,000 validation, 1,000 test images. For comparing the performance of our model

¹Github: <https://github.com/GeondoPark/MHSAN>

²We use the bidirectional SRU in the text extractor for this results.

Method	Sentence Retrieval			Image Retrieval		
	R@1	R@5	R@10	R@1	R@5	R@10
2-Way Net [2]	49.8	67.5	-	36.0	55.6	-
VSE++ [4]	52.9	80.5	87.2	39.6	70.1	79.5
DAN [20]	55.0	81.8	89.0	39.4	69.2	79.1
NAN [18]	55.1	80.3	89.6	39.4	68.8	79.9
DPC [32]	55.6	81.9	89.5	39.1	69.2	80.9
SCO [10]	55.5	82.0	89.3	41.1	70.5	80.1
Ours (1-head)	66.2	89.8	93.9	49.0	79.0	86.5
Ours (5-head)	67.5	89.6	94.8	51.8	80.0	87.8
Ours (10-head) ²	69.6	92.2	95.8	52.4	80.7	88.6

Table 2: Results of experiments on Flickr30K, the best in bold.

with state-of-the-art methods in image-text retrieval, we report the same evaluation metric with previous works [4, 3, 14, 28, 10]. Recall@K of the text retrieval is the percentage of the correct text being ranked within the top-K retrieved results, and vice versa. For MS-COCO dataset, we additionally report averaging R@K results, which are performed 5 times on 1,000-image subsets of the test set and testing on the full 5,000 test images.

4.1. Comparison with the state-of-the-art

MS-COCO Retrieval Task. Table 1 shows the quantitative results of our model on MS-COCO dataset. We denote our models with the number of attention weight vectors (#-head), which is the same in the image and text encoders. Comparing with recently proposed methods, our model with 10-head attention network achieves the new state-of-the-art performance on MS-COCO datasets. On the 1K test set experiment, our models outperform SCO [10], the current state-of-the-art based on the visual-semantic embedding method, by 5.2% in sentence retrieval (R@1) and by 4.2% in image retrieval (R@1). Also, our best results on 5k test set experiment show 15.2% and 9.4% relative improvement in sentence retrieval (R@1) and in image retrieval (R@1), respectively.

Flickr30K Retrieval Task. We also report the experimental results on Flickr30K, which is a relatively smaller dataset in Table 2. The network architecture and dimension of the joint embedding are the same as those used for MS-COCO. Compared to the recently proposed methods, our model achieves the best performance across all measures with a large gap. Our model with 10-head relatively improves 25.4% on sentence retrieval (R@1) and 27.5% on image retrieval (R@1) compared with the current state-of-the-art score from SCO [10]. This result also demonstrates that our model works well for small dataset, which can cause overfitting emerged in the previous two-path architectures [3, 28].

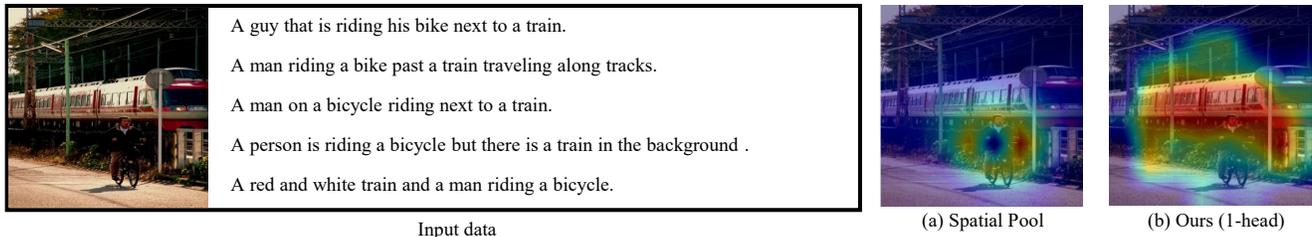


Figure 2: Comparison of single-head (self-attention) with spatial pooling (without attention)

Image Encoder	Text Encoder	Sentence Retrieval(R@1)	Image Retrieval(R@1)
W/O attention [3]	W/O attention [3]	69.8	55.9
Our image encoder(1-head)	W/O attention [3]	70.1	55.8
W/O attention [3]	Our text encoder(1-head)	69.9	56.2

Table 3: Effect of self-attention (1-head) in each encoder. We denote the image encoder using selective spatial pooling and the text encoder using the last hidden state which are used in Engilberg *et al.* [3] as W/O attention. We report only R@1 for effective comparison.

4.2. Comparison of Self-Attention and Spatial Pooling

In this subsection, we show that the single-head self-attention improved the performance of retrieval tasks compared to the selective spatial pooling [3] that did not use the attention mechanism. For this experiment, we use the MSCOCO dataset. To see the effect of the single-head self-attention mechanism at each encoder, we apply the single-head attention to each encoder separately. Table 3 shows that our model with a single-head attention at each encoder slightly improves the performance in the image-text retrieval task.

Visualization. We validate qualitatively why the self-attention performs better than the network using selective spatial pooling [3] through visualization of the attention heatmaps. In Figure 2, we present a comparison of attention heatmap computed by our model with the single-head and the selective spatial pooling heatmaps. For visualizing which region is considered for embedding in the selective spatial pooling, we employ the segmentation method pro-

posed in [3] substituting a phrase with a full sentence. We observe that self-attention with a single-head captures the missing parts on the selective spatial pooling, e.g. the train contained in descriptions of the image as shown in Figure 2. Through the visualization, we demonstrate that our model is more suitable for visual semantic embedding, where the encoding of rich details is necessary.

4.3. Comparison of Multi-Head Self-Attention and Single-Head Self-Attention

We improve the performance of image-text retrieval tasks by extending our model to multi-head self-attention as shown in Table 1 and Table 2. To observe the effect of the number of attention weight vectors in each encoder, we experiment by fixing one encoder without self-attention and changing the number of attention maps in the other. Table 4 shows that more attention maps result in better performance and we achieve the state-of-the-art results on image-text retrieval by using 10-head self-attention in both encoders. Also, we observe that applying the multi-head self-attention

Image Encoder	Text Encoder	Sentence Retrieval(R@1)	Image Retrieval(R@1)
Our image encoder(1-head)	W/O attention	70.1	55.8
Our image encoder(5-head)	W/O attention	70.9	56.2
Our image encoder(10-head)	W/O attention	71.1	57.1
W/O attention	Our text encoder(1-head)	69.9	56.2
W/O attention	Our text encoder(5-head)	71.2	57.4
W/O attention	Our text encoder(10-head)	71.7	58.0
Our image encoder(1-head)	Our text encoder(1-head)	72.4	57.5
Our image encoder(10-head)	Our text encoder(10-head)	73.5	59.1

Table 4: Effect of the number of heads in each encoder. We denote the image encoder using selective spatial pooling and the text encoder using the last hidden state which are used in Engilberg *et al.* [3] as W/O attention. We report only R@1 for effective comparison.

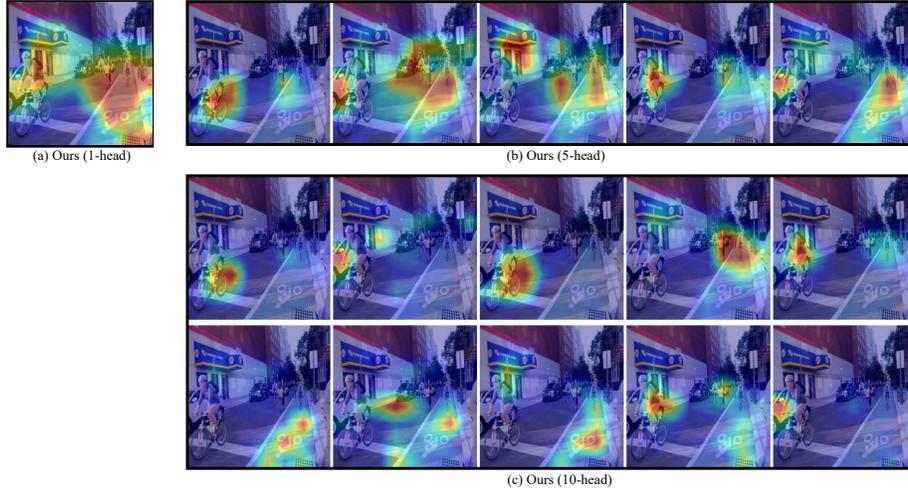


Figure 3: An example of image attention heatmaps shows the comparison of ours with varying numbers of attention (multi-head). While self-attention with single-head (a) is globally collapsing the street, the crowd, and the person riding the bike in one attention map, the multi-head self-attention (b),(c) show that they are attended to separately in multiple attention maps.

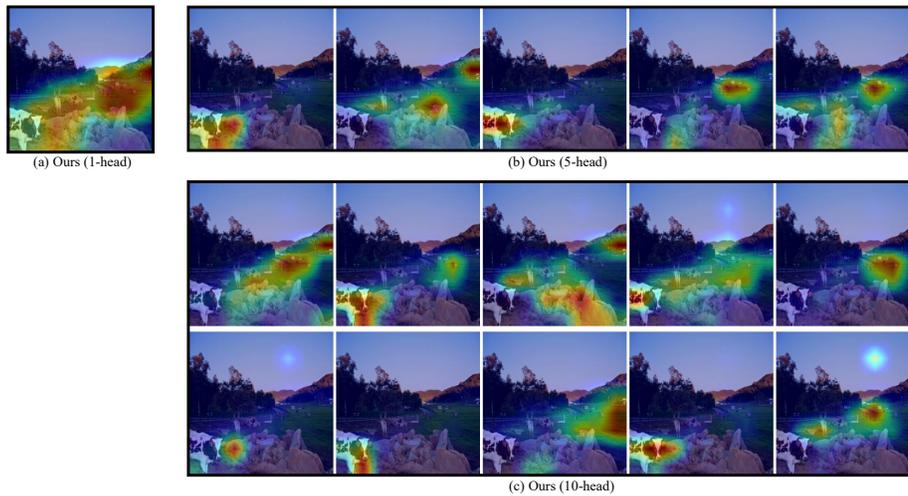


Figure 4: An example of image attention heatmaps show the comparison of ours with varying numbers of attention (multi-head). While self-attention with single-head (a) is globally collapsing the green field with other objects such as the cow in one attention map, the multi-head self-attention (b),(c) show that they are attended to separately in multiple attention maps.

only to the text encoder is more effective than the image encoder only.

Visualization. To observe how the image is encoded into the joint space with our proposed model, we visualize the attention heatmaps learned by the model by varying the number of heads in Figure 3, 4. Although self-attention with a single-head captures the globally important region of the image, the attention weights are distributed over all important region of the image as shown in Figure 3, 4 (a). On the other hand, the model with 5-head separately provide the information of different semantic regions by generating five attention weight vectors as shown in Figure 3, 4

(b). Comparing ours (10-head) to ours (5-head), we show clearly that the object-level regions in the image are concentrated in each attention heatmap as visualized in Figure 3, 4 (c). This result verifies that our model with the multi-head self-attention provides the diverse aspects of semantics in the image to visual-semantic embedding, resulting in better performance than the single-head self-attention. Moreover, we are able to interpret the meaning of each vector in the image embedding using attention heatmaps. For example, as shown in the first attention map in Figure 3 (c), the first row of the embedding matrix (F) in 10-head means the semantics of the bike in the image and the fourth row of

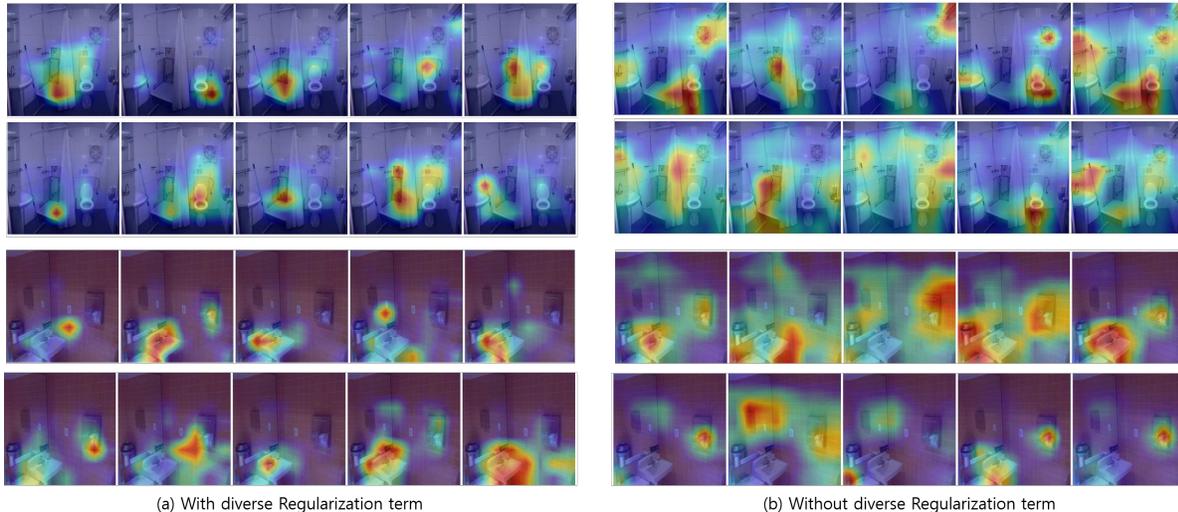


Figure 5: Effect of diversity regularization term. We train our model (10-head) (a) with diverse regularization term in loss (coefficient=0.1) and (b) without diverse regularization term in the loss.

Diversity Regularization	Sentence Retrieval	Image Retrieval	Diversity Loss
W/	73.5	59.1	5.59
W/O	72.3	58.4	7.36

Table 5: Impact of diversity regularization on performance of retrieval. The lower the diversity loss, the more diverse the local regions focused by the multiple attention maps.

embedding matrix (F) means the person riding the bike as shown in the fourth attention map.

4.4. Effect of Diversity Regularization in Loss

In this subsection, we experiment to observe the effect of diversity regularization in the loss. For comparison, we report the average diversity loss value described in equation 5 on the 5k test set and performance of image-text retrieval tasks come from our model (10-head) trained with and without diverse regularization term in Table 5. These results in Table 5 show that diversity regularization has a great effect on image-text retrieval performance. We conjecture that diversity regularization encourages encoded vectors to contain various semantic meanings in an image, reducing the redundancy of the encoded vectors and attention weight vectors.

Visualization. We qualitatively present the multiple attention maps focusing on more distinct local regions by introducing the diversity regularization. As shown in Figure 5, the diversity regularization term in the loss function makes attention matrix sparse, which encourages different regions to be captured by each attention map. Therefore, the image embedding encodes more distinct components in an image, resulting in an improvement of performance on the image-text retrieval tasks.

5. Conclusion

We presented an interpretable and powerful multi-head self-attention network for visual-semantic embedding. Different from existing visual semantic embedding models, our proposed network represents visual and textual data in multiple vectors generating multiple attention weights where each weight encodes a different part in the data. Furthermore, we enhance the model with diversity regularization loss to address the redundancy problem where the multiple attention mechanism provides similar representations for the data. Experiments on two benchmark datasets show that our model outperforms state-of-the-art visual semantic embedding models by a large margin. Through the visualization of attention maps, we demonstrate that our proposed method provides various aspects of visual and textual data by encoding semantically distinct parts of the image and the text separately. Notably, our model is able to achieve superior performance without any other resources such as bounding box or image segmentation datasets. Since we achieve the significant improvement by applying the multi-head attention module independently in a modality encoder, we expect that our method can be widely applied for other computer vision tasks such as image classification, visual question answering, and semantic segmentation.

References

- [1] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- [2] A. Eisenschlat and L. Wolf. Linking image and text with 2-way nets. In *arXiv preprint arXiv:1608.07973*, 2016.
- [3] M. Engilberge, L. Chevallier, P. Pérez, and M. Cord. Finding beans in burgers: Deep semantic-visual embedding with localization. In *CVPR*, 2018.
- [4] F. Faghri, D. J. Fleet, R. Kiros, and S. Fidler. VSE++: improved visual-semantic embeddings. In *BMVC*, 2018.
- [5] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. *EMNLP*, 2016.
- [6] R. B. Girshick. Fast R-CNN. In *ICCV*, 2015.
- [7] J. Gu, J. Cai, S. R. Joty, L. Niu, and G. Wang. Look, imagine and match: Improving textual-visual cross-modal retrieval with generative models. In *CVPR*, 2017.
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [9] S. Hochreiter and J. Schmidhuber. Long short-term memory. In *Neural computation*, pages 1735–1780, 1997.
- [10] Y. Huang, Q. Wu, and L. Wang. Learning semantic concepts and order for image and sentence matching. In *CVPR*, 2018.
- [11] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, pages 3128–3137, 2015.
- [12] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. In *NeurIPS*, 2014.
- [13] R. Kiros, Y. Zhu, R. Salakhutdinov, R. S. Zemel, A. Torralba, R. Urtasun, and S. Fidler. Skip-thought vectors. In *NeurIPS*, 2015.
- [14] K. Lee, X. Chen, G. Hua, H. Hu, and X. He. Stacked cross attention for image-text matching. In *ECCV*, 2018.
- [15] T. Lei, Y. Zhang, S. I. Wang, H. Dai, and Y. Artzi. Simple recurrent units for highly parallelizable recurrence. In *EMNLP*, 2018.
- [16] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, 2014.
- [17] Z. Lin, M. Feng, C. N. dos Santos, M. Yu, B. Xiang, B. Zhou, and Y. Bengio. A structured self-attentive sentence embedding. In *ICLR*, 2017.
- [18] C. Liu, Z. Mao, W. Zang, and B. Wang. A neighbor-aware approach for image-text matching. In *ICASSP*, 2019.
- [19] Y. Liu, C. Sun, L. Lin, and X. Wang. Learning natural language inference using bidirectional LSTM model and inner-attention. In *arXiv preprint arXiv:1605.09090*, 2016.
- [20] H. Nam, J. Ha, and J. Kim. Dual attention networks for multimodal reasoning and matching. *CVPR*, 2017.
- [21] J. Park, S. Woo, J. Lee, and I. S. Kweon. BAM: bottleneck attention module. In *BMVC*, 2018.
- [22] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, pages 211–252, 2015.
- [23] S. Sukhbaatar, A. Szlam, J. Weston, and R. Fergus. Weakly supervised memory networks. In *NeurIPS*, 2015.
- [24] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, pages 1–9, 2015.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [26] I. Vendrov, R. Kiros, S. Fidler, and R. Urtasun. Order-embeddings of images and language. *ICLR*, 2015.
- [27] L. Wang, Y. Li, and S. Lazebnik. Learning two-branch neural networks for image-text matching tasks. In *TPAMI*, 2018.
- [28] J. Wehrmann and R. C. Barros. Bidirectional retrieval made simple. In *CVPR*, pages 7718–7726, 2018.
- [29] S. Woo, J. Park, J. Lee, and I. S. Kweon. CBAM: convolutional block attention module. In *ECCV*, 2018.
- [30] H. Wu, J. Mao, Y. Zhang, Y. Jiang, L. Li, W. Sun, and W. Ma. Univse: Robust visual semantic embeddings via structured semantic representations. *CVPR*, 2019.
- [31] P. Young, A. Lai, M. Hodosh, and J. Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. In *TACL*, pages 67–78. MIT Press, 2014.
- [32] Z. Zheng, L. Zheng, M. Garrett, Y. Yang, and Y. Shen. Dual-path convolutional image-text embedding. In *Arxiv Preprint Arxiv:1711.05535*, 2017.