

Asymmetric Visual Semantic Embedding Framework for Efficient Vision-Language Alignment

Yang Liu^{1, 2}, Mengyuan Liu^{3*}, Shudong Huang^{1, 2}, Jiancheng Lv^{1, 2}

¹ College of Computer Science, Sichuan University, Chengdu, 610065, China

² Engineering Research Center of Machine Learning and Industry Intelligence, Ministry of Education, Chengdu, China

³ State Key Laboratory of General Artificial Intelligence, Peking University, Shenzhen Graduate School, Shenzhen, China
liuyyy111@gmail.com, liumengyuan@pku.edu.cn, huangsd@scu.edu.cn, lvjiancheng@scu.edu.cn

Abstract

Learning visual semantic similarity is a critical challenge in bridging the gap between images and texts. However, there exist inherent variations between vision and language data, such as information density, i.e., images can contain textual information from multiple different views, which makes it difficult to compute the similarity between these two modalities accurately and efficiently. In this paper, we propose a novel framework called Asymmetric Visual Semantic Embedding (AVSE) to dynamically select features from various regions of images tailored to different textual inputs for similarity calculation. To capture information from different views in the image, we design a radial bias sampling module to sample image patches and obtain image features from various views. Furthermore, AVSE introduces a novel module for efficient computation of visual semantic similarity between asymmetric image and text embeddings. Central to this module is the presumption of foundational semantic units within the embeddings, denoted as “meta-semantic embeddings.” It segments all embeddings into meta-semantic embeddings with the same dimension and calculates visual semantic similarity by finding the optimal match of meta-semantic embeddings of two modalities. Our proposed AVSE model is extensively evaluated on the large-scale MS-COCO and Flickr30K datasets, demonstrating its superiority over recent state-of-the-art methods.

Code — <https://github.com/liuyyy111/AVSE>

Introduction

Understanding the correspondence between the visual world and human language is one of the fundamental capabilities of artificial intelligence (Gong, Yuan, and Bao 2021; Liu et al. 2023; Liu and Tsang 2015). It motivates much research on vision-and-language tasks. As a foundation task in the vision-and-language domain, image-text matching (Qin et al. 2023; Liu et al. 2022) is devoted to bridging the semantic gap between these two different modalities, which aims to search images for a given textual description or vice versa. The key challenge of image-text matching is to measure the semantic similarity between images and texts.

*Corresponding author.

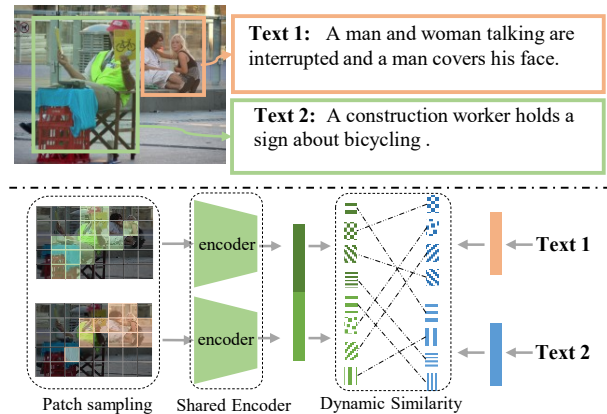


Figure 1: (Top) Information density varies in vision and language data, e.g., an image can be described from multiple different views using language. (Bottom) The conceptual diagram of our proposed method. We first sample different patches (*Radial Bias Sampling*) to the shared encoder (*Vision Transformers*) to get two group embedding, and compute dynamic similarity for different text by selecting different parts of the image feature (*Asymmetric Embedding Optimal Matching*). (Best viewed in color).

To bridge the semantic gap between images and texts, mainstream image-text matching methods follow a common procedure of encoding images and texts into a shared embedding space and measuring the similarity between them, typically optimizing the model using a triplet loss that enforces a greater similarity between matched image-text pairs compared to unmatched pairs. For global-level matching methods, most works (Faghri et al. 2018; Li et al. 2019; Chen et al. 2021) calculate the visual semantic similarity via the inner product (cosine similarity). For local-level matching methods, many works (Lee et al. 2018; Liu et al. 2020; Diao et al. 2021) adopt cross-modal attention mechanism to compute the similarity. Therefore, global-level matching methods have faster inference speed, while local matching methods obtain higher accuracy through cross-modal interaction but also bring high inference costs.

However, the inherent information density difference between images and text, that is, visual data usually contains

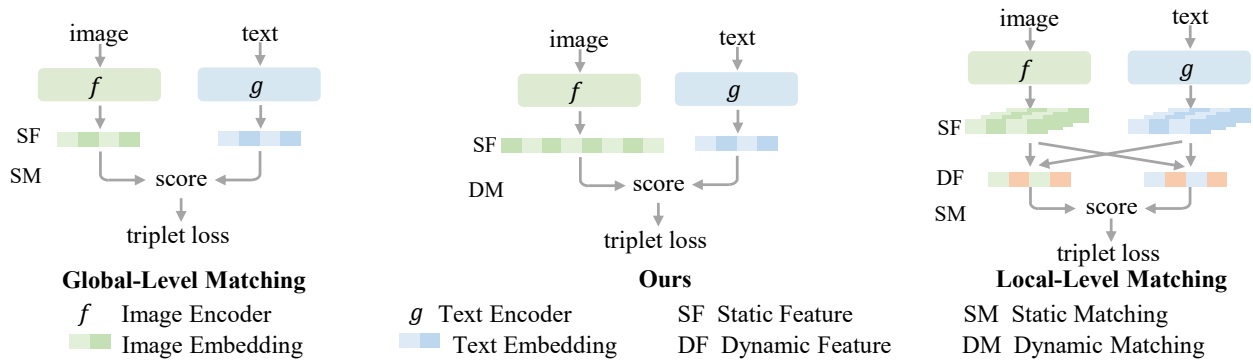


Figure 2: Difference between our proposed AVSE framework and previous methods. “Dynamic features” refer to hybrid features with cross-modal attention, which is a computationally expensive operation. “Dynamic matching” means that when calculating similarity, the meta-semantic embedding of the sentence is used to find the most similar meta-semantic embedding in the image. The process is very simple and only requires cosine similarity.

more information than text data, is a challenge to measure the similarity between the two. As shown in the figure1, human language can be used to describe an image from multiple different perspectives. Existing methods are difficult to solve this problem effectively. The global matching methods always extract static image and text features and calculate the static cosine similarity between the two, which makes it difficult to solve the problem of multiple texts corresponding to an image. In contrast, local-level matching methods use the cross-modal attention mechanism to extract dynamic features and perform cosine similarity on them, which solves this problem to a certain extent. Therefore, the retrieval performance of the local matching method is better. However, the computational complexity of cross-modal attention (Lee et al. 2018) is $O(n^2)$, and the high cost limits its practical application in real-world scenarios.

We analyze that the key to solving the different information densities of images and texts lies in dynamically calculating the similarity between image-text features and obtaining dynamic features will introduce cross-modal interactions with $O(n^2)$ complexity. *So can we dynamically calculate the similarity between the two based on static features?*

Inspired by this, we introduce a novel framework, Asymmetric Visual Semantic Embedding (AVSE), to learn visual semantic similarity by explicitly considering the information density difference between the two modalities, and implementing an $O(n)$ complexity algorithm to select features from various regions of the image tailored to different textual inputs for similarity calculation. Specifically, we introduce an asymmetric feature extraction approach aimed at explicitly capturing richer image information. This method employs a radial bias sampling algorithm to partition image patches into two distinct groups. Subsequently, an encoder is utilized to extract embeddings from each group, which are then concatenated to form comprehensive image embeddings. These embeddings are larger and more information-dense compared to text embeddings, enabling the acquisition of multi-view image features. To effectively compute visual semantic similarity while accounting for the dimen-

sional disparities between text and image embeddings, we propose a novel Asymmetric Embedding Optimal Matching (AEOM) module. Central to this module is the assumption that embeddings comprise fundamental semantic units, termed “meta-semantic embeddings.” The AEOM module decomposes visual and text embeddings into these meta-semantic embeddings. Visual semantic similarity is subsequently determined by aligning each text meta-semantic embedding with its most closely related image meta-semantic embedding, which can dynamically select different parts of image features. Furthermore, we introduce a dimensionality regularization loss function specifically designed to constrain features extracted from various views. This innovative loss function is instrumental in augmenting the AEOM module’s performance by enriching the semantic content of each view and facilitating optimal alignment between meta-embeddings.

Our main contributions are summarized as follows:

- In light of the distinct information densities inherent in images and text, we propose a novel framework to dynamically compute the similarity between static image and text features with $O(n)$ complexity.
- We propose the Asymmetric Embedding Optimal Matching (AEOM) module for dynamically calculating similarity between asymmetric multi-view image and text embeddings. Additionally, we introduce a dimension-wise regularizing loss to enhance the semantic richness of each view and facilitate optimal matching between meta-embeddings.
- We evaluate AVSE with existing fine-grained methods on diverse model backbones. AVSE outperforms the latest state-of-the-art methods on image text retrieval on Flickr30K and MS-COCO, and is also significantly faster than Local-level Matching Methods.

Related Work

Image-text retrieval is a popular research area aimed at bridging the gap between vision and language domains. A

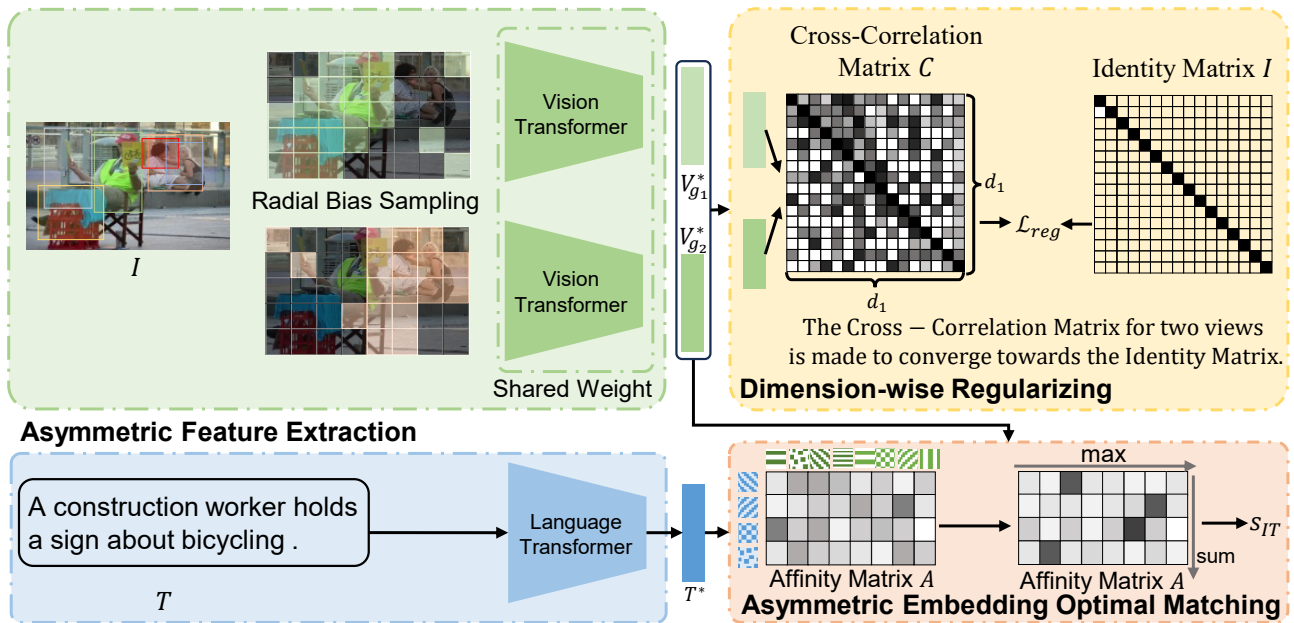


Figure 3: An overview of Asymmetric Visual Semantic Embedding. *Asymmetric Feature Extraction* extract image features from different views to bridge the inherent differences in information density between images and texts. *Asymmetric Embedding Optimal Matching* attempts to learn meta-semantic embeddings of different modalities and calculate similarity through the optimal matching of meta-semantic embeddings between images and texts. *Dimension-wise Regularization* regularizes the embeddings of different image views to assist in learning meta-semantic embeddings.

major challenge in this task is measuring similarities between images and texts. Based on the granularity of the matching similarity, existing works can be broadly classified into two categories: global-level matching methods and local-level matching methods. The difference between previous methods and AVSE as shown in Figure 2.

Global-level matching methods. Early methods (Kiros, Salakhutdinov, and Zemel 2014; Faghri et al. 2018; Vendrov et al. 2016) directly projected images and texts into a shared embedding space, computed similarity within that space, and used a triplet loss for optimization. The goal was to bring semantically aligned image-text pairs closer together while pushing apart those that were not. However, global-level matching methods rely on high-quality image text embeddings to compute visual semantic similarity due to the lack of interaction between images and texts. Recent related work typically uses complex networks to extract image and text features, such as graph convolutional network (Li et al. 2019; Wang et al. 2020, 2022; Cheng et al. 2022), self-attention mechanisms (Wu et al. 2019), or special pooling functions (Chen et al. 2021; Li et al. 2022). But all the global-level methods compute the static similarity with the cosine function, in contrast, *our proposed method is the global-matching method with computing dynamic similarity.*

Local-level matching methods. To learn the latent fine-grained correspondence between images and texts, local-level matching methods calculate visual semantic similarity by aligning the sub-fragments, i.e., regions in images and words in sentences. With the success of bottom-up at-

tention (Anderson et al. 2018) in image captioning and visual question answering, (Lee et al. 2018) propose a stacked cross attention network to attend to image regions with respect to each word in sentences and versa. Subsequently, there are many follow-up works (Qin et al. 2022; Liu et al. 2019, 2021a; Wang et al. 2019; Chen et al. 2020; Liu et al. 2020; Diao et al. 2021; Zhang et al. 2022a,b; Pan, Wu, and Zhang 2023; Fu et al. 2024) that use such stack attention mechanism to learn more fine-grained region-word correspondences. Although such an attention mechanism can extract features dynamically, it also incurs $O(n^2)$ complexity.

Asymmetric Visual Semantic Embedding

In this section, we present our proposed Asymmetric Visual Semantic Embedding (AVSE) model in a formal manner. The objective of this model is to compute the visual semantic similarity of an image-text pair by taking into account the difference in information density between the two modalities. The framework of our proposed AVSE model is illustrated in Figure 3.

Asymmetric Feature Extraction

Given the varied information density in images and text, we expect visual embeddings to hold more data than textual ones. Unlike previous methods, our approach uses asymmetric encoders, allowing image features to be more informative than text features.

Visual Representation With Multiple Views. To devise an image encoder that can capture a multi-view visual repre-

sentation, which contains a richer set of visual information, we propose a novel approach to learn it.

RADIAL BIAS SAMPLING. We introduce *Radial Bias Sampling (RBS)*, a technique that emphasizes sampling points near a randomly selected central point in a 2D space. The key idea behind RBS is to enhance the sampling density around a specific location, thus allowing for finer detail analysis where it matters most. Unlike traditional sampling methods, RBS dynamically adjusts the sampling probability based on the distance from a central point, allowing for a more focused analysis on regions of interest. The process begins by randomly selecting a central point (x_c, y_c) in the given 2D array, which represents the area of primary interest. Therefore, we calculate the Euclidean distance between the central point and every other point in the array.

To implement the bias, we define a weight matrix W where each element $W(i, j)$ represents the weight or importance of the point (i, j) . The weight is inversely proportional to the distance from the central point, often calculated using an exponential decay function:

$$W(i, j) = e^{-\alpha \cdot d(i, j)} \quad (1)$$

where $d(i, j)$ is the Euclidean distance from (i, j) to (x_c, y_c) , and α is a hyperparameter controlling the rate of decay. The higher the value of α , the more localized the sampling becomes around the central point. After constructing the weight matrix, it is normalized to form a probability distribution $P(i, j)$:

$$P(i, j) = \frac{W(i, j)}{\sum_{i, j} W(i, j)} \quad (2)$$

This probability distribution is then used to sample points from the 2D array, ensuring that points closer to the central point have a higher chance of being selected.

MULTI-VIEW IMAGE ENCODER. Given an image I , we use RBS to sample two group patches in Vision Transformers. Then we get two group image feature $V_{g_1}^* \in \mathcal{R}^{d_1}$ and $V_{g_2}^* \in \mathcal{R}^{d_1}$. We can think of these two sets of features as containing information from two different perspectives.

Textual Representation. Given a sentence T , we utilize the standard sequence transformer, BERT(Devlin et al. 2018) as the textual encoder. The sentence undergoes tokenization into linguistic words and is fed to encoder. We get the sentence feature $T^* \in \mathcal{R}^{d_1}$.

Asymmetric Embedding Optimal Matching

To dynamically calculate similarity, we want to select different parts of image features for different texts to calculate similarity. We propose a more fine-grained method to calculate the similarity between image and text modalities. The core of this method is the assumption of "meta-semantic embedding", that is, there are smaller semantic units in the embedding vector, representing more fine-grained semantic information. Visual semantic similarity is calculated by finding the best match of meta-semantic embeddings between different modalities.

Concretely, we first split the visual embedding v and text embedding t into a set of d_2 -dimensional meta-semantic embedding, represented as $v_u = [v_{u_1}, v_{u_2}, \dots, v_{u_p}]$ and $t_u =$

$[t_{u_1}, t_{u_2}, \dots, t_{u_q}]$, respectively, where $p = \frac{n \times d_1}{d_2}$ and $q = \frac{d_1}{d_2}$. Then we compute the affinity matrix A of meta-semantic embeddings of two modalities as follows:

$$A_{i, j} = \frac{v_{u_i}^\top t_{u_j}}{\|v_{u_i}\| \cdot \|t_{u_j}\|}, i \in [1, p], j \in [1, q] \quad (3)$$

where $A_{i, j}$ indicate the affinity score of visual meta-embedding v_{u_i} and text meta-embedding unit t_{u_j} .

In the end, we more precisely compute the correspondence between meta-semantic embeddings of two modalities. Inspired by (Karpathy and Fei-Fei 2015; Lee et al. 2018), we employ the max-sum pooling to calculate the similarity between image I and text T by computing the max over the columns and then summing:

$$S(I, T) = \sum_{j=1}^q \max_{i \in [1, p]} A_{i, j} \quad (4)$$

Dimension-wise Regularization

In the Asymmetric Embedding Optimal Matching module, we calculate visual semantic similarity by finding the optimal match of meta-semantic embeddings between different modalities. This process has an underlying assumption that different views should have same distribution in dimension-wise, which means that corresponding dimensions between different views should have the same semantics. To obtain the optimal match between the text meta-semantic embedding t_u and the visual meta-semantic embedding containing different views v_u , we propose that the image embeddings of different views should contain equivalent amounts of information. While they have different view information, they still should contain a small amount of contextual information about the environment. As such, it is essential that the image embeddings of different views satisfy the same feature distribution, i.e., each channel of the embedding in all group represent the same semantics. By regularizing the channels, it is possible to facilitate the learning of the meta-semantic embedding units and enhance the optimal matching between meta-semantic embeddings.

Specifically, we calculate the cross-correlation matrix C between any two group v_{g_1} and v_{g_2} along the batch dimension as follows:

$$C_{ij} = \frac{\sum_b v_{b,i}^A v_{b,i}^B}{\sqrt{\sum_b (v_{b,i}^A)^2} \sqrt{\sum_b (v_{b,i}^B)^2}} \quad (5)$$

where b indicates the index of batch sample and i, j indicate the vector dimension of embedding, C is a matrix with size of $d_1 \times d_1$. Then, we regularize the embeddings of different views by trying to equate the diagonal elements of the cross-correlation matrix to 1 and equate the off-diagonal elements of the cross-correlation matrix to 0. We define the loss between any two group v_{g_i} and v_{g_j} as:

$$\mathcal{L}_{reg} = \sum_i (1 - C_{ii})^2 + \lambda \sum_i \sum_{j \neq i} C_{ij}^2 \quad (6)$$

where λ is a positive constant to balance the weight of two terms of the loss function.

Method	FG	Flickr30k 1K Test						MS-COCO 5-fold 1K Test						MS-COCO 5K Test					
		Text Retrieval			Image Retrieval			Text Retrieval			Image Retrieval			Text Retrieval			Image Retrieval		
		R1	R5	R10	R1	R5	R10	R1	R5	R10	R1	R5	R10	R1	R5	R10	R1	R5	R10
Pre-trained ResNet-152 on IN & GRU																			
NAAF* (Zhang et al. 2022b)	✓	81.9	96.1	98.3	61.0	85.3	90.6	80.5	96.5	98.8	64.1	90.7	96.5	58.9	85.2	92.0	42.5	70.9	81.4
HREM* (Fu et al. 2023)	✗	81.4	96.5	98.5	60.9	85.6	91.3	81.2	96.5	98.9	63.7	90.7	96.0	60.6	86.4	92.5	41.3	71.9	82.4
CHAN (Pan et al. 2023)	✓	79.7	94.5	97.3	60.2	85.3	90.7	79.7	96.7	98.7	63.8	90.4	95.8	60.2	85.9	92.4	41.7	71.5	81.7
Ours: AVSE*	✗	82.3	96.2	98.5	62.8	88.0	92.8	81.5	96.6	98.8	65.3	91.6	96.4	61.3	86.5	92.8	43.6	73.3	83.4
ViT-Base-224 + BERT-base, 14×14 patches																			
VSE++(Faghri et al. 2018)	✗	71.8	92.8	96.5	59.4	84.7	90.9	75.0	94.6	98.0	62.7	89.4	94.9	62.3	87.6	93.4	43.9	73.6	83.3
SCAN(Lee et al. 2018)	✓	69.5	90.9	95.6	56.4	83.1	90.0	76.0	95.4	98.1	64.5	90.8	95.8	53.9	81.8	90.0	42.9	72.3	82.5
SGR(Diao et al. 2021)	✓	69.7	90.8	95.2	59.1	84.1	89.9	77.2	95.0	98.0	65.1	90.7	95.8	54.9	82.8	90.5	42.8	72.2	82.5
CHAN(Pan et al. 2023)	✓	69.2	91.8	95.0	58.4	84.9	90.6	77.1	95.1	98.1	65.0	91.0	96.0	56.3	83.2	90.1	43.0	72.6	82.8
LAPS(Fu et al. 2024)	✓	74.0	93.4	97.4	62.5	87.3	92.7	78.7	95.5	98.3	66.2	91.3	96.2	57.5	84.0	90.8	44.5	74.0	83.6
Ours: AVSE	✗	76.0	94.6	97.5	62.7	88.4	93.1	79.8	95.6	98.3	67.0	91.5	96.3	58.8	84.3	91.0	45.1	74.3	83.9
ViT-Base-384 + BERT-base, 24×24 patches																			
VSE++(Faghri et al. 2018)	✗	77.1	95.7	97.5	65.8	90.2	94.3	77.0	95.7	98.4	64.6	91.1	96.2	54.9	82.8	90.4	42.4	72.4	82.8
SCAN(Lee et al. 2018)	✓	75.4	94.4	96.9	63.6	88.6	93.5	76.1	95.5	98.5	65.1	91.6	96.3	53.3	81.8	90.0	42.6	72.6	82.9
SGR(Diao et al. 2021)	✓	76.9	94.9	98.1	64.2	88.4	93.3	75.8	95.7	98.6	65.6	92.0	96.5	53.3	81.0	89.6	42.9	73.1	83.7
CHAN(Pan et al. 2023)	✓	75.4	94.5	97.6	63.2	88.6	93.1	78.1	95.8	98.6	66.1	92.1	96.6	55.6	83.8	91.2	43.4	73.6	83.5
LAPS(Fu et al. 2024)	✓	79.0	96.0	98.1	67.3	90.5	94.5	78.6	96.3	98.9	68.0	92.4	96.8	57.4	84.9	92.5	46.4	75.8	85.2
Ours: AVSE	✗	80.3	96.4	98.7	67.9	91.2	94.7	81.1	97.1	99.0	68.3	92.7	97.0	61.2	86.8	93.2	46.2	75.9	85.0
Swin-Base-224 + BERT-base, 7×7 patches																			
VSE++(Faghri et al. 2018)	✗	82.5	96.5	98.9	70.0	91.4	95.1	83.3	97.5	99.3	71.0	93.0	96.7	64.0	88.2	94.2	49.9	78.0	86.6
SCAN(Lee et al. 2018)	✓	79.0	95.9	98.2	67.7	90.6	94.9	80.9	97.0	99.1	69.7	93.1	97.1	60.7	86.6	93.2	48.1	77.1	86.1
SGR(Diao et al. 2021)	✓	80.4	97.0	98.7	66.9	90.2	94.5	81.2	97.1	99.1	69.9	93.2	97.2	61.0	86.7	93.2	48.6	77.2	86.3
CHAN(Pan et al. 2023)	✓	81.4	97.0	98.6	68.5	90.6	94.5	81.6	97.2	99.3	70.6	93.7	97.6	64.1	87.9	93.5	49.1	77.3	86.1
LAPS(Fu et al. 2024)	✓	82.4	97.4	99.5	70.0	91.7	95.4	84.0	97.6	99.3	72.1	93.7	97.3	64.5	89.2	94.4	51.6	78.9	87.2
Ours: AVSE	✗	83.9	97.4	99.4	70.0	92.4	95.6	84.9	98.0	99.3	72.1	94.0	97.4	66.2	89.8	94.7	51.7	79.2	87.3
Swin-Base-384 + BERT-base, 24×24 patches																			
VSE++(Faghri et al. 2018)	✗	83.3	97.5	99.2	71.1	93.2	96.2	82.9	97.7	99.4	71.3	93.5	97.3	63.0	88.5	94.3	50.1	78.9	87.4
SCAN(Lee et al. 2018)	✓	81.9	96.9	98.9	70.0	92.7	95.8	81.6	96.8	99.1	69.1	92.7	96.7	61.1	87.3	93.3	47.8	76.9	85.9
SGR(Diao et al. 2021)	✓	80.7	96.8	99.0	69.9	91.7	95.3	81.9	96.7	99.1	69.3	92.8	96.7	62.8	87.0	92.9	48.1	77.0	86.0
CHAN(Pan et al. 2023)	✓	81.2	96.7	98.8	70.3	92.2	95.9	83.1	97.3	99.2	70.4	93.1	97.1	63.4	88.4	94.1	49.2	77.9	86.6
LAPS(Fu et al. 2024)	✓	85.1	97.7	99.2	74.0	93.0	96.3	84.1	97.4	99.2	72.1	93.9	97.4	67.1	88.6	94.3	53.0	79.5	87.6
Ours: AVSE	✗	87.1	98.3	99.2	73.6	93.5	96.5	85.1	98.2	99.5	71.6	94.0	97.5	68.6	90.2	95.6	52.2	79.6	87.8

Table 1: Comparisons of experimental results on MS-COCO and Flickr30k datasets. “FG” indicates whether it is the fine-grained cross-modal alignment. The best performances are marked **bold**.

Objective Function

To align images and texts, we adopt hinge-based triplet loss (Kiros, Salakhutdinov, and Zemel 2014) using the hardest negative samples (Faghri et al. 2018; Lee et al. 2018; Chen et al. 2021). The loss function is defined as follows:

$$\mathcal{L}_m = \sum_{(I,T) \in \mathcal{D}} \left[\alpha - S(I, T) + S(I, \hat{T}) \right]^+ + \left[\alpha - S(I, T) + S(\hat{I}, T) \right]^+ \quad (7)$$

where α serves as a margin parameter and $[x]^+ \equiv \max(x, 0)$. In dataset \mathcal{D} , the visual semantic similarity in a positive pair

$S(I, T)$ should be higher than that in the hardest negative pairs $S(\hat{I}, T)$ and $S(I, \hat{T})$ by a margin α .

In summary, the final loss function is defined as follows to perform joint optimization of the two objectives.

$$\mathcal{L} = \mathcal{L}_m + \mathcal{L}_{reg} \quad (8)$$

Experiments

Dataset and Settings

Datasets. Following previous works (Faghri et al. 2018), we use two widely used benchmark datasets MS-COCO (Lin et al. 2014) and Flickr30K (Plummer et al. 2015) for our ex-

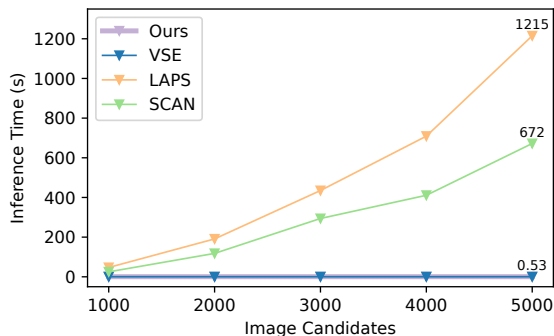


Figure 4: Inference time for image-text retrieval on GPU (lower the better). Our AVSE method is almost the same as VSE in calculating similarity speed, and much faster than the local-level matching method, especially when the number of images grows large.

periment. MS-COCO is a dataset that contains 123287 images and five text captions are annotated for each image. Following (Faghri et al. 2018), all data are split into training set, validation set and testing set which contains 113287, 5000, 5000 images respectively. Flickr30K is composed of 31783 images and each image has 5 corresponding descriptions. We follow the split in (Faghri et al. 2018), using 29000 images for training, 1000 images for validation, and 1000 images for testing.

Evaluation Metrics. The commonly used evaluation metrics for image-text matching are Recall@K (where $K = 1, 5, 10$), abbreviated as $R@1$, $R@5$, and $R@10$. These metrics represent the percentage of ground truth instances that are retrieved within the top 1, 5, or 10 results, respectively.

Implementation details

For feature extractions, we use the conventional region features and also use the recent popular Vision Transformer as the backbone, e.g., ViT (Dosovitskiy et al. 2020), Swin-Transformer (Liu et al. 2021b) we set the dimension of the shared embedding space d_1 as 512. For the asymmetric embedding optimal matching module, we set the dimension of blocks d_2 as 256. For objective function, we set λ_1 as $\frac{1}{d_1-1}$ to balance two terms of \mathcal{L}_{reg} and α is set to 0.2 as margin parameter. We train the proposed model for 25 epochs with set the mini-batch size as 128 using AdamW (Loshchilov and Hutter 2019) optimizer. The learning rate is set as 0.0005 for the first epochs and then decreased to 0.00005 for the rest 10 epochs. *We provide a more detailed implementation setup in the supplementary material.*

Comparisons with Recent State-of-the-arts

We compare AVSE with the most recent state-of-the-art approaches on two widely used datasets, *i.e.*, Flickr30K and MS-COCO. It is noted that many state-of-the-arts methods adopt the ensemble strategy (Lee et al. 2018; Li et al. 2019; Diao et al. 2021). For fair comparison, we also provide the ensemble version in results. We not only experiment with

Method	TR		IR	
	R@1	R@10	R@1	R@10
Asymmetric Embedding Optimal Matching				
w/o AEOM (1 views)	73.2	96.5	60.4	90.6
w/o AEOM (2 views)	74.0	96.9	59.9	91.5
w/o AEOM (4 views)	74.1	96.7	60.5	91.2
w/ AEOM (2 views)	76.0	97.5	62.7	93.1
w/ AEOM (4 views)	75.9	97.8	62.6	93.0
Dimension of Block d_2				
initial $d_2 = 64$	74.5	97.1	61.9	92.1
initial $d_2 = 128$	75.8	97.4	62.8	92.9
initial $d_2 = 256$	76.0	97.5	62.7	93.1
Sample Strategy d_2				
Uniform sampling	73.4	96.6	61.0	90.4
Gaussian sampling	74.9	97.1	61.8	91.2
Radial bias sampling	76.0	97.5	62.7	93.1
Dimension-wise Regularizing Loss				
w/o \mathcal{L}_{reg}	74.8	97.2	61.8	92.5
w/ \mathcal{L}_{reg}	76.0	97.5	62.7	93.1

Table 2: Ablation studies examining various parameters were conducted on the Flickr30K dataset using the ViT-base-224 model as the backbone.

the classic backbone, but in order to verify the potential of our model, we use the updated backbone to test the effect of our model. In addition, our AVSE is a general framework that is not only applicable to the latest Vision Transformer, but also to previous backbones. *More experimental results are in the supplementary materials.*

The experimental results of Flickr30K and the larger and complicated MS-COCO 5-fold 1K test set and MS-COCO 5K test set are shown in Table 1 AVSE outperforms all state-of-the-art methods by an impressive margin. Compared to the recent state-of-the-art method LAPS (Fu et al. 2024), our AVSE method outperforms it across nearly all metrics and backbones. For instance, on the Flickr30K dataset, our approach achieves a 2.0% improvement in $R@1$ for text retrieval when using the ViT-Base-224 backbone, and a 1.5% improvement in text retrieval when using the Swin-Base-224 backbone. Similar improvements are observed on the MS-COCO dataset.

Inference efficiency analysis. In evaluating the performance of models used in search engines with large-scale databases for image or text queries, the efficiency at the inference stage is of equal importance as the accuracy of caption or image retrieval. However, recent state-of-the-art method usually rely on complex cross-modal attention mechanism, which significantly harm the inference speed of the model. As shown in Figure 4, we compare inference time for image-text retrieval on a single GPU with recent state-of-the-art. We conduct the ablation study on MS-COCO 5K test set, where each image has 5 captions. It is obvious that global-level matching methods, VSE, is much faster than local-level matching methods (SCAN and LAPS). When the




Figure 5: Visualization of the radial bias sampling strategy, which can effectively extract multi-view information.

number of candidate images is small, typically 1000, the local-level matching method seems to be reasonably fast. Moreover, the time advantage of global-level matching increases with the number of candidate images.

Ablation Study

To verify the effectiveness of each component of our AVSE, we conduct extensive ablation studies on Flickr30K dataset. *There are also some ablation experiments on why we need to use multiple views for text in the supplementary material.*

The impact of AEOM. The AEOM is a critical module in our AVSE, to validate the superiority of AEOM, we first compare it with the conventional similarity functions, i.e., Cosine similarity, we also extract multi-view image features and feed it into a mean pooling to have multi-view information. As shown in Table 2, we list the results of 1,2,4 views of using Cosine similarity. It prove that multi-view information can directly improve the retrieval accuracy. And it is clearly to see that our AEOM achieves the best performance on all metrics when using 2 view image features. Although increasing views does not improve retrieval performance, our AEOM outperforms conventional methods by a large margin of 2.8% on R@1. At the same time, we also found that 4 views did not improve significantly, which may be due to the large amount of redundant information between too many views. Therefore, the more views, the better.

The impact of different block dimension. The dimension of the visual and text semantic blocks d_2 is a sensitive parameter in AEOM, which determines the ability to learn the visual-semantic similarities. Hence, an appropriate parameter is important in our proposed model. To validate the impact of the blocks' dimension d_2 , we conduct extensive experiments on Flickr30K dataset. Here, we study the matching performance by setting d_2 to 64, 128, and 256, and the results are reported in Table 2. Obviously, when the block dimension d_2 is set to 128 and 256, the model performs best. Since $d_2 = 256$ has less computation, we choose to set it to 256.

The impact of dimension-wise regularizing loss. We validate the positive effect of this dimension-wise regularizing loss on our proposed method. We can clearly find that by using \mathcal{L}_{reg} , the retrieval performance is improved in all metrics. Especially for R@1, our method obtain 1.2% relative improvement both for image retrieval and text retrieval.

Effective of Radial Bias Sampling. As shown in Table2, we compare the proposed Radial bias sampling with various classic sampling methods such as Gaussian sampling and uniform sampling. We find that the proposed Radial bias sampling has the best effect, while uniform sampling is too uniform, resulting in small differences between views, which is not conducive to the proposed AVSE framework. In addition, we also visualize the sampling effect of the Radia bias sampling module. As shown in Figure??, our Radia bias sampling can better mask the vicinity of the center point, which can be more conducive to our extraction of multi-view features.

Conclusion

In this paper, we propose a novel Asymmetric Visual Semantic Embedding (AVSE) for efficient image-text matching. The key insight is that the difference information density between vision and language is crucial for image-text retrieval. To better exploit the information density difference between the two modality data, let the model learn an asymmetric visual semantic embedding and make full use of the information density difference by a novel similarity learning module. Central to this module is the presumption of foundational semantic units within the embeddings, denoted as "meta-semantic embeddings". It segments all embeddings into meta-semantic embeddings with the same dimension and calculates visual semantic similarity by finding the optimal match of meta-semantic embeddings of two modalities. Comprehensive experiments on two widely-used benchmarks validate the effectiveness of the proposed method, leading to state-of-the-art performance.

Acknowledgments

This work was supported by Shenzhen Innovation in Science and Technology Foundation for The Excellent Youth Scholars (No. RCYX20231211090248064), National Natural Science Foundation of China (No. 62203476).

References

- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 6077–6086.
- Chen, H.; Ding, G.; Liu, X.; Lin, Z.; Liu, J.; and Han, J. 2020. IMRAM: Iterative Matching with Recurrent Attention Memory for Cross-Modal Image-Text Retrieval. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 12655–12663.
- Chen, J.; Hu, H.; Wu, H.; Jiang, Y.; and Wang, C. 2021. Learning the best pooling strategy for visual semantic embedding. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 15789–15798.
- Cheng, Y.; Zhu, X.; Qian, J.; Wen, F.; and Liu, P. 2022. Cross-modal Graph Matching Network for Image-text Retrieval. *ACM Trans. Multim. Comput. Commun. Appl.*, 18(4): 1–23.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Diao, H.; Zhang, Y.; Ma, L.; and Lu, H. 2021. Similarity reasoning and filtration for image-text matching. In *Proc. AAAI Conf. Artif. Intell.*, volume 35, 1218–1226.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Faghri, F.; Fleet, D. J.; Kiros, J. R.; and Fidler, S. 2018. Vse++: Improving visual-semantic embeddings with hard negatives. In *BMVC*.
- Fu, Z.; Mao, Z.; Song, Y.; and Zhang, Y. 2023. Learning Semantic Relationship Among Instances for Image-Text Matching. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 15159–15168.
- Fu, Z.; Zhang, L.; Xia, H.; and Mao, Z. 2024. Linguistic-Aware Patch Slimming Framework for Fine-grained Cross-Modal Alignment. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 26307–26316.
- Gong, X.; Yuan, D.; and Bao, W. 2021. Understanding partial multi-label learning via mutual information. *Advances in Neural Information Processing Systems*, 34: 4147–4156.
- Karpathy, A.; and Fei-Fei, L. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 3128–3137.
- Kiros, R.; Salakhutdinov, R.; and Zemel, R. S. 2014. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*.
- Lee, K.-H.; Chen, X.; Hua, G.; Hu, H.; and He, X. 2018. Stacked cross attention for image-text matching. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 201–216.
- Li, K.; Zhang, Y.; Li, K.; Li, Y.; and Fu, Y. 2019. Visual semantic reasoning for image-text matching. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 4654–4662.
- Li, Z.; Guo, C.; Feng, Z.; Hwang, J.-N.; and Xue, X. 2022. Multi-View Visual Semantic Embedding. In *Proc. Int. Joint Conf. on Artif. Intell.*
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 740–755.
- Liu, C.; Mao, Z.; Liu, A.-A.; Zhang, T.; Wang, B.; and Zhang, Y. 2019. Focus your attention: A bidirectional focal attention network for image-text matching. In *Proc. 27th ACM Int. Conf. Multimedia.*, 3–11.
- Liu, C.; Mao, Z.; Zhang, T.; Xie, H.; Wang, B.; and Zhang, Y. 2020. Graph structured network for image-text matching. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 10921–10930.
- Liu, W.; and Tsang, I. 2015. On the optimality of classifier chain for multi-label classification. *Advances in Neural Information Processing Systems*, 28.
- Liu, Y.; Chen, C.; Wang, C.; King, X.; and Liu, M. 2023. Regress before construct: Regress autoencoder for point cloud self-supervised learning. In *Proceedings of the 31st ACM International Conference on Multimedia*, 1738–1749.
- Liu, Y.; Liu, H.; Wang, H.; and Liu, M. 2022. Regularizing Visual Semantic Embedding with Contrastive Learning for Image-Text Matching. *IEEE Sign. Process. Letters*.
- Liu, Y.; Wang, H.; Meng, F.; Liu, M.; and Liu, H. 2021a. Attend, Correct And Focus: A Bidirectional Correct Attention Network For Image-Text Matching. In *Proc. IEEE Int. Conf. on Image Process.*, 2673–2677.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 10012–10022.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled weight decay regularization. In *Proc. Int. Conf. on Learn. Representations*.
- Pan, Z.; Wu, F.; and Zhang, B. 2023. Fine-Grained Image-Text Matching by Cross-Modal Hard Aligning Network. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 19275–19284.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2641–2649.
- Qin, Y.; Peng, D.; Peng, X.; Wang, X.; and Hu, P. 2022. Deep evidential learning with noisy correspondence for

- cross-modal retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4948–4956.
- Qin, Y.; Sun, Y.; Peng, D.; Zhou, J. T.; Peng, X.; and Hu, P. 2023. Cross-modal Active Complementary Learning with Self-refining Correspondence. *arXiv preprint arXiv:2310.17468*.
- Vendrov, I.; Kiros, R.; Fidler, S.; and Urtasun, R. 2016. Order-embeddings of images and language. In *Proc. Int. Conf. on Learn. Representations*.
- Wang, H.; He, D.; Wu, W.; Xia, B.; Yang, M.; Li, F.; Yu, Y.; Ji, Z.; Ding, E.; and Wang, J. 2022. CODER: Coupled Diversity-Sensitive Momentum Contrastive Learning for Image-Text Retrieval. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*.
- Wang, H.; Zhang, Y.; Ji, Z.; Pang, Y.; and Ma, L. 2020. Consensus-aware visual-semantic embedding for image-text matching. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 18–34.
- Wang, Y.; Yang, H.; Qian, X.; Ma, L.; Lu, J.; Li, B.; and Fan, X. 2019. Position Focused Attention Network for Image-Text Matching. In *Proc. Int. Joint Conf. on Artif. Intell.*, 3792–3798.
- Wu, Y.; Wang, S.; Song, G.; and Huang, Q. 2019. Learning fragment self-attention embeddings for image-text matching. In *Proc. 27th ACM Int. Conf. Multimedia.*, 2088–2096.
- Zhang, H.; Mao, Z.; Zhang, K.; and Zhang, Y. 2022a. Show Your Faith: Cross-Modal Confidence-Aware Network for Image-Text Matching. In *Proc. AAAI Conf. Artif. Intell.*
- Zhang, K.; Mao, Z.; Wang, Q.; and Zhang, Y. 2022b. Negative-Aware Attention Framework for Image-Text Matching. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 15661–15670.