







HADA: A Graph-Based Amalgamation Framework in Image-Text Retrieval

Manh-Duy Nguyen¹  , Binh T. Nguyen^{2,3,4} , and Cathal Gurrin¹ 

¹ School of Computing, Dublin City University, Dublin, Ireland
manh.nguyen5@mail.dcu.ie

² VNU-HCM, University of Science, Ho Chi Minh City, Vietnam

³ Vietnam National University, Ho Chi Minh City, Vietnam

⁴ AISIA Lab, Ho Chi Minh City, Vietnam

Abstract. Many models have been proposed for vision and language tasks, especially the image-text retrieval task. State-of-the-art (SOTA) models in this challenge contain hundreds of millions of parameters. They also were pretrained on large external datasets that have been proven to significantly improve overall performance. However, it is not easy to propose a new model with a novel architecture and intensively train it on a massive dataset with many GPUs to surpass many SOTA models already available to use on the Internet. In this paper, we propose a compact graph-based framework named HADA, which can combine pretrained models to produce a better result rather than starting from scratch. Firstly, we created a graph structure in which the nodes were the features extracted from the pretrained models and the edges connecting them. The graph structure was employed to capture and fuse the information from every pretrained model. Then a graph neural network was applied to update the connection between the nodes to get the representative embedding vector for an image and text. Finally, we employed cosine similarity to match images with their relevant texts and vice versa to ensure a low inference time. Our experiments show that, although HADA contained a tiny number of trainable parameters, it could increase baseline performance by more than 3.6% in terms of evaluation metrics on the Flickr30k dataset. Additionally, the proposed model did not train on any external dataset and only required a single GPU to train due to the small number of parameters required. The source code is available at <https://github.com/m2man/HADA>.

Keywords: Image-text retrieval · Graph neural network · Fusion model

1 Introduction

Image-text retrieval is one of the most popular challenges in vision and language tasks, with many state-of-the-art (SOTA) models recently introduced [3, 10, 17–19, 25, 28]. This challenge includes two subtasks, which are image-to-text retrieval and text-to-image retrieval. The former subtask utilises an image

query to retrieve relevant texts in a multimodal dataset, while the latter is concerned with text queries for ranked videos.

Most of the SOTA models in this research field share two things in common: (1) they were built on transformer-based cross-modality attention architectures [3, 19] and (2) they were pretrained on the large-scale multimodal data crawled from the Internet [13, 17–19, 28]. However, these things have their own disadvantages. The attention structure between two modalities could achieve an accurate result, but it costs a large amount of inference time due to the massive computation required. For instance, UNITER [3] contained roughly 303 million parameters, and it took a decent amount of time (more than 12s for each query on a dataset with 30000 images [31]) to perform the retrieval in real-time. Many recent works have resolved this model-related problem by introducing joint-encoding learning methods. They can learn visual and semantic information from both modalities without using any cross-attention modules, which can be applied later to rerank the initial result [18, 25, 31]. Figure 1 illustrates the architecture of these pipelines. Regarding the data perspective, the large collected data usually comes with noisy annotations, which could impact on the models trained on it. Several techniques have been proposed to mitigate this issue [17–19]. However, training on a massive dataset still burdens computation, such as the number of GPUs required to train the model successfully and efficiently [28].

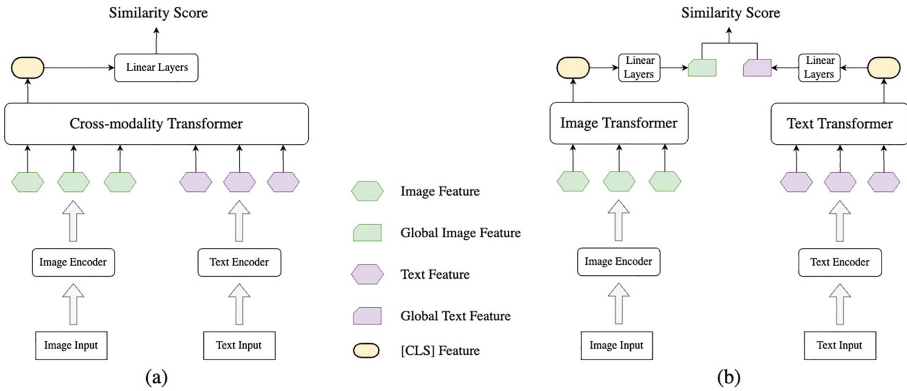


Fig. 1. Two most popular pipelines of the SOTA for image-text retrieval challenge. (a) A cross-modality transformer network is applied to measure the similarity between an image and a text based on their features. (b) Each modality used their own transformer network to get its global embedding.

It has motivated us to answer the question: *Can we combine many SOTA models, which are currently available to use, to get a better-unified model without intensive training using many GPUs?* In this paper, we introduce a graph-based amalgamation framework, called **HADA**, which utilises a graph-based structure

to fuse the features produced by other pretrained models. We did not use any time-consuming cross-modality attention network to ensure fast retrieval speed. A graph neural network was employed to extract visual and textual embedded vectors from fused graph-based structures of images and texts, where we can measure their cosine similarity. To the best of our knowledge, the graph structure has been widely applied in the image-text retrieval challenge [7, 21, 26, 27, 35]. Nevertheless, it was utilized to capture the interaction between objects or align local and global information within images. HADA is the first approach that applies this data structure to combine SOTA pretrained models by fusing their features in each modality. We trained HADA on the Flickr30k dataset without using any large-scale datasets. Then, we applied the Momentum Distillation technique [18], which has been shown to mitigate not only the harmful effect of noise annotation, but also improve accuracy on a clean dataset. Our experiments showed that HADA, with the tiny extra number of training parameters, could improve total recall by 3.64% compared to the input SOTA, without training with millions of additional image-text pairs as other models require. This is the most crucial contribution since it is expensive to utilise multiple GPU, especially for small and medium businesses or start-up companies. Therefore, we believe that HADA can be applied in both academic and industrial domains.

Our main contribution can be summarised as follows: (1) We introduced HADA, a compact pipeline that can combine two or many SOTA pretrained models to address the image-text retrieval challenge. (2) We proposed a way to fuse the information between input pretrained models by using graph structures. (3) We evaluated the performance of HADA on the well-known Flickr30k dataset [37] and MSCOCO dataset [20] without using any other large-scale dataset but still improved the accuracy compared to the baseline input models.

2 Related Work

A typical vision-and-language model, including an image-text retrieval task, was built using transformer-based encoders. In specific, OSCAR [19], UNITER [3], and VILLA [10] firstly employed Faster-RCNN [29], and BERT [6] to extract visual and text features from images and texts. These features were then fed into a cross-modality transformer block to learn the contextualized embedding that captured the relations between regional features from images and word pieces from texts. An additional fully connected layer was used to classify whether the images and texts were relevant to each other or not based on the embedding vectors. Although achieving superior results, these approaches had a drawback in applying them to real-time use cases. It required a huge amount of time to perform the online retrieval, since models had to process the intensive cross-attention transformer architecture many times for each query [31].

Recently, some works have proposed an approach to resolve that problem by utilizing two distinct encoders for images and text. Data from each modality can now be embedded offline and hence improve retrieval speed [13, 17, 18, 25, 28, 31]. In terms of architecture, all approaches used the similar BERT-based encoder for

semantic data but different image encoders. While LightningDOT [31] encoded images with detected objects extracted by the Faster-RCNN model, FastnSlow [25] applied the conventional Resnet network to embed images. On the other side, ALBEF [18] and BLIP [17] employed the Vision Transformer backbone [8] to get the visual features corresponding to their patches. Because these SOTA did not use the cross-attention structure, which was a critical point to achieve high accuracy, they applied different strategies to increase performance. Specifically, pretraining a model on a large dataset can significantly improve the result [13, 18, 19]. For instance, CLIP [28] and ALIGN [13] were pretrained on 400 million and 1.8 billion image-text pairs, respectively. Another way was that they ran another cross-modality image-text retrieval model to rerank the initial output and get a more accurate result [18, 31].

Regarding graph structures, SGM [35] introduced a visual graph encoder and a textual graph encoder to capture the interaction between objects appearing in images and between the entities in text. LGSGM [26] proposed a graph embedding network on top of SGM to learn both local and global information about the graphs. Similarly, GSMN [21] presented a novel technique to assess the correspondence of nodes and edges of graphs extracted from images and texts separately. SGRAF [7] built a reasoning and filtration graph network to refine and remove irrelevant interactions between objects in both modalities.

Although there are many SOTAs with different approaches for image-text retrieval problems, there is no work that tries combining these models, rather they introduce a new architecture and pretrain on massive datasets instead. Training an entirely new model from scratch on the dataset is a challenging task since it will create a burden on the computation facilities such as GPUs. In this paper, we introduced a simple method that combined the features extracted from the pretrained SOTA by applying graph structures. Unlike other methods that also used this data structure, we employed graphs to fuse the information between the input features, which was then fed into a conventional graph neural network to obtain the embedding for each modality. Our HADA consisted of a small number of trainable parameters, hence can be easily trained on a small dataset but still obtained higher results than the input models.

3 Methodology

This section will describe how our HADA addressed the retrieval challenge by combining any available pretrained models. Figure 2 depicted the workflow of HADA. We started with only two models ($N_{models} = 2$) as illustrated in Fig. 2 for simplicity. Nevertheless, HADA can be extended with a larger N_{models} . HADA began using some pretrained models to extract the features from each modality. We then built a graph structure to connect the extracted features together, which were fed into a graph neural network (GNN) later to update them. The outputs of the GNN were concatenated with the original global features produced by the pretrained models. Finally, simple linear layers were employed to get the final

representation embedding features for images and texts, which can be used to measure similarity to perform the retrieval. For evaluation, we could extract our representation features offline to guarantee high-speed inference time.

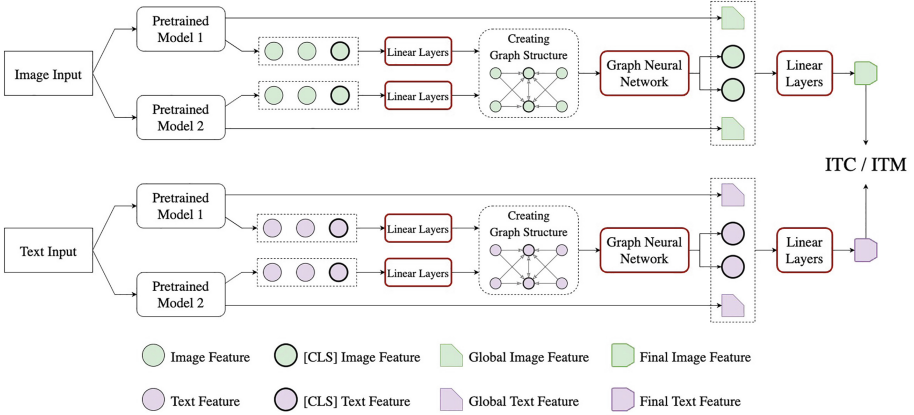


Fig. 2. The pipeline of the proposed HADA. The red borders indicated trainable components. The ITM and ITC inferred the training tasks which will be discussed later. (Color figure online)

3.1 Revisit State-of-the-Art Models

We only used the pretrained models without using the cross-modality transformer structure to extract features as depicted in Fig. 1 in order to reduce the number of computations and ensure the high-speed inference time. Basically, they used a unimodal encoder to get the features of an image or a text followed by a transformer network to embed them and obtain the $[CLS]$ embedding. This $[CLS]$ token was updated by one or many fully connected layers to become a representative global feature that can be compared with that of the remaining modality to get the similarity score.

HADA began with the output of the transformer layer from the pretrained models. In detail, for an input image \mathbf{I} , we obtained the sequence of patch tokens from each model i denoted as $\mathbf{v}^{(i)} = \{v_{cls}^{(i)}, v_1^{(i)}, v_2^{(i)}, \dots, v_{N_i}^{(i)}\}$, where $v_j^{(i)} \in \mathbb{R}^{d_v^{(i)}}$ and N_i was the length of the sequence. This length depended on the architecture of the image encoder network employed in the pretrained model. For example, it could be the number of patches if the image encoder was a Vision Transformer (ViT) network [8], or the number of detected objects or regions of interest if the encoder was a Faster-RCNN model [29]. Additionally, we also extracted the global visual representation feature $h_v^{(i)} \in \mathbb{R}^{d_h^{(i)}}$ from $v_{cls}^{(i)}$ as illustrated in Fig. 1. Regarding the semantic modality, we used the same process as that of the visual modality. Specifically, we extracted the sequence of patch tokens $\mathbf{w}^{(i)} = \{w_{cls}^{(i)}, w_1^{(i)}, w_2^{(i)}, \dots, w_L^{(i)}\}$ where $w_j^{(i)} \in \mathbb{R}^{d_w^{(i)}}$ and L was the length of the text,

and the global textual representation embedding $h_w^{(i)} \in \mathbb{R}^{d_h^{(i)}}$ for an input text \mathbf{T} using the pretrained model i . The input model i matched a pair of an image \mathbf{I} and a text \mathbf{T} by calculating the dot product $\langle h_v^{(i)}, h_w^{(i)} \rangle$ of their global features. However, HADA used not only the global embedding but also the intermediate transformer tokens to make the prediction. We used our learned $[CLS]$ tokens to improve the global features. In contrast, using the original global features could ensure high performance of the pretrained models and mitigate the effect of unhelpful tokens.

3.2 Create Graph Structure

Each pretrained model i produced different $[CLS]$ features $v_{cls}^{(i)}$ and $w_{cls}^{(i)}$ for an image and text, respectively. Since our purpose was to combine the models, we needed to fuse these $[CLS]$ tokens to obtain the unified ones for each modality separately. In each modality, for example, the visual modality, HADA not only updated $v_{cls}^{(i)}$ based on $\mathbf{v}^{(i)}$ solely but also on those of the remaining pretrained models $\{\mathbf{v}^{(j)} \mid j \neq i\}$. Because these \mathbf{v} came from different models, their dimensions would not be similar to each other. Therefore, we applied a list of linear layers $f_v^{(i)} : \mathbb{R}^{d_v^{(i)}} \rightarrow \mathbb{R}^{d_p}$ to map them in the same dimensional space:

$$\mathbf{p}^{(i)} = \{f_v^{(i)}(x) \mid x \in \mathbf{v}^{(i)}\} = \{p_{cls}^{(i)}, p_1^{(i)}, p_2^{(i)}, \dots, p_{N_i}^{(i)}\}$$

We performed a similar process for the textual modality to obtain:

$$\mathbf{s}^{(i)} = \{f_w^{(i)}(x) \mid x \in \mathbf{w}^{(i)}\} = \{s_{cls}^{(i)}, s_1^{(i)}, s_2^{(i)}, \dots, s_L^{(i)}\}, \text{ where } f_w^{(i)} : \mathbb{R}^{d_w^{(i)}} \rightarrow \mathbb{R}^{d_s}$$

We then used graph structures $\mathcal{G}_p = \{\mathcal{V}_p, \mathcal{E}_p\}$ and $\mathcal{G}_s = \{\mathcal{V}_s, \mathcal{E}_s\}$ to connect these mapped features together, where \mathcal{V} and \mathcal{E} denoted the list of nodes and edges in the graph \mathcal{G} accordingly. In our HADA, nodes indicated the mapped features. Specifically, $\mathcal{V}_p = \{\mathbf{p}^{(i)}\}$ and $\mathcal{V}_s = \{\mathbf{s}^{(i)}\}$ for all $i \in [1, N_{models}]$. Regarding edges, we symbolized $e_{a \rightarrow b}$ as a directed edge from node a to node b in the graph, thus the set of edges of the visual graph \mathcal{E}_p and the textual graph \mathcal{E}_s were:

$$\mathcal{E}_p = \{e_{x \rightarrow p_{cls}^{(j)}} \mid x \in \mathbf{p}^{(i)} \text{ and } i, j \in [1, N_{models}]\}$$

$$\mathcal{E}_s = \{e_{x \rightarrow s_{cls}^{(j)}} \mid x \in \mathbf{s}^{(i)} \text{ and } i, j \in [1, N_{models}]\}$$

To be more detailed, we created directed edges that went from every patch feature to the $[CLS]$ feature, including from the $[CLS]$ itself, for all pretrained models but not in the reverse direction, as shown in Fig. 2. The reason was that $[CLS]$ was originally introduced as a representation of all input data, so it would summarize all patch tokens [2, 6, 8]. Therefore, it would be the node that received information from other nodes in the graph. This connection structure ensured that HADA could update the $[CLS]$ tokens based on the patch tokens from all pretrained models in a fine-grained manner.

3.3 Graph Neural Network

Graph neural networks (GNN) have witnessed an increase in popularity over the past few years, with many GNN structures having been introduced recently [1, 5, 11, 15, 30, 34]. HADA applied the modified Graph Attention Network (GATv2), which was recommended to be used as a baseline whenever employing a GNN [1], to fuse the patch features from different pretrained models together to get the unified $[CLS]$ features. Let $\mathcal{N}_k = \{x \in \mathcal{V} \mid e_{x \rightarrow k} \in \mathcal{E}\}$ be the set of neighbor nodes from which there was an edge connecting to node k in the graph \mathcal{G} . GATv2 used a scoring function se to weight every edge indicating the importance of the neighbor nodes x in \mathcal{N}_k before updating the node $k \in \mathbb{R}^d$:

$$se(e_{x \rightarrow k}) = \mathbf{A}^\top \text{LeakyRELU}(\mathbf{W}_1 x + \mathbf{W}_2 k)$$

where $\mathbf{A} \in \mathbb{R}^{d'}$, $\mathbf{W}_1 \in \mathbb{R}^{d' \times d}$, and $\mathbf{W}_2 \in \mathbb{R}^{d' \times d}$ were learnable parameters. These weights were then normalized across all neighbor nodes in \mathcal{N}_k by using a softmax function to get the attention scores:

$$\alpha_{e_{x \rightarrow k}} = \frac{\exp(se(e_{x \rightarrow k}))}{\sum_{y \in \mathcal{N}_k} \exp(se(e_{y \rightarrow k}))}$$

The updated node $k' \in \mathbb{R}^{d'}$ was then calculated based on its neighbors in \mathcal{N}_k , including k if we add an edge connect it to itself:

$$k' = \sigma\left(\sum_{x \in \mathcal{N}_k} \alpha_{e_{x \rightarrow k}} \cdot \mathbf{W}_1 x\right),$$

where σ was a nonlinear activate function. Furthermore, this GATv2 network could be enlarged by applying a multi-head attention structure, and improved performance [34]. The output now was a concatenation of each head output, which was similar to Transformer architecture [33]. An extra linear layer was used at the end to convert these concatenated nodes to the desired dimensions.

We used distinct GATv2 structures with H attention heads for each modality in this stage, as illustrated in Fig. 2. HADA took the input graphs \mathcal{G}_p and \mathcal{G}_s with nodes \mathcal{V}_p and \mathcal{V}_s in the vector space of d_p and d_s dimensions and updated them to $\mathcal{V}'_p = \{\mathbf{p}'^{(i)}\}$ and $\mathcal{V}'_s = \{\mathbf{s}'^{(i)}\}$ with dimensions of d'_p and d'_s . We then concatenated the updated $[CLS]$ nodes p'_{cls} and s'_{cls} from all pretrained models with their corresponding original global embedding h_v and h_w . Finally, we fed them into a list of linear layers to get our normalized global representation $h_p \in \mathbb{R}^{d_h}$ and $h_s \in \mathbb{R}^{d_h}$.

3.4 Training Tasks

Image-Text Contrastive Learning. HADA encoded the input image \mathbf{I} and text \mathbf{T} to h_p and h_s , accordingly. We used a similarity function that was a dot product $S(\mathbf{I}, \mathbf{T}) = \langle h_p, h_s \rangle = h_p^\top h_s$ to ensure that a pair of relevant image-text (positive pair) would have a higher similar representation compared to irrelevant

pairs (negative pairs). The contrastive loss for image-to-text (i2t) retrieval and text-to-image (t2i) retrieval for the mini-batch of M relevant pairs $(\mathbf{I}_m, \mathbf{T}_m)$ were:

$$\mathcal{L}_{i2t}(\mathbf{I}_m) = -\log \frac{\exp(S(\mathbf{I}_m, \mathbf{T}_m)/\tau)}{\sum_{i=1}^M \exp(S(\mathbf{I}_m, \mathbf{T}_i)/\tau)}$$

$$\mathcal{L}_{t2i}(\mathbf{T}_m) = -\log \frac{\exp(S(\mathbf{T}_m, \mathbf{I}_m)/\tau)}{\sum_{i=1}^M \exp(S(\mathbf{T}_m, \mathbf{I}_i)/\tau)}$$

where τ was a temperature parameter that could be learned during training. Such contrastive learning has been used in many vision-and-language models and has been proven to be effective [17, 18, 28, 31]. In our experiment, we trained HADA with the loss that optimized both subtasks:

$$\mathcal{L}_{ITC} = \frac{1}{M} \sum_{m=1}^M (\mathcal{L}_{i2t}(\mathbf{I}_m) + \mathcal{L}_{t2i}(\mathbf{T}_m))$$

Inspired by ALBEF [18], we also applied momentum contrast (MoCo) [12] and their momentum distillation strategy for this unsupervised representation learning to cope with the problem of noisy information in the dataset and improve accuracy.

Image-Text Matching. This objective was a binary classification task to distinguish irrelevant image-text pairs that had similar representations. This task would ensure that they were different in fine-grained details. We implemented an additional discriminator layer $dc : \mathbb{R}^{4d_h} \rightarrow \mathbb{R}$ on top of the final embedding features h_p and h_s to classify whether the image \mathbf{I} and the text \mathbf{T} is a positive pair or not:

$$dc(h_p, h_s) = \text{sigmoid}(\mathbf{W}^\top [h_p \| h_s \| \text{abs}(h_p - h_s) \| h_p \odot h_s])$$

where $\mathbf{W} \in \mathbb{R}^{4d_h}$ was trainable parameters, $\|$ indicated the concatenation, $\text{abs}(\cdot)$ was the absolute value, and \odot denoted elementwise multiplication. We used binary cross-entropy loss for this ordinary classification task:

$$\mathcal{L}_{itm}(\mathbf{I}, \mathbf{T}) = y \log(dc(h_p, d_s)) + (1 - y) \log(1 - dc(h_p, d_s))$$

where y was the one-hot vector representing the ground truth label of the pair.

For each positive pair in the minibatch of M positive pairs, we sampled 1 hard negative text for the image and 1 hard negative image for the text. These negative samples were chosen from the current mini-batch in which they were not relevant based on the ground-truth labels, but have the highest similarity dot product score. Therefore, the objective for this task was:

$$\mathcal{L}_{ITM} = \frac{1}{3M} \sum_{m=1}^M (\mathcal{L}_{itm}(\mathbf{I}_m, \mathbf{T}_m) + \mathcal{L}_{itm}(\mathbf{I}_m, \mathbf{T}'_m) + \mathcal{L}_{itm}(\mathbf{I}'_m, \mathbf{T}_m))$$

where \mathbf{T}'_m and \mathbf{I}'_m were the hard negative text and image samples in the mini-batch that were corresponding with the \mathbf{I}_m and \mathbf{T}_m , respectively. The final loss function in HADA was:

$$\mathcal{L} = \mathcal{L}_{ITC} + \mathcal{L}_{ITM}$$

4 Experiment

4.1 Dataset and Evaluation Metrics

We trained and evaluated HADA on two different common datasets in the image-text retrieval task, which are Flickr30k [37] and MSCOCO [20]. The Flickr30k dataset consists of 31K images collected on the Flickr website, while MSCOCO comprises 123K images. Each image contains five relevant texts or captions that describe the image. We used Karpathy’s split [14], which has been widely applied by all models in the image-text retrieval task, to split each dataset into train/evaluate/test on 29K/1K/1K and 113K/5K/5K images on Flickr30k and MSCOCO, respectively.

The common evaluation metric in this task was the Recall at K ($R@K$) because many SOTA works used this metric [3, 10, 13, 17–19, 28, 31]. This metric scores the proportion of the number of queries that we found the correct relevant output in the top K of the retrieved ranked list:

$$R@K = \frac{1}{N_q} \sum_{q=1}^{N_q} \mathbf{1}(q, K)$$

where N_q is the number of queries and $\mathbf{1}(q, K)$ is a binary function returning 1 if the model finds the correct answer of the query q in the top K of the retrieved output. In particular, for the image-to-text subtask, $R@K$ is the percentage of the number of images where we found relevant texts in the top K of the output result. In our experiment, we used $R@1$, $R@5$, $R@10$, and RSum, which was the sum of them.

4.2 Implementation Details

In our experiment, we combined two SOTA models that had available pretrained weights fine-tuned on the Flickr30k dataset: ALBEF¹ and LightningDOT². None of them used the cross-modality transformer structure when retrieved to ensure the fast inference speed³. Although they used the same BERT architecture to encode a text, the former model employed the ViT network to encode an image, while the latter model applied the Faster-RCNN model. We chose these two

¹ <https://github.com/salesforce/ALBEF>.

² <https://github.com/intersun/LightningDOT>.

³ Indeed, these two models applied the cross-modality transformer network to rerank the initial result in the subsequent step. However, we did not focus on this stage.

models because we wanted to combine different models with distinct embedding backbones to utilize the advantages of each of them.

Regarding ALBEF, their ViT network encoded an image to 577 patch tokens including the $[CLS]$ one ($N_{ALB} = 576$ and $d_v^{(ALB)} = 768$). This $[CLS]$ was projected to the lower dimension to obtain the global feature ($d_h^{(ALB)} = 256$). Because LightningDOT encoded an image based on the detected objects produced by the Faster-RCNN model, its N_{DOT} varied depending on the number of objects in the image. The graph neural network, unlike other conventional CNNs, can address this inconsistent number of inputs due to the flexible graph structure with nodes and edges. Unlike ALBEF, the dimensions of image features and global features from LightningDOT were the same with $d_v^{(DOT)} = d_h^{(DOT)} = 768$. In terms of text encoder, the output of both models was similar since they used the same BERT network: $d_w^{(ALB)} = d_w^{(DOT)} = 768$. We projected these features to a latent space where $d_p = d_s = 512$, which was the average of their original dimensions. We used a 1-layer GATv2 network with $H = 4$ multi-head attention to update the graph features while still keeping the input dimensions of $d'_p = d'_s = 512$. We also applied Dropout with $p = 0.7$ in linear layers and graph neural networks. In total, our HADA contained roughly 10M trainable parameters.

The input pretrained models were pretrained on several large external datasets. For example, ALBEF was pretrained on 14M images compared to only 29K images on Flickr30k that we used to train HADA. We used this advantage in our prediction instead of training HADA in millions of samples. We modified the similarity score to a weighted sum of our predictions and the original prediction of the input models. Therefore, the weighted similarity score that we used was:

$$S(\mathbf{I}, \mathbf{T}) = (1 - \alpha)\langle h_p, h_s \rangle + \alpha\langle h_v^{(ALB)}, h_w^{(ALB)} \rangle$$

where α was a trainable parameter. We did not include the original result of the LightningDOT model since its result was lower than ALBEF by a large margin and, therefore, could have a negative impact on overall performance⁴.

We trained HADA for 50 epochs (early stopping⁵ was implemented) using the batch size of 20 on one NVIDIA RTX3080Ti GPU. We used the AdamW [23] optimizer with a weight decay of 0.02. The learning rate was set at $1e^{-4}$ and decayed to $5e^{-6}$ following cosine annealing [22]. Similarly to ALBEF, we also applied RandAugment [4] for data augmentation. The initial temperature parameter was 0.07 [36], and we kept it in the range of [0.001, 0.5] during training. To mitigate the dominant effect of ALBEF global features on our weighted similarity score, we first trained HADA with $\alpha = 0$. After the model had converged, we continued to train but initially set $\alpha = 0.5$ and kept it in the range of [0.1, 0.9].

⁴ We tried including the LightningDOT in the weighted similarity score, but the result was lower than using only ALBEF.

⁵ In our experiment, it converged after roughly 20 epochs.

4.3 Baselines

We built two baselines that also integrated ALBEF and LightningDOT as input to show the advantages of using graph structures to fuse these input models.

Baseline B1. We calculated the average of the original ranking results obtained from ALBEF and LightningDOT and considered them as the distance between images and text. This meant that the relevant pairs should be ranked at the top, whilst irrelevant pairs would rank lower.

Baseline B2. Instead of using a graph structure to fuse the features extracted from the pretrained models, we only concatenated their global embedding and fed them into the last linear layers to obtain the unified features. We trained this baseline B2 following the same strategy as described in Sect. 4.2 using the weighted similarity score.

4.4 Comparison to Baseline

Table 1 illustrated the evaluation metrics of the different models in the Flickr30k dataset. Similarly to LightningDOT, our main target was to introduce an image-text retrieval model that did not implement a cross-modality transformer module to ensure that it can perform in real time without any delay. Thus, we only reported the result from LightningDOT and ALBEF that did not use the time-consuming compartment to rerank in the subsequent step. If the model has a better initial result, it can have a better-reranked result by using the cross-modality transformer later. We also added UNITER [3], and VILLA [10] to our comparison. These approaches both applied cross-modality transformer architecture.

Table 1. Performance of models on Flickr30k Dataset. The symbol † indicated the results were originally reported in their research, while others were from our re-implementation using their public pretrained checkpoints. The column ΔR showed the difference compared to ALBEF.

Methods	Image-to-Text				Text-to-Image				Total	ΔR
	R@1	R@5	R@10	RSum	R@1	R@5	R@10	RSum	RSum	
UNITER†	87.3	98	99.2	284.5	75.56	94.08	96.76	266.4	550.9	↓13.68
VILLA†	87.9	97.2	98.8	283.9	76.26	94.24	96.84	267.34	551.24	↓13.34
LightningDOT	83.6	96	98.2	277.8	69.2	90.72	94.54	254.46	532.26	↓32.32
LightningDOT†	83.9	97.2	98.6	279.7	69.9	91.1	95.2	256.2	535.9	↓28.68
ALBEF	92.6	99.3	99.9	291.8	79.76	95.3	97.72	272.78	564.58	0
B1	90.7	99	99.6	289.3	79.08	94.5	96.94	270.52	559.82	↓4.76
B2	91.4	99.5	99.7	290.6	79.64	95.34	97.46	272.44	563.04	↓1.54
HADA	93.3	99.6	100	292.9	81.36	95.94	98.02	275.32	568.22	↑3.64

It was clear that our HADA obtained the highest metrics at all recall values compared to others. HADA achieved a slightly better R@5 and R@10 in Image-to-Text (I2T) and Text-to-Image (T2I) subtasks than ALBEF. However, the gap became more significant at R@1. We improved the R@1 of I2T by 0.7% (92.96 \rightarrow 93.3) and the R@1 of T2I by 1.6% (79.76 \rightarrow 81.36). In total, our RSum was 3.64% higher than that of ALBEF (564.58 \rightarrow 568.22).

The experiment also showed that LightningDOT, which encoded images using Faster-RCNN, performed worse than ALBEF when its total RSum was lower than that of ALBEF by approximately 30%. The reason might be that the object detector was not as powerful as the ViT network, and LightningDOT was pretrained on 4M images compared to 14M images used to train ALBEF. Although also using object detectors as the backbone but applying a cross-modality network, UNITER and VILLA surpassed LightningDOT by a large margin at 15%. It proved that this intensive architecture made a large impact on multimodal retrieval.

Regarding our two baselines, B1 and B2, both of them failed to get better results than the input model ALBEF. Model B1, using the simple strategy of taking the average ranking results and having no learnable parameters, performed worse than model B2, which used a trainable linear layer to fuse the pretrained features. Nevertheless, the RSum of B2 was lower than HADA by 5.18%. It showed the advantages of using a graph structure to fuse the information between models to obtain a better result.

4.5 HADA with Other Input Models

To show the stable performance of HADA, we used it to combine two other different pretrained models, including BLIP [17] and CLIP [28]. While CLIP is well-known for its application in many retrieval challenges [9, 24, 31, 32], BLIP is the enhanced version of ALBEF with the bootstrapping technique in the training process. We used the same configuration as described in 4.2 to train and evaluate HADA in Flickr30k and MSCOCO datasets. We used the pretrained BLIP and CLIP from the LAVIS library [16]. It was noted that the CLIP we used in this experiment was the zero-shot model since the fine-tuned CLIP for these datasets is not available yet.

Table 2. Performance of models on the test set in Flickr30k and MSCOCO datasets. The column ΔR showed the difference compared to BLIP in that dataset.

Dataset	Methods	Image-to-Text				Text-to-Image				Total	ΔR
		R@1	R@5	R@10	RSum	R@1	R@5	R@10	RSum	RSum	
Flickr30k	BLIP	94.3	99.5	99.9	293.7	83.54	96.66	98.32	278.52	572.22	0
	CLIP	88	98.7	99.4	286.1	68.7	90.6	95.2	254.5	540.6	↓31.62
	HADA	95.2	99.7	100	294.9	85.3	97.24	98.72	281.26	576.16	↑3.94
MSCOCO	BLIP	75.76	93.8	96.62	266.18	57.32	81.84	88.92	228.08	494.26	0
	CLIP	57.84	81.22	87.78	226.84	37.02	61.66	71.5	170.18	397.02	↓97.24
	HADA	75.36	92.98	96.44	264.78	58.46	82.85	89.66	230.97	495.75	↑1.49

Table 2 showed the comparison between HADA and the input models. CLIP performed worst on both Flickr30k and MSCOCO with huge differences compared to BLIP and HADA because CLIP was not fine-tuned for these datasets. Regarding the Flickr30k dataset, HADA managed to improve the RSum by more than 3.9% compared to that of BLIP. Additionally, HADA obtained the highest scores in all metrics for both subtasks. Our proposed framework also increased the RSum of BLIP by 1.49% in the MSCOCO dataset. However, BLIP performed slightly better than HADA in the I2T subtask, while HADA achieved higher performance in the T2I subtask.

5 Conclusion

In this research, we proposed a simple graph-based framework, called HADA, to combine two pretrained models to address the image-text retrieval problem. We created a graph structure to fuse the extracted features obtained from the pretrained models, followed by the GATv2 network to update them. Our proposed HADA only contained roughly 10M learnable parameters, helping it become easy to train using only one GPU. Our experiments showed the promise of the proposed method. Compared to input models, we managed to increase total recall by more than 3.6%. Additionally, we implemented two other simple baselines to show the advantage of using the graph structures. This result helped us to make two contributions: (1) to increase the performance of SOTA models in image-text retrieval tasks and (2), to not require many GPUs to train on any large-scale external dataset. It has opened the possibility of applying HADA in the industry where large-scale GPU utilisation may be considered too costly in financial or environmental terms.

Although we achieved a better result compared to the baselines, there are still rooms to improve the performance of HADA. Firstly, it can be extended not only by two pretrained models as proposed in this research but can be used with more than that number. Secondly, the use of different graph neural networks, such as the graph transformer [30], can be investigated in future work. Third, the edge feature in the graph is also considered. Currently, HADA did not implement the edge feature in our experiment, but they can be learnable parameters in graph neural networks. Last but not least, pretraining HADA on a large-scale external dataset as other SOTA have done might enhance its performance.

Acknowledgement. This publication has emanated from research supported in part by research grants from Science Foundation Ireland under grant numbers SFI/12/RC/2289, SFI/13/RC/2106, and 18/CRT/6223.

References

1. Brody, S., Alon, U., Yahav, E.: How attentive are graph attention networks? arXiv preprint [arXiv:2105.14491](https://arxiv.org/abs/2105.14491) (2021)
2. Chen, C.F.R., Fan, Q., Panda, R.: Crossvit: cross-attention multi-scale vision transformer for image classification. In: Proceedings of the IEEE/CVF International Conference on Computer VisionM pp. 357–366 (2021)

3. Chen, Y.-C., et al.: UNITER: UNiversal image-TExt representation learning. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12375, pp. 104–120. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58577-8_7
4. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: practical automated data augmentation with a reduced search space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 702–703 (2020)
5. Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional neural networks on graphs with fast localized spectral filtering. *Adv. Neural Inf. Process. Syst.* **29** (2016)
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
7. Diao, H., Zhang, Y., Ma, L., Lu, H.: Similarity reasoning and filtration for image-text matching. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, pp. 1218–1226 (2021)
8. Dosovitskiy, A., et al.: An image is worth 16×16 words: transformers for image recognition at scale. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929) (2020)
9. Dzabraev, M., Kalashnikov, M., Komkov, S., Petiushko, A.: Mdmmt: multidomain multimodal transformer for video retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3354–3363 (2021)
10. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. *Adv. Neural Inf. Process. Syst.* **33**, 6616–6628 (2020)
11. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. *Adv. Neural Inf. Process. Syst.* **30**, 1–11 (2017)
12. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9729–9738 (2020)
13. Jia, C., et al.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International Conference on Machine Learning, pp. 4904–4916. PMLR (2021)
14. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3128–3137 (2015)
15. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint [arXiv:1609.02907](https://arxiv.org/abs/1609.02907) (2016)
16. Li, D., Li, J., Le, H., Wang, G., Savarese, S., Hoi, S.C.H.: Lavis: a library for language-vision intelligence. arXiv preprint [arXiv:2209.09019](https://arxiv.org/abs/2209.09019) (2022)
17. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: bootstrapping language-image pre-training for unified vision-language understanding and generation. arXiv preprint [arXiv:2201.12086](https://arxiv.org/abs/2201.12086) (2022)
18. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: vision and language representation learning with momentum distillation. *Adv. Neural Inf. Process. Syst.* **34**, 9694–9705 (2021)
19. Li, X., et al.: OSCAR: object-semantics aligned pre-training for vision-language tasks. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12375, pp. 121–137. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58577-8_8

20. Lin, T.-Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
21. Liu, C., Mao, Z., Zhang, T., Xie, H., Wang, B., Zhang, Y.: Graph structured network for image-text matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10921–10930 (2020)
22. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with warm restarts. arXiv preprint [arXiv:1608.03983](https://arxiv.org/abs/1608.03983) (2016)
23. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint [arXiv:1711.05101](https://arxiv.org/abs/1711.05101) (2017)
24. Luo, H., et al.: Clip4clip: an empirical study of clip for end to end video clip retrieval. arXiv preprint [arXiv:2104.08860](https://arxiv.org/abs/2104.08860) (2021)
25. Miech, A., Alayrac, J.B., Laptev, I., Sivic, J., Zisserman, A.: Thinking fast and slow: efficient text-to-visual retrieval with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9826–9836 (2021)
26. Nguyen, M.D., Nguyen, B.T., Gurrin, C.: A deep local and global scene-graph matching for image-text retrieval. arXiv preprint [arXiv:2106.02400](https://arxiv.org/abs/2106.02400) (2021)
27. Nguyen, M.-D., Nguyen, B.T., Gurrin, C.: Graph-based indexing and retrieval of lifelog data. In: Lokoč, J., et al. (eds.) MMM 2021. LNCS, vol. 12573, pp. 256–267. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-67835-7_22
28. Radford, A., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning, pp. 8748–8763. PMLR (2021)
29. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: towards real-time object detection with region proposal networks. *Adv. Neural Inf. Process. Syst.* **28** (2015)
30. Shi, Y., Huang, Z., Feng, S., Zhong, H., Wang, W., Sun, Y.: Masked label prediction: unified message passing model for semi-supervised classification. arXiv preprint [arXiv:2009.03509](https://arxiv.org/abs/2009.03509) (2020)
31. Sun, S., Chen, Y.C., Li, L., Wang, S., Fang, Y., Liu, J.: Lightningdot: pre-training visual-semantic embeddings for real-time image-text retrieval. In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 982–997 (2021)
32. Tran, L.D., Nguyen, M.D., Nguyen, B., Lee, H., Zhou, L., Gurrin, C.: E-myscéal: embedding-based interactive lifelog retrieval system for lsc’22. In: Proceedings of the 5th Annual on Lifelog Search Challenge, pp. 32–37 (2022)
33. Vaswani, A., et al.: Attention is all you need. *Adv. Neural Inf. Process. Syst.* **30** (2017)
34. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint [arXiv:1710.10903](https://arxiv.org/abs/1710.10903) (2017)
35. Wang, S., Wang, R., Yao, Z., Shan, S., Chen, X.: Cross-modal scene graph matching for relationship-aware image-text retrieval. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 1508–1517 (2020)
36. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3733–3742 (2018)
37. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: From image descriptions to visual denotations: new similarity metrics for semantic inference over event descriptions. *Trans. Assoc. Comput. Linguisti.* **2**, 67–78 (2014)