



# Improving visual-semantic embeddings by learning semantically-enhanced hard negatives for cross-modal information retrieval

Yan Gong\*, Georgina Cosma

Department of Computer Science, School of Science, Loughborough University, Epinal Way, Loughborough, LE11 3TU, Leicestershire, UK

## ARTICLE INFO

### Article history:

Received 8 July 2022

Revised 3 November 2022

Accepted 17 December 2022

Available online 18 December 2022

### Keywords:

Visual semantic embedding network

Cross-modal

Information retrieval

Hard negatives

## ABSTRACT

Visual Semantic Embedding (VSE) networks aim to extract the semantics of images and their descriptions and embed them into the same latent space for cross-modal information retrieval. Most existing VSE networks are trained by adopting a hard negatives loss function which learns an objective margin between the similarity of relevant and irrelevant image–description embedding pairs. However, the objective margin in the hard negatives loss function is set as a fixed hyperparameter that ignores the semantic differences of the irrelevant image–description pairs. To address the challenge of measuring the optimal similarities between image–description pairs before obtaining the trained VSE networks, this paper presents a novel approach that comprises two main parts: (1) finds the underlying semantics of image descriptions; and (2) proposes a novel semantically-enhanced hard negatives loss function, where the learning objective is dynamically determined based on the optimal similarity scores between irrelevant image–description pairs. Extensive experiments were carried out by integrating the proposed methods into five state-of-the-art VSE networks that were applied to three benchmark datasets for cross-modal information retrieval tasks. The results revealed that the proposed methods achieved the best performance and can also be adopted by existing and future VSE networks.

© 2022 The Authors. Published by Elsevier Ltd.

This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

## 1. Introduction

In information retrieval, Visual Semantic Embedding (VSE) networks aim to create joint representations of images and textual descriptions and map these in a joint embedding space (i.e. same latent space) to enable various information retrieval-related tasks, such as image–text retrieval, image captioning, and visual question answering [1]. Within the shared embedding space, the aim is to position the relevant image–description pairs far away from the irrelevant pairs [2]. Currently, VSE literature can be summarised into: (1) approaches that extend the cross-modal encoder-decoder network for improving the learning of latent representations crossing images and descriptions [3]; (2) specifically designed attention architectures that improve the performance of networks [4]; and (3) networks that are modified based on generative adversarial methods for learning the common representation of images and

descriptions [5]. The above-mentioned studies aim to improve the VSE networks for information retrieval and have been evaluated using the benchmark MS-COCO [6] and Flickr30K [7] datasets. Few studies focus on exploring the learning potential of VSE networks. The hard negatives loss function [3] defines the learning objective of VSE networks, and it is commonly adopted by the current VSE architectures [8].

Furthermore, the hard negatives loss function learns a fixed margin that is the optimal difference between the similarity of the relevant image–description embedding pair and that of the irrelevant embedding pair. However, the fixed margin ignores the semantic differences between the irrelevant image–description pairs. The hard negatives loss function does not consider the distance of the irrelevant items to the query and sets the same learning objective (i.e. fixed margin) for both pairs, image–D1 and image–D2 (sample from Fig. 1), even though the semantic differences of the irrelevant training pairs are useful for training an information retrieval model [9]. To illustrate this point, consider Fig. 1, where in the irrelevant image–D1 pair the image and description are semantically closer than those of the irrelevant image–D2 pair, but the

\* Corresponding author.

E-mail addresses: [y.gong2@lboro.ac.uk](mailto:y.gong2@lboro.ac.uk) (Y. Gong), [g.cosma@lboro.ac.uk](mailto:g.cosma@lboro.ac.uk) (G. Cosma).

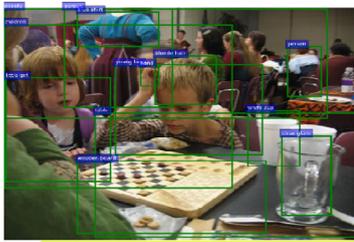


Image content: *Two kids are sitting at a table playing chess.*  
 Irrelevant D1: *Two kids are sitting at a table eating.*  
 Irrelevant D2: *Two people are climbing a rock wall.*

Fig. 1. Sample of irrelevant image–description pairs. Description D1 is the one semantically closer to the image.

hard negatives loss function sets the same learning objective for both pairs, i.e. image–D1 and image–D2, and this is not suitable. To solve the limitations of the fixed margin, Wei et al. [10] introduced a polynomial loss function with an adaptive objective margin, but their method does not consider the optimal semantic information from irrelevant image–description pairs.

Our paper aims to semantically enhance the hard negatives loss function for exploring the learning potential of VSE networks. This paper (1) proposes a new loss function for improving the learning efficiency and the cross-modal information retrieval performance of VSE networks; (2) embeds the proposed loss function within state-of-the-art VSE networks, and (3) evaluates its efficiency using benchmark datasets suitable for the task of cross-modal information retrieval. The contributions of our paper are as follows.

- A novel approach that infers the semantics of image descriptions by finding the underlying meaning of descriptions using eigendecomposition and dimensionality reduction (i.e. Singular Value Decomposition). The derived descriptions are then utilised by a proposed semantically-enhanced hard negatives loss function, entitled LSEH, when computing the optimal similarities between irrelevant image–description pairs.
- A semantically-enhanced hard negatives loss function that redefines the learning objective for VSE networks. The proposed loss function dynamically adjusts the learning objective according to the semantic similarities between irrelevant image–description pairs. Ambiguous training pairs with larger optimal similarity scores obtain larger gradients that are utilised by the proposed loss function to improve training efficiency.
- The proposed approach and loss function can be integrated into other VSE networks that improve learning efficiency and cross-modal information retrieval. Extensive experiments were carried out by integrating the proposed methods into five state-of-the-art VSE networks that were applied to the Flickr30K, MSCOCO, and IAPR TC12 datasets, and the results showed that the proposed methods achieved the best performance.

## 2. Related work

**VSE Networks.** VSE networks aim to align embeddings of relevant images and descriptions in the same latent space for cross-modal information retrieval [1]. Faghri et al. [3] proposed an Improved Visual Semantic Embedding architecture (VSE++). Image region features extracted by the faster R-CNN [11] and their descriptions were embedded into the same latent space by using a fully connected neural network and a Gated Recurrent Units (GRU) network [12]. Most state-of-the-art VSE networks improve upon VSE++. Li et al. [13] introduced a Visual Semantic Reasoning Network (VSRN) to enhance image features with image region relationships extracted by a Graph Convolution Network (GCN) [14]; Liu et al. [15] applied a Graph Structured Matching Network

(GSMN) to build a graph of image features and words and learn the fine-grained correspondence between image features and words; Diao et al. [4] proposed the Similarity Graph Reasoning and Attention Filtration network (SGRAF) that extends the attention mechanisms of image and description sets. SGRAF also provides two individual sub-networks to process the attention results between the image features and the description – where a Similarity Graph Reasoning network (SGR) builds a graph of the attention results for reasoning, and a Similarity Attention Filtration network (SAF) filters the important information from the attention results. Chen et al. [8] proposed a variation of the VSE network,  $VSE_{\infty}$ , that benefits from a generalized pooling operator which discovers the best strategy for pooling image and description embeddings. Recently, vision transformer-based networks, that are not relying on the hard negatives loss function, have become popular for cross-modal information retrieval [16]. However, compared to traditional VSE networks, vision transformer-based cross-modal retrieval networks require a large amount of data for training and the time they require for retrieving the results of a query makes them unsuitable for real-world applications [1]. The hashing-based network is another active solution for cross-modal information retrieval [17]. For example, Liu et al. [18] firstly proposed a hashing framework for learning varying hash codes of different lengths for the comparison between images and descriptions, and the learned modality-specific hash codes contain more semantics. Hashing-based networks are concerned with reducing data storage costs and improving retrieval speed. Such networks are out of scope for this paper because the focus herein is on VSE networks which mostly aim to explore the local information alignment between images and descriptions for improved retrieval performance.

**Loss Functions for Cross-modal Information Retrieval.** One of the earliest and most used cross-modal information retrieval loss functions is the Sum of Hinges Loss (LSH) [19]. LSH is also known as a negatives loss function, and it learns a fixed margin between the similarities of the relevant image–description embedding pairs and those of the irrelevant embedding pairs. A more recent hard negatives loss function, the Max of Hinges Loss (LMH) [3], is adopted in most recent VSE networks, due to its ability to outperform LSH [20,21]. An improved version of LSH, LMH only focuses on learning the hard negatives, which are the irrelevant image–description embedding pairs that are nearest to the relevant pairs. Song et al. [9] presented a margin-adaptive triplet loss for the task of cross-modal information retrieval that uses a hashing-based method which embeds the image and text into a low-dimensional Hamming space. Liu et al. [22] applied a variant triplet loss function into their novel VSE network for cross-modal information retrieval, where the input text embedding for the loss is replaced by the reconstructed image embedding of the network. Recently, Wei et al. [10] proposed a polynomial [23] based Universal Weighting Metric Loss (LUWM) with flexible objective margins, and that has been shown to outperform existing hard negatives loss functions.

A summary of the limitations of the existing loss functions are as follows. (1) The learning objectives of LSH [19] and LMH [3] are not flexible because of their fixed margins. (2) The adaptive margin in [9] is not optimal, because it relies on the computed similarities between irrelevant image–description embedding pairs by the training network which is optimising. (3) The modified ranking loss of [22] cannot be integrated into other networks. (4) LUWM [10] does not consider the optimal semantic information from irrelevant image–description pairs.

**VSE Networks with a Negatives Loss Function.** LMH was proposed by Faghri et al. [3], and thereafter other VSE networks adopted LMH. For improving the attention mechanism, Lee et al. [24] proposed an approach to align the image region features with keywords of the relevant image–description pair; and Diao et al.

[4] built an architecture that deeply extends the attention mechanisms of image-image, text-text, and image-text tasks. For extracting high-level semantics, Li et al. [13] utilised GCN [14] to explore the relations of image objects. For aggregating the image and description embeddings, Chen et al. [8] proposed a special pooling operator. Wang et al. [25] proposed an end-to-end VSE network without relying on a pre-trained CNN for image feature extraction tasks.

**Methods for Finding the Underlying Meaning of Descriptions.** BERT [26] is a supervised and widely used deep neural network for NLP tasks [29]. Singular Value Decomposition (SVD) [27] is an unsupervised matrix decomposition method and an established approach in NLP and information retrieval [28]. BERT and SVD both have dimensionality reduction capabilities that enable them to find the underlying semantic similarity between texts (e.g. sentences, image captions, documents).

### 3. Proposed semantically-enhanced hard negatives loss function

Let  $X = \{(I_i, D_i) | i = 1 \dots n\}$  denote a training set containing paired images and descriptions, where each image  $I_i$  corresponds to its relevant description  $D_i$ ;  $i$  is the index and  $n$  is the size of set  $X$ . Let  $X' = \{(v_i, u_i) | i = 1 \dots n\}$  be a set of image-description embedding pairs output by a VSE network, where each  $i$ th relevant pair consists of an image embedding  $v_i$  and its relevant description embedding  $u_i$ . Let  $\hat{v}_i = \{v_j | j = 1 \dots n, j \neq i\}$  denote a set of all image embeddings from  $X'$  irrelevant to  $u_i$ , and  $\hat{u}_i = \{u_j | j = 1 \dots n, j \neq i\}$  denote a set of all description embeddings from  $X'$  that are irrelevant to  $v_i$ . LSH, LMH, and the proposed approach and loss function that are used during the training of VSE networks are computed as follows.

#### 3.1. Related methods and notation

**LSH Description.** The basic negatives loss function, LSH, is shown in Eq. (1):

$$\text{LSH}(v_i, u_i) = \sum_{\hat{u}_i} [\alpha + s(v_i, \hat{u}_i) - s(v_i, u_i)]_+ + \sum_{\hat{v}_i} [\alpha + s(u_i, \hat{v}_i) - s(v_i, u_i)]_+ \quad (1)$$

where  $[x]_+ \equiv \max(x, 0)$ , and  $\alpha$  serves as a margin parameter. Let  $s(v_i, u_i)$  be the similarity score between the relevant image embedding  $v_i$  and description embedding  $u_i$ ; let  $s(v_i, \hat{u}_i)$  be the set of similarity scores of the image embedding  $v_i$  with its all irrelevant description embeddings  $\hat{u}_i$ ; and let  $s(u_i, \hat{v}_i)$  be the set of similarity scores of the description embedding  $u_i$  with its all irrelevant image embeddings  $\hat{v}_i$ . Given a relevant pair of image-description embeddings  $(v_i, u_i)$ , the result of the function takes the sum from irrelevant pairs  $s(v_i, \hat{u}_i)$  and  $s(u_i, \hat{v}_i)$  respectively.

**LMH Description.** The hard negatives loss function, LMH, is an improved version of LSH, that only focuses on the hard negatives [3].

$$\text{LMH}(v_i, u_i) = \max_{\hat{u}_i} [\alpha + s(v_i, \hat{u}_i) - s(v_i, u_i)]_+ + \max_{\hat{v}_i} [\alpha + s(u_i, \hat{v}_i) - s(v_i, u_i)]_+ \quad (2)$$

As shown in Eq. (2), given a relevant image-description pair  $(v_i, u_i)$ , the result of the function only takes the max value of the irrelevant pairs  $s(v_i, \hat{u}_i)$  and  $s(u_i, \hat{v}_i)$  respectively.

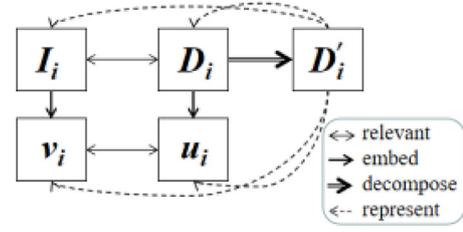


Fig. 2. Based on the joint relations between  $D_i$  with  $I_i$ ,  $v_i$ , and  $u_i$ ,  $D'_i$  can simultaneously represent the semantics of  $I_i$ ,  $v_i$ , and  $u_i$ .

#### 3.2. Proposed LSEH loss function

The proposed Semantically-Enhanced Hard negatives Loss function (LSEH) is an improved version of LMH and it is defined in Eq. (3):

$$\text{LSEH}(v_i, u_i) = \max_{\hat{u}_i} [\alpha + (s(v_i, \hat{u}_i) + f(v_i, \hat{u}_i)) - s(v_i, u_i)]_+ + \max_{\hat{v}_i} [\alpha + (s(u_i, \hat{v}_i) + f(u_i, \hat{v}_i)) - s(v_i, u_i)]_+ \quad (3)$$

LSEH introduces two sets of semantic factors  $f(v_i, \hat{u}_i)$  and  $f(u_i, \hat{v}_i)$  for the image-description embedding pairs  $(v_i, \hat{u}_i)$  and the description-image embedding pairs  $(u_i, \hat{v}_i)$  respectively, and  $f(v_i, \hat{u}_i)$  and  $f(u_i, \hat{v}_i)$  can be obtained via Eq. (4):

$$f(v_i, \hat{u}_i) = \lambda \times S(v_i, \hat{u}_i), \quad f(u_i, \hat{v}_i) = \lambda \times S(u_i, \hat{v}_i) \quad (4)$$

where  $\lambda$  serves as a temperature hyperparameter, and let  $S(v_i, \hat{u}_i)$  denote the optimal semantic similarity scores of the irrelevant image-description embedding pairs  $(v_i, \hat{u}_i)$ , and  $S(u_i, \hat{v}_i)$  denote the optimal semantic similarity scores of the irrelevant description-image embedding pairs  $(u_i, \hat{v}_i)$ . Therefore, the question is to compute the semantic factors  $f(v_i, \hat{u}_i)$  and  $f(u_i, \hat{v}_i)$ .

The semantic factors are computed by finding the underlying meaning of descriptions using SVD. After pre-processing [30], the descriptions set  $\{D_i | i = 1 \dots n\}$  is converted to a matrix  $\mathbf{A}$  of size  $n \times w$ , where  $n$  is the number of descriptions,  $w$  is the total number of unique terms found in the set of descriptions, and each  $i$ th row of  $\mathbf{A}$  corresponds to each  $i$ th description  $D_i$ . Then the truncated SVD is applied as shown in Eq. (5) [31]:

$$\mathbf{A}_{n \times w} \approx \mathbf{U}_{n \times k} \mathbf{\Lambda}_{k \times k} \mathbf{V}_{k \times w}^T, \quad \mathbf{B}_{n \times k} = \mathbf{A}_{n \times w} \mathbf{V}_{w \times k} \quad (5)$$

where  $k$  is the number of singular values. The reduced matrix  $\mathbf{B}$  containing  $n$  rows of description vectors ( $k$  dimension) is obtained by multiplying the original descriptions matrix  $\mathbf{A}_{n \times w}$  with matrix  $\mathbf{V}_{w \times k}$ .

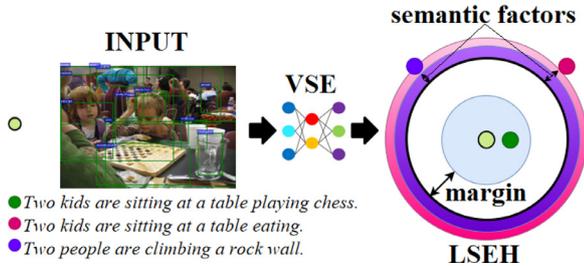
Let a set  $C = \{D'_i | i = 1 \dots n\}$  contain the reduced description vectors, and be derived from matrix  $\mathbf{B}_{n \times k}$ , where each  $i$ th element  $D'_i$  is each  $i$ th row vector of  $\mathbf{B}$ , therefore  $D'_i$  represents the extracted semantic of each  $i$ th description  $D_i$ . Also, as shown in Fig. 2, description  $D_i$  is relevant to image  $I_i$ , and embeddings  $u_i$  and  $v_i$  are output from  $D_i$  and  $I_i$  respectively, hence  $D'_i$  can also simultaneously represent the semantics of  $I_i$ ,  $v_i$ , and  $u_i$ .

Therefore, let  $\hat{D}'_i = \{D'_j | j = 1 \dots n, j \neq i\}$  denote a set of reduced vectors of descriptions from set  $C$ , where each  $j$ th vector  $D'_j$  simultaneously represents the optimal semantics of  $v_j$  and  $u_j$  from sets  $\hat{v}_i$  and  $\hat{u}_i$  respectively, then  $S(v_i, \hat{u}_i)$  and  $S(u_i, \hat{v}_i)$  can be alternatively calculated using Eq. (6):

$$S(v_i, \hat{u}_i) = S(u_i, \hat{v}_i) = s(D'_i, \hat{D}'_i) \quad (6)$$

where  $s(D'_i, \hat{D}'_i) \in [-1, 1]$  computes a set of cosine similarity scores of  $D'_i$  with  $\hat{D}'_i$ , thus:

$$f(v_i, \hat{u}_i) = f(u_i, \hat{v}_i) = \lambda \times s(D'_i, \hat{D}'_i) \quad (7)$$



**Fig. 3.** Illustrations of LSEH. LSEH dynamically adjusts the learning objective for the VSE network, and takes the maximum gradients from the ambiguous image-description pairs.

Furthermore, as a recent popular pre-trained language model, BERT [26] can also find the underlying meaning of descriptions. BERT relies on a large amount of data for training and is known to perform well in processing long documents [26]. The experiment in Section 4.7 compares the performance of LSEH when using SVD and when using BERT. Finally, integrating Eq. (7) into Eq. (3), LSEH is computed as Eq. (8):

$$\begin{aligned} \text{LSEH}(v_i, u_i) = & \max_{\hat{u}_i} [\alpha + s(v_i, \hat{u}_i) + \lambda \times s(D'_i, \hat{D}'_i) - s(v_i, u_i)]_+ \\ & + \max_{\hat{v}_i} [\alpha + s(u_i, \hat{v}_i) + \lambda \times s(D'_i, \hat{D}'_i) - s(v_i, u_i)]_+ \end{aligned} \quad (8)$$

**Illustration of LSEH.** As shown in Fig. 3, the learning objective of LSEH, defined by margin  $\alpha$ , is dynamically adjusted for every irrelevant image-description pair based on their optimal semantic similarity (see Eq. (7)). LSEH has two purposes: (1) it dynamically adjusts the learning objective for the VSE network for flexible learning; and (2) performs efficient training by taking the maximum gradients from the ambiguous image-description pairs because their large optimal similarity scores that cause the large gradients.

### 3.3. Process of training VSE networks using LSEH

Algorithm (1) shows the process of using LSEH to train a VSE

---

#### Algorithm 1 Pseudocode of training a VSE network using LSEH.

---

**Input:**  $X$  training set (image-description pairs),  
 $E$  training epochs,  $Vstep$  validation step  
**Output:** LSEH

- 1: Obtain set  $C$  based on set  $X$   $\triangleright$  Formula 5
- 2: Apply the VSE network
- 3: **for** epoch in  $E$  **do**
- 4:   Split sets  $(X, C)$  to mini-batches
- 5:   **for** mini-batch  $(X_p, C_p)$  in  $(X, C)$  **do**
- 6:     Embed set  $X_p$  to set  $X'_p$
- 7:     **for**  $v_i, u_i$  in  $X'_p$  **do**
- 8:       Use  $C_p$  to compute the semantic factors  $\triangleright$  Formula 7
- 9:       Use the semantic factors to compute LSEH  $\triangleright$  Formula 8
- 10:     **end for**
- 11:     Backpropagate LSEH for updating network
- 12:     **if** mini-batches number ==  $Vstep$  **then**
- 13:       Validate the model
- 14:       **if** the overall performance is the best **then**
- 15:         Save the current model
- 16:       **end if**
- 17:     **end if**
- 18:   **end for**
- 19: **end for**

---

network. Initially, the descriptions from the train set  $X$  are represented as a set of reduced vectors  $C$  (line 1). Thereafter the VSE

**Table 1**

Dataset split of Flickr30K, MS-COCO, and IAPR TC-12.

Dataset	Train	Validate	Test
Flickr30K	29,000	1014	1000
MS-COCO	113,287	1000	1000 and 5000
IAPRTC-12	18,000	1000	1000

network is applied as follows (line 2). In every epoch, the train set  $X$  is split into mini-batches and then the corresponding set of vectors are obtained from  $C$  such that  $\{(X, C) = (X_p, C_p) | p = 1 \dots m\}$ , where  $p$  is the index and  $m$  is the number of mini-batches (lines 3–4).

For each mini-batch  $(X_p, C_p)$ , the network outputs a set  $X'_p$  of embeddings from  $X_p$  (lines 5–6). Thereafter, for every relevant image-description embedding pair  $(v_i, u_i)$  in set  $X'_p$ , the semantic factors are computed using set  $C_p$  (Eq. (7)), then the LSEH value is computed (Eq. (8)) (lines 7–10). Finally, the LSEH value from the mini-batch is used for backpropagation (line 11). Furthermore, if the number of mini-batches reaches the validation step  $Vstep$ , the network is validated and the model that achieved the best overall performance is saved (lines 12–17).

## 4. Experiments

### 4.1. Experiment setup

**Datasets and Protocols.** The datasets utilised in the experiments are Flickr30K [7] and MS-COCO [6], and these datasets are typically used for evaluating the performance of VSE networks [3]. This paper also utilises the IAPR TC-12 dataset [32]. The datasets were split into train, test and validation sets as shown in Table 1 [3]. In the Flickr30K and MS-COCO datasets, every image is associated with five relevant textual descriptions. IAPR TC-12 includes pictures of different sports and actions, photographs of people, animals, cities, landscapes and many other aspects of contemporary life. Each image in IAPR TC-12 is associated with one relevant textual description [32].

The descriptions were pre-processed with lowercase, stemming, removing punctuation and alphabetic, filtering out stop words and short words, and the term frequency-inverse document frequency (TFIDF) text vectorizer was applied to transform the descriptions into usable vectors [33]. The number of singular values  $k$  for SVD was set as 400 for all datasets. The process of representing descriptions in a reduced dimensional space is described in Section 3. SVD is computed once for each dataset. LSEH takes the vector from matrix  $\mathbf{B}_{n \times k}$  (see Eq. (5)) to compute the cosine similarity between vectors from the matrix. Therefore, the cost for LSEH is only the cosine similarity computation. Section 4.6 compares LSEH, LMH, and LUWM on computation time.

**Network Implementations.** The experiments tested VSE++ [3], VSRN [13], SGRAF (sub-networks SGRAF-SAF and SGRAF-SGR) [4], VSE  $\infty$  [8], and GSMN [15]. For consistency of comparisons, VSE++, VSRN, SGRAF, VSE  $\infty$ , and GSMN have been tuned using the image region feature (size of  $36 \times 2048$ ) that was pre-extracted by the modified faster R-CNN [11], where the architecture of the faster R-CNN includes a backbone of ResNet-101 that was pre-trained on ImageNet [34] and Visual Genome [35]. The architectures of VSE++, VSRN, and SGRAF (sub-networks SGRAF-SAF and SGRAF-SGR) follow those described in [3,4,13] respectively. For VSE  $\infty$ , the image region mode of architecture was used, and it follows the one described in [8]. For GSMN the sparse mode of architecture was used, and it follows the one described in [15]. The source-code files of the abovementioned networks when using the proposed LSEH loss

**Table 2**

Hyperparameters with different settings for LMH1, LMH2, LSEH, and LUWM, in terms of Margin, Learning Rate (LR), and epochs to update learning rate (Update).

Loss	Flickr30K			MS-COCO			IAPR TC-12		
	Margin	LR	Update	Margin	LR	Update	Margin	LR	Update
	VSE+								
LMH1	0.20	0.0002	15	0.20	0.0002	15	0.20	0.0002	25
LMH2	0.16	0.0020	3	0.16	0.0030	4	0.16	0.0008	5
LSEH	0.16-0.21	0.0020	3	0.16-0.21	0.0030	4	0.16-0.21	0.0008	5
	VSRN								
LMH1	0.20	0.0002	10	0.20	0.0002	15	0.20	0.0005	20
LMH2	0.16	0.0004	5	0.16	0.0004	6	0.16	0.0005	20
LSEH	0.16-0.21	0.0004	5	0.16-0.21	0.0004	6	0.16-0.21	0.0005	20
	SGRAF-SAF								
LMH1	0.20	0.0002	20	0.20	0.0002	10	0.20	0.0002	20
LMH2	0.16	0.0004	16	0.16	0.0004	6	0.16	0.0004	15
LSEH	0.16-0.21	0.0004	16	0.16-0.21	0.0004	6	0.16-0.21	0.0004	15
	SGRAF-SGR								
LMH1	0.20	0.0002	30	0.20	0.0002	10	0.20	0.0002	30
LMH2	0.16	0.0004	20	0.16	0.0004	6	0.16	0.0004	20
LSEH	0.16-0.21	0.0004	20	0.16-0.21	0.0004	6	0.16-0.21	0.0004	20
	VSE $\infty$								
LMH1	0.20	0.0005	15	0.20	0.0005	15	0.20	0.0005	15
LMH2	0.16	0.0008	10	0.16	0.0010	10	0.16	0.0009	10
LSEH	0.16-0.21	0.0008	10	0.16-0.21	0.0010	10	0.16-0.21	0.0009	10
	GSMN								
LMH1	0.20	0.0002	15	0.20	0.0005	5	-	-	-
LUWM	-	0.0002	15	-	0.0005	5	-	-	-
LSEH	0.16-0.21	0.0004	10	0.16-0.21	0.0006	4	-	-	-

function and their loss functions are provided in our GitHub repository<sup>1</sup>.

**Network Training Hyperparameters.** The benchmark hyperparameter settings for each network are described in [3,4,8,10,13] respectively, where the benchmark hyperparameter settings on the new IAPRTC-12 refer to Flickr30K because they contain a similar number of images. Let LMH1 denote LMH with the benchmark hyperparameters of each network and LMH2 denote LMH with LSEH's hyperparameter settings, then all the hyperparameters except those shown in Table 2 for using LMH1, LMH2, LUWM, and LSEH refer to each network's benchmark-settings.

The hyperparameters shown in Table 2 were selected experimentally and tuned as follows: (1) for LSEH the learning and update learning rates can be set to a larger value at an earlier epoch than for LMH1; and (2) LSEH sets the margin  $\alpha$  as 0.185 and the semantic factor temperature hyperparameter  $\lambda$  as 0.025. All of the experiments were conducted on a workstation with NVIDIA Titan GPU. For a fair comparison, all software programs were set with the same random seed.

**Evaluation Measures.** To evaluate the performance of the networks two evaluation measures were used: Recall@k and M-Recall.

**Recall@k.** The evaluation measure for the cross-modal information retrieval experiments is the commonly used Recall at rank k (Recall@k), which is defined as the percentage of relevant items in the top k retrieved results [3]. The experiments evaluate the performance of the network in retrieving any one of the relevant items from the list of relevant items and computed the average Recall of the results of the test queries [3].

**M-Recall.** Defined as Eq. (9) [3]:

$$M - \text{Recall} = \frac{1}{6} \sum_k^{1,5,10} (\text{Recall}_{I2T}@k + \text{Recall}_{T2I}@k) \quad (9)$$

where M-Recall is the Mean of average Recall@1, 5, and 10 from both image-to-text (I2T) and text-to-image (T2I) retrieval. M-Recall is used for evaluating the overall performance of a network during the validation stages.

#### 4.2. Comparison of LMH and LSEH on training efficiency

This section compares the learning performance of VSE++, VSRN, SGRAF-SAF, SGRAF-SGR, and VSE $\infty$  when using LMH1, LMH2, and LSEH using the train and validation sets of the Flickr30K, MS-COCO, and IAPR TC-12 datasets. Obtaining the optimal trained model for the VSE networks is through validating the middle-trained model on the overall performance using the M-Recall evaluation measure. Validation occurs at every 1000 mini-batches for SGRAF [4], and at every 500 mini-batches for the VSE $\infty$  [8], VSRN [13], and VSE++ [3].

**Graphic and Quantitative Results.** Fig. 4 compares the training performance of networks using LMH1, LMH2, and LSEH. The comparison considers the number of epochs that each network needs to reach its largest M-Recall value. The comparison is illustrated in Fig. 4.

Table 3 quantifies the improvement in training efficiency of each network when using LMH and the proposed LSEH loss functions. In Table 3, for the comparison of LMH1 and LSEH  $M\text{-Recall}_{LMH1}$  is LMH1's largest M-Recall value.  $\text{epochs}_{LMH1}$  is the number of epochs needed to achieve  $M\text{-Recall}_{LMH1}$  by LMH1.  $\text{epochs}_{LSEH}$  is the number of epochs needed to achieve  $M\text{-Recall}_{LMH1}$  by LSEH. Same interpretation applies for the comparison of LMH2 and LSEH (right side of Table 3).

Difference is computed using Eq. (10).

$$\text{Difference} = \left( \frac{\text{epochs}_{LSEH} - \text{epochs}_{LMH}}{\text{epochs}_{LMH}} \right) \%, \quad (10)$$

where LMH is either LMH1 or LMH2. A negative Difference value denotes an improvement in performance when using LSEH. As shown in Table 3, when using LSEH, the epochs needed to exceed the largest M-Recall values with LMH1 and LMH2 could be reduced by 53.2% and 74.7% on Flickr30K, by 43.0% and 52.1% on MS-COCO, and by 48.5% and 69.8% on IAPR TC-12 respectively with LSEH. The largest improvement was for VSE++, where the training efficiency on IAPR TC-12 was improved by approximately 80.5% with LSEH.

<sup>1</sup> <https://github.com/yangong23/VSEnetworksLSEH>

**Table 3**  
Comparison of training efficiency of each network when using LMH and the proposed LSEH loss function. Difference is calculated as shown in Eq. (10).

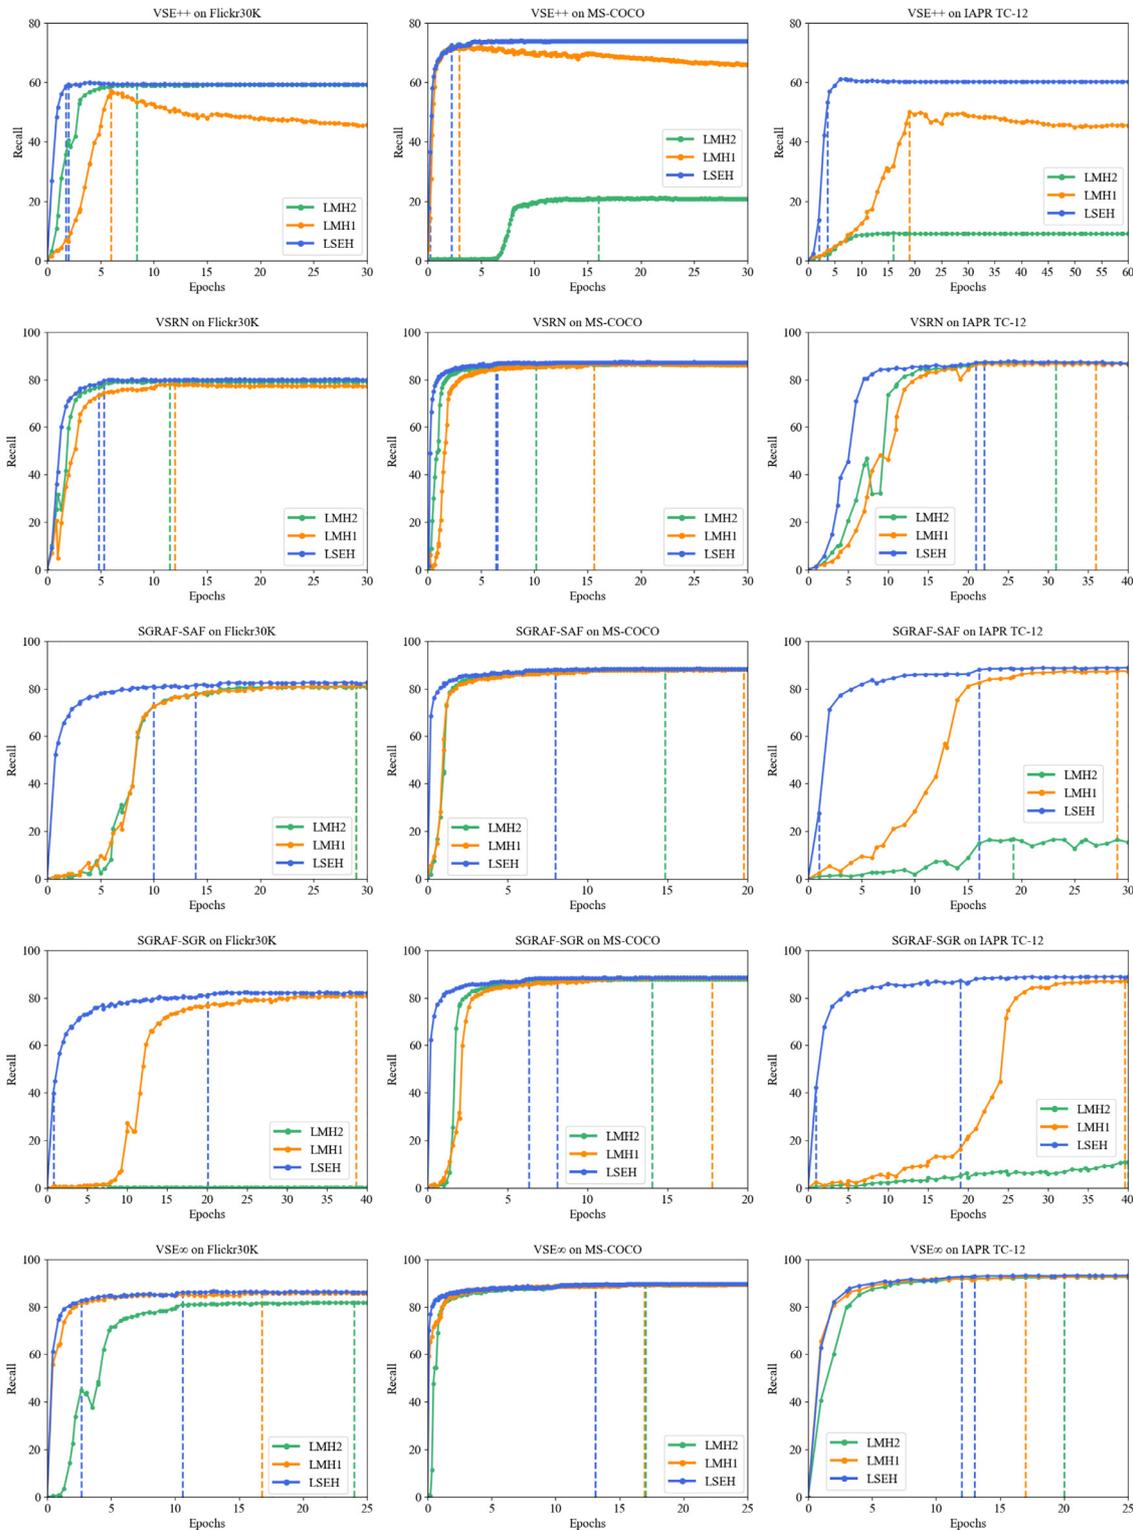
Network	Comparison of LMH1 and LSEH				Comparison of LMH2 and LSEH			
	M-Recall <sub>LMH1</sub>	epochs <sub>LMH1</sub>	epochs <sub>LSEH</sub>	Difference	M-Recall <sub>LMH2</sub>	epochs <sub>LMH2</sub>	epochs <sub>LSEH</sub>	Difference
Flickr30K								
VSE+	57.1	6.0	1.8	-4.2 (-70.0%)	59.1	8.4	2.0	-6.4 (-76.2%)
VSRN	78.0	12.0	4.9	-7.1 (-59.2%)	79.4	11.5	5.3	-6.2 (-53.9%)
SGRAF-SAF	81.4	29.0	13.9	-15.1 (-52.1%)	80.9	29.0	10.0	-19.0 (-65.5%)
SGRAF-SGR	81.0	38.6	20.1	-18.5 (-47.9%)	1.6	7.0	0.8	-6.2 (-88.6%)
VSE <sub>∞</sub>	85.9	16.8	10.6	-6.2 (-36.9%)	81.8	24.0	2.6	-21.4 (-89.2%)
Average				-10.2 (-53.2%)				-11.8 (-74.7%)
MS-COCO								
VSE+	72.5	2.9	2.3	-0.6 (-20.7%)	21.3	16.0	0.2	-15.8 (-98.8%)
VSRN	86.8	15.6	6.6	-9.0 (-57.7%)	86.7	10.2	6.4	-3.8 (-37.3%)
SGRAF-SAF	88.0	19.8	8.0	-11.8 (-59.6%)	88.0	14.8	8.0	-6.8 (-45.9%)
SGRAF-SGR	88.4	17.8	8.1	-9.7 (-54.5%)	87.8	14.0	6.3	-7.7 (-55.0%)
VSE <sub>∞</sub>	89.4	16.9	13.1	-3.8 (-22.5%)	89.5	17.1	13.1	-4.0 (-23.4%)
Average				-7.0 (-43.0%)				-7.6 (-52.1%)
IAPR TC-12								
VSE+	50.0	19.0	3.7	-15.3 (-80.5%)	9.3	16.0	2.0	-14.0 (-87.5%)
VSRN	86.7	36.0	21.0	-15.0 (-41.7%)	87.2	31.0	22.0	-9.0 (-29.0%)
SGRAF-SAF	87.3	29.0	16.0	-13.0 (-44.8%)	16.8	19.2	1.0	-18.2 (-94.8%)
SGRAF-SGR	87.1	39.6	19.0	-20.6 (-52.0%)	10.9	40.0	1.0	-39.0 (-97.5%)
VSE <sub>∞</sub>	92.7	17.0	13.0	-4.0 (-23.5%)	92.5	20.0	12.0	-8.0 (-40.0%)
Average				-13.6 (-48.5%)				-17.6 (-69.8%)

**Table 4**  
Results of cross-modal information retrieval by LMH1, LMH2, and LSEH on the Flickr30K dataset in terms of average Recall@k (%).

Network	Loss	Image-to-Text Retrieval				Text-to-Image Retrieval			
		R@1	R@5	R@10	Mean	R@1	R@5	R@10	Mean
VSE+	LMH1	38.9	68.0	78.5	61.8 (±16.8)	29.7	58.7	69.8	52.7 (±16.9)
	LMH2	41.3	71.0	82.2	64.8 (±17.3)	<b>33.3</b>	61.5	72.0	55.6 (±16.3)
	LSEH	<b>45.9</b>	<b>74.0</b>	<b>82.7</b>	<b>67.5 (±15.7)</b>	33.2	<b>62.2</b>	<b>73.3</b>	<b>56.2 (±16.9)</b>
VSRN	LMH1	69.8	89.0	94.4	84.4 (±10.6)	52.1	78.8	86.6	72.5 (±14.8)
	LMH2	71.1	91.5	<b>95.8</b>	86.1 (±10.8)	54.8	81.0	87.4	74.4 (±14.1)
	LSEH	<b>73.0</b>	<b>92.8</b>	95.7	<b>87.2 (±10.1)</b>	<b>55.8</b>	<b>81.9</b>	<b>88.8</b>	<b>75.5 (±14.2)</b>
SGRAF-SAF	LMH1	75.5	93.6	97.0	88.7 (±9.4)	55.3	81.7	88.5	75.2 (±14.3)
	LMH2	74.0	93.6	97.0	88.2 (±10.1)	54.6	80.6	87.0	74.1 (±14.0)
	LSEH	<b>76.2</b>	<b>93.8</b>	<b>97.2</b>	<b>89.1 (±9.2)</b>	<b>57.7</b>	<b>82.3</b>	<b>88.8</b>	<b>76.3 (±13.4)</b>
SGRAF-SGR	LMH1	74.2	92.5	96.5	87.7 (±9.7)	55.4	80.6	85.9	74.0 (±13.3)
	LMH2	0.3	1.3	2.0	1.2 (±0.7)	0.4	1.5	2.8	1.6 (±1.0)
	LSEH	<b>78.2</b>	<b>94.3</b>	<b>96.8</b>	<b>89.8 (±8.2)</b>	<b>57.6</b>	<b>82.5</b>	<b>87.9</b>	<b>76.0 (±13.2)</b>
VSE <sub>∞</sub>	LMH1	80.8	<b>96.4</b>	98.3	91.8 (±7.8)	62.6	86.9	91.7	80.4 (±12.7)
	LMH2	74.4	91.7	96.3	87.5 (±9.4)	56.1	82.3	89.0	75.8 (±14.2)
	LSEH	<b>82.4</b>	96.0	<b>98.6</b>	<b>92.3 (±7.1)</b>	<b>63.7</b>	<b>87.1</b>	<b>92.5</b>	<b>81.1 (±12.5)</b>

**Table 5**  
Results of cross-modal information retrieval by LMH1, LMH2, and LSEH on the MS-COCO1K dataset in terms of average Recall@k (%).

Network	Loss	Image-to-Text Retrieval				Text-to-Image Retrieval			
		R@1	R@5	R@10	Mean	R@1	R@5	R@10	Mean
VSE+	LMH1	54.5	84.9	91.7	77.0 (±16.2)	40.8	75.6	86.3	67.6 (±19.4)
	LMH2	11.9	38.5	53.0	34.5 (±17.0)	2.5	8.6	13.8	8.3 (±4.6)
	LSEH	<b>56.3</b>	<b>85.3</b>	<b>92.2</b>	<b>77.9 (±15.6)</b>	<b>41.2</b>	<b>76.1</b>	<b>86.6</b>	<b>68.0 (±19.4)</b>
VSRN	LMH1	76.4	94.2	97.6	89.4 (±9.3)	<b>63.1</b>	89.4	94.3	82.3 (±13.7)
	LMH2	73.5	94.7	98.1	88.8 (±10.9)	61.7	89.3	94.8	81.9 (±14.5)
	LSEH	<b>76.5</b>	<b>95.7</b>	<b>98.4</b>	<b>90.2 (±9.7)</b>	62.0	<b>90.1</b>	<b>95.1</b>	<b>82.4 (±14.6)</b>
SGRAF-SAF	LMH1	<b>80.2</b>	96.6	98.7	91.8 (±8.3)	64.5	90.5	96.0	83.7 (±13.7)
	LMH2	78.8	96.3	98.6	91.2 (±8.8)	64.0	90.6	95.5	83.4 (±13.8)
	LSEH	80.1	<b>97.3</b>	<b>98.7</b>	<b>92.0 (±8.5)</b>	<b>64.6</b>	<b>90.9</b>	<b>96.0</b>	<b>83.8 (±13.8)</b>
SGRAF-SGR	LMH1	79.7	97.0	98.8	91.8 (±8.6)	<b>64.0</b>	90.5	95.7	83.4 (±13.9)
	LMH2	79.7	96.3	98.6	91.5 (±8.4)	63.7	90.5	95.5	83.2 (±14.0)
	LSEH	<b>80.1</b>	<b>97.7</b>	<b>99.1</b>	<b>92.3 (±8.6)</b>	63.9	<b>90.7</b>	<b>95.9</b>	<b>83.5 (±14.0)</b>
VSE <sub>∞</sub>	LMH1	80.8	<b>97.0</b>	<b>99.0</b>	92.3 (±8.1)	66.0	91.7	96.1	84.6 (±13.3)
	LMH2	81.0	96.8	99.0	92.3 (±8.0)	66.4	<b>92.0</b>	96.0	84.8 (±13.1)
	LSEH	<b>82.2</b>	96.9	98.6	<b>92.6 (±7.4)</b>	<b>66.5</b>	91.9	<b>96.2</b>	<b>84.9 (±13.1)</b>



**Fig. 4.** Comparison of training efficiency when using LMH1, LMH2, and LSEH to train various VSE networks on Flickr30K, MS-COCO, and IAPR TC-12. Recall (in Y-axis) is M-Recall (%). The largest M-Recall values of LMH1 and LMH2 are denoted by the vertical dashed lines of orange and green respectively. The vertical dashed blue lines indicate that LSEH’s M-Recall has exceeded the best performance of LMH1 or LHM2. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

### 4.3. Comparison of LMH and LSEH on cross-modal information retrieval

This section compares the cross-modal information retrieval performance between the optimal models of VSE++, VSRN, SGRAF-

SAF, SGRAF-SGR, and VSE $\infty$  with using LMH1, LMH2, and LSEH on the Flickr30K, MS-COCO1K, MS-COCO5K, and IAPR TC-12 test sets. Tables 4, 5, 6, and 7 show the results of the comparisons of LMH1, LMH2, and LSEH when integrated into various networks and applied to the datasets of Flickr30K (see Table 4), MS-COCO1K (see

**Table 6**

Results of cross-modal information retrieval by LMH1, LMH2, and LSEH on the MS-COCO5K dataset in terms of average Recall@k (%).

Network	Loss	Image-to-Text Retrieval				Text-to-Image Retrieval			
		R@1	R@5	R@10	Mean	R@1	R@5	R@10	Mean
VSE+	LMH1	27.6	56.6	69.5	51.2 ( $\pm 17.5$ )	<b>20.0</b>	45.7	59.0	41.6 ( $\pm 16.2$ )
	LMH2	3.0	12.8	20.9	12.2 ( $\pm 7.3$ )	0.4	2.2	3.7	2.1 ( $\pm 1.3$ )
	LSEH	<b>28.5</b>	<b>57.6</b>	<b>70.3</b>	<b>52.1 (<math>\pm 17.5</math>)</b>	19.3	<b>46.2</b>	<b>60.0</b>	<b>41.8 (<math>\pm 16.9</math>)</b>
VSRN	LMH1	49.4	79.3	88.5	72.4 ( $\pm 16.7$ )	38.4	69.3	79.8	62.5 ( $\pm 17.6$ )
	LMH2	49.3	79.4	88.4	72.4 ( $\pm 16.7$ )	37.5	68.3	79.7	61.8 ( $\pm 17.8$ )
	LSEH	<b>50.3</b>	<b>80.3</b>	<b>88.6</b>	<b>73.1 (<math>\pm 16.5</math>)</b>	<b>38.6</b>	<b>69.6</b>	<b>80.6</b>	<b>62.9 (<math>\pm 17.8</math>)</b>
SGRAF-SAF	LMH1	54.4	82.9	<b>91.0</b>	76.1 ( $\pm 15.7$ )	40.1	69.7	80.3	63.4 ( $\pm 17.0$ )
	LMH2	54.8	82.8	90.2	75.9 ( $\pm 15.2$ )	39.9	69.3	80.0	63.1 ( $\pm 17.0$ )
	LSEH	<b>56.4</b>	<b>83.5</b>	90.7	<b>76.9 (<math>\pm 14.8</math>)</b>	<b>40.6</b>	<b>69.7</b>	<b>80.7</b>	<b>63.7 (<math>\pm 16.9</math>)</b>
SGRAF-SGR	LMH1	56.8	83.1	<b>91.1</b>	77.0 ( $\pm 14.7$ )	<b>40.8</b>	69.7	80.6	63.7 ( $\pm 16.8$ )
	LMH2	55.9	83.4	90.8	76.7 ( $\pm 15.0$ )	39.6	68.9	79.7	62.7 ( $\pm 16.9$ )
	LSEH	<b>57.8</b>	<b>83.7</b>	91.0	<b>77.5 (<math>\pm 14.2</math>)</b>	40.6	<b>69.8</b>	<b>80.6</b>	<b>63.7 (<math>\pm 16.9</math>)</b>
VSE $\infty$	LMH1	58.3	84.7	91.8	78.3 ( $\pm 14.4$ )	42.5	72.7	83.0	66.1 ( $\pm 17.2$ )
	LMH2	<b>58.9</b>	84.9	92.2	78.7 ( $\pm 14.3$ )	42.2	72.4	82.8	65.8 ( $\pm 17.2$ )
	LSEH	58.8	<b>85.4</b>	<b>92.6</b>	<b>78.9 (<math>\pm 14.5</math>)</b>	<b>43.1</b>	<b>73.2</b>	<b>83.2</b>	<b>66.5 (<math>\pm 17.0</math>)</b>

**Table 7**

Results of cross-modal information retrieval by LMH1, LMH2, and LSEH on the IAPR TC-12 dataset in terms of average Recall@k (%).

Network	Loss	Image-to-Text Retrieval				Text-to-Image Retrieval			
		R@1	R@5	R@10	Mean	R@1	R@5	R@10	Mean
VSE+	LMH1	25.0	56.5	70.7	50.7 ( $\pm 19.1$ )	27.4	60.9	73.5	53.9 ( $\pm 19.5$ )
	LMH2	5.0	14.2	23.1	14.1 ( $\pm 7.4$ )	0.7	3.5	6.1	3.4 ( $\pm 2.2$ )
	LSEH	<b>35.7</b>	<b>69.6</b>	<b>82.5</b>	<b>62.6 (<math>\pm 19.7</math>)</b>	<b>36.2</b>	<b>71.3</b>	<b>83.2</b>	<b>63.6 (<math>\pm 20.0</math>)</b>
VSRN	LMH1	71.5	92.9	96.4	86.9 ( $\pm 11.0$ )	71.8	92.8	96.0	86.9 ( $\pm 10.7$ )
	LMH2	71.4	93.4	<b>96.5</b>	87.1 ( $\pm 11.2$ )	70.8	<b>93.3</b>	<b>96.5</b>	86.9 ( $\pm 11.4$ )
	LSEH	<b>73.9</b>	<b>93.8</b>	96.2	<b>88.0 (<math>\pm 10.0</math>)</b>	<b>73.3</b>	93.2	96.0	<b>87.5 (<math>\pm 10.1</math>)</b>
SGRAF-SAF	LMH1	70.7	94.4	97.4	87.5 ( $\pm 11.9$ )	73.1	93.4	97.2	87.9 ( $\pm 10.6$ )
	LMH2	9.8	25.5	37.1	24.1 ( $\pm 11.2$ )	1.9	8.7	14.0	8.2 ( $\pm 5.0$ )
	LSEH	<b>74.5</b>	<b>95.1</b>	<b>97.8</b>	<b>89.1 (<math>\pm 10.4</math>)</b>	<b>73.7</b>	<b>94.5</b>	<b>97.7</b>	<b>88.6 (<math>\pm 10.6</math>)</b>
SGRAF-SGR	LMH1	70.9	93.3	97.8	87.3 ( $\pm 11.8$ )	72.1	92.8	96.7	87.2 ( $\pm 10.8$ )
	LMH2	4.3	12.7	19.0	12.0 ( $\pm 6.0$ )	2.3	9.7	15.2	9.1 ( $\pm 5.3$ )
	LSEH	<b>75.1</b>	<b>95.6</b>	<b>98.1</b>	<b>89.6 (<math>\pm 10.3</math>)</b>	<b>75.1</b>	<b>94.9</b>	<b>98.0</b>	<b>89.3 (<math>\pm 10.1</math>)</b>
VSE $\infty$	LMH1	81.5	97.3	<b>99.0</b>	92.6 ( $\pm 7.9$ )	79.1	96.4	<b>98.8</b>	91.4 ( $\pm 8.8$ )
	LMH2	81.0	96.7	98.8	92.2 ( $\pm 7.9$ )	80.7	95.8	97.9	91.5 ( $\pm 7.7$ )
	LSEH	<b>83.7</b>	<b>97.7</b>	98.8	<b>93.4 (<math>\pm 6.9</math>)</b>	<b>81.4</b>	<b>96.9</b>	98.5	<b>92.3 (<math>\pm 7.7</math>)</b>

**Table 8**

Mean average Recall of each network for each dataset. A summary derived from Tables 4, 5, 6, and 7.

Method	Flickr30k		MS-COCO1K	
	Image-Text	Text-Image	Image-Text	Text-Image
LMH1	82.9 ( $\pm 15.6$ )	71.0 ( $\pm 17.3$ )	88.8 ( $\pm 10.9$ )	80.3 ( $\pm 16.3$ )
LMH2	65.6 ( $\pm 35.1$ )	56.3 ( $\pm 31.3$ )	79.7 ( $\pm 25.2$ )	68.3 ( $\pm 32.6$ )
LSEH	<b>85.2 (<math>\pm 13.8</math>)</b>	<b>73.0 (<math>\pm 16.6</math>)</b>	<b>89.0 (<math>\pm 11.8</math>)</b>	<b>80.5 (<math>\pm 16.4</math>)</b>
Method	MS-COCO5K		IAPR TC-12	
	Image-Text	Text-Image	Image-Text	Text-Image
LMH1	71.0 ( $\pm 18.8$ )	59.4 ( $\pm 19.2$ )	81.0 ( $\pm 20.0$ )	81.5 ( $\pm 18.8$ )
LMH2	63.2 ( $\pm 29.2$ )	51.1 ( $\pm 29.0$ )	45.9 ( $\pm 37.1$ )	39.8 ( $\pm 41.0$ )
LSEH	<b>71.7 (<math>\pm 18.5</math>)</b>	<b>59.7 (<math>\pm 19.3</math>)</b>	<b>84.5 (<math>\pm 16.5</math>)</b>	<b>84.3 (<math>\pm 16.3</math>)</b>

Table 5), MS-COCO5K (see Table 6), and IAPR TC-12 (see Table 7) respectively. Recall values @1, @5, and @10 show the average Recall values across all queries found in each of the test sets (see Table 1). The last column of Tables 4, 5, 6, and 7 shows the mean of the average Recall values (i.e. the average of columns 3–5). Table 8 summarises the mean average Recall across five networks for each dataset from Tables 4, 5, 6, and 7, and below is a summary of the main findings for each dataset.

(1) **Flickr30K**. The proposed LSEH reached a Recall of 85.2% for image-to-text, and a Recall of 73.0% for text-to-image retrieval. LSEH outperformed LMH1 by 2.3% and 2.0% for image-to-text and text-to-image retrieval, respectively. LSEH outperformed LMH2 by

19.6% for image-to-text retrieval and by 16.7% for text-to-image retrieval.

(2) **MS-COCO1K**. For LSEH, Recall reached 89.0% for image-to-text retrieval and 80.5% for text-to-image retrieval, and outperformed LMH1 by 0.2% and 0.2% for those tasks, respectively. LSEH also outperformed LMH2 for image-to-text and text-to-image retrieval by 9.3% and 12.2%, respectively.

(3) **MS-COCO5K**. For image-to-text retrieval, LSEH's Recall reached 71.7% which outperformed LMH1 by 0.7% and LMH2 by 8.5%. For text-to-image retrieval, the Recall of the proposed LSEH reached 59.7% which outperformed LMH1 by 0.3% and LMH2 by 8.6%.

(4) **IAPR TC-12**. LSEH's Recall values for image-to-text and text-to-image retrieval were 84.5% and 84.3% respectively, and these outperformed the results of LMH1 by 3.5% and 2.8% respectively. The proposed LSEH outperformed LMH2 for image-to-text and text-to-image retrieval by 38.6% and 44.5% respectively.

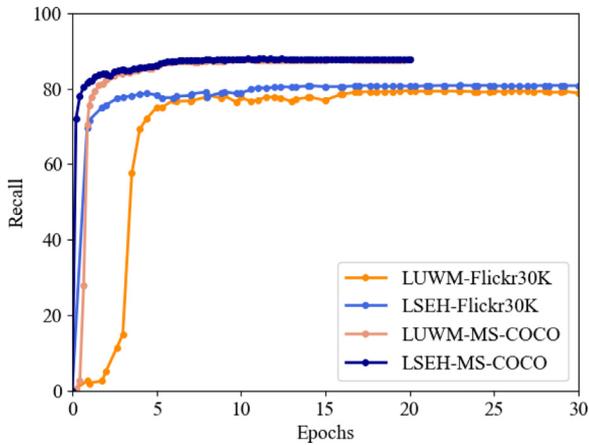
#### 4.4. Comparison of LSEH with polynomial loss

This section compares the training efficiency and cross-modal information retrieval performance of the proposed LSEH with the LUWM polynomial loss function [10]. In Wei et al. [10] LUWM was tested using GSMN, and hence the experiments included herein adopt the GSMN. This will facilitate the comparison of the performance of GSMN when using the LUWM and the proposed LSEH.

**Training Efficiency.** Fig. 5 compares the training efficiency of GSMN using LUWM and LSEH on the Flickr30K and MS-COCO

**Table 9**  
Average Recall@k (%) of GSMN when using the LUWM and LSEH loss functions for cross-modal information retrieval across datasets.

Network	Loss	Image-to-Text Retrieval				Text-to-Image Retrieval			
		R@1	R@5	R@10	Mean	R@1	R@5	R@10	Mean
Flickr30K									
GSMN [15]	LMH1	71.4	92.0	96.1	86.5 ( $\pm 10.8$ )	53.9	79.7	87.1	73.6 ( $\pm 14.2$ )
GSMN [10]	LUWM	73.1	92.7	<b>96.8</b>	87.5 ( $\pm 10.3$ )	54.2	79.9	<b>87.3</b>	73.8 ( $\pm 14.2$ )
GSMN	LSEH	<b>74.1</b>	<b>93.3</b>	96.5	<b>88.0 (<math>\pm 9.9</math>)</b>	<b>55.4</b>	<b>81.2</b>	87.2	<b>74.6 (<math>\pm 13.8</math>)</b>
MS-COCO1K									
GSMN [15]	LMH1	76.1	95.6	98.3	90.0 ( $\pm 9.9$ )	60.4	88.7	95.0	81.4 ( $\pm 15.0$ )
GSMN [10]	LUWM	76.8	96.2	98.5	90.5 ( $\pm 9.7$ )	60.9	89.0	<b>95.5</b>	81.8 ( $\pm 15.0$ )
GSMN	LSEH	<b>79.7</b>	<b>96.4</b>	<b>98.9</b>	<b>91.7 (<math>\pm 8.5</math>)</b>	<b>63.2</b>	<b>90.3</b>	95.4	<b>83.0 (<math>\pm 14.1</math>)</b>
MS-COCO5K									
GSMN [10]	LUWM	<b>54.7</b>	82.2	89.7	75.5 ( $\pm 15.0$ )	38.5	67.6	78.9	61.7 ( $\pm 17.0$ )
GSMN	LSEH	54.3	<b>82.4</b>	<b>90.2</b>	<b>75.6 (<math>\pm 15.4</math>)</b>	<b>39.0</b>	<b>68.4</b>	<b>79.5</b>	<b>62.3 (<math>\pm 17.1</math>)</b>



**Fig. 5.** Comparison of training efficiency when using LUWM and LSEH to train GSMN on Flickr30K and MS-COCO respectively.

datasets, where the validation step for GSMN is at every 1000 mini-batches. Fig. 5 shows that GSMN with LSEH achieved a larger M-Recall than with LUWM, at fewer epochs on both datasets.

**Cross-modal Information Retrieval Performance.** Table 9 shows the results of the comparisons of LUWM and LSEH when integrated into GSMN and applied to the datasets. The last column of Table 9 shows the mean average Recall values (i.e. the average of columns 3–5), and the main findings from each dataset are summarised as follows.

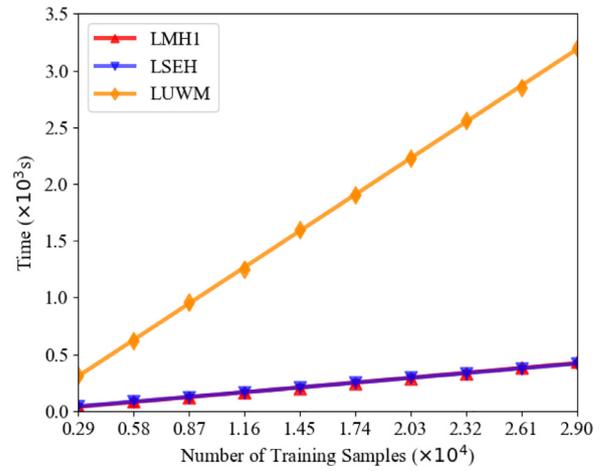
(1) **Flickr30K.** The proposed LSEH reached a Recall of 88.0% for image-to-text, a Recall of 74.6% for text-to-image retrieval, and outperformed LUWM by 0.5% and 0.8% for those tasks, respectively.

(2) **MS-COCO1K.** For image-to-text retrieval, LSEH’s Recall reached 91.7% which outperformed LMH1 by 1.2%. For text-to-image retrieval, LSEH’s Recall reached 83.0% which outperformed LUWM by 1.2%.

(3) **MS-COCO5K.** For LSEH, Recall reached 75.6% for image-to-text and 62.3% for text-to-image retrieval. LSEH outperformed LMH1 by 0.1% and 0.6% for image-to-text and text-to-image retrieval, respectively.

4.5. Discussion on quantitative results of graph-based VSE networks

Graph-based VSE networks such as SGRAF-SGR, VSRN, and GSMN consider the local information of images and descriptions for cross-modal information retrieval. The results described in Section 4.2–4.4 revealed that the above graph-based VSE networks performed better than when using the LMH and LUWM loss functions. When using LMH1, SGRAF-SGR needed 38.6 training epochs to achieve its higher M-Recall value (81.0%) on the Flickr30K



**Fig. 6.** Computation time at 1 epoch when using LMH1, LUWM, and LSEH to train VSE∞.

dataset, compared to needing 20.1 epochs when using LSEH. Hence when using LSEH the number of training epochs needed by SGRAF-SGR were reduced by 47.9%. In terms of retrieval performance using the mean average Recall evaluation metric, (1) VSRN using LSEH reached the Recall of 87.2% and 75.5% for image-to-text and text-to-image retrieval respectively on the Flickr30K dataset, and outperformed VSRN when using LMH1 by 2.8% and 3.0% for those tasks, respectively; and (2) GSMN using LSEH reached a Recall of 91.7% for image-to-text, and a Recall of 83.0% for text-to-image retrieval on the MS-COCO1K dataset, and outperformed GSMN when using LUWM by 1.2% and 1.2% for those tasks, respectively.

4.6. Comparisons on computation time

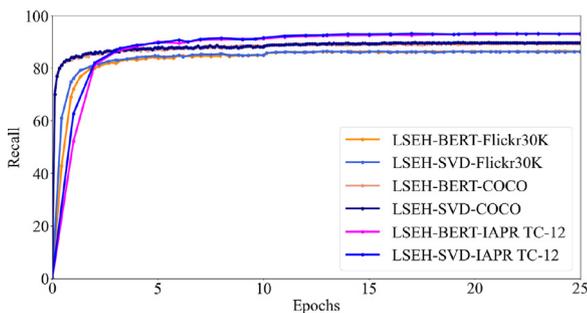
Fig. 6 shows the computation time of one epoch when training VSE∞ [8] using LMH1, LUWM, and LSEH. These experiments utilised the training set from Flickr30K [7]. In Fig. 6, for each loss function the network is trained 10 times with various numbers of training samples (increased from 2900 to 29,000 in steps of 2900), and the computation time is recorded. The three lines that fit the data points of LMH1, LUWM, and LSEH follow the equations of  $T_{LMH1}(n) = 147.8n - 6.4$ ,  $T_{LUWM}(n) = 1104.6n - 14.7$ , and  $T_{LSEH}(n) = 144.4n - 1.7$  respectively. Note that the lines of LSEH and LMH1 are almost aligned. The computation time of LSEH is almost the same as that of LMH1, and six times faster than LUWM.

4.7. Comparison of LSEH-SVD and LSEH-BERT

This section compares LSEH when using SVD (LSEH-SVD) and when using BERT (LSEH-BERT) for cross-modal information re-

**Table 10**  
Average Recall@k (%) of VSE $_{\infty}$  when using the LSEH-SVD and LSEH-BERT for cross-modal information retrieval across the datasets.

Loss	Method	Image-to-Text Retrieval				Text-to-Image Retrieval			
		R@1	R@5	R@10	Mean	R@1	R@5	R@10	Mean
					Flickr30K				
LMH1	-	80.8	<b>96.4</b>	98.3	91.8 ( $\pm 7.8$ )	62.6	86.9	91.7	80.4 ( $\pm 12.7$ )
LSEH	BERT	81.0	96.1	97.7	91.6 ( $\pm 7.5$ )	63.1	87.1	<b>92.7</b>	81.0 ( $\pm 12.8$ )
LSEH	SVD	<b>82.4</b>	96.0	<b>98.6</b>	<b>92.3 (<math>\pm 7.1</math>)</b>	<b>63.7</b>	<b>87.1</b>	92.5	<b>81.1 (<math>\pm 12.5</math>)</b>
					MS-COCO1K				
LMH1	-	80.8	97.0	<b>99.0</b>	92.3 ( $\pm 8.1$ )	66.0	91.7	96.1	84.6 ( $\pm 13.3$ )
LSEH	BERT	82.0	<b>97.1</b>	98.7	92.6 ( $\pm 7.5$ )	66.2	<b>92.0</b>	96.2	84.8 ( $\pm 13.3$ )
LSEH	SVD	<b>82.2</b>	96.9	98.6	<b>92.6 (<math>\pm 8.5</math>)</b>	<b>66.5</b>	91.9	<b>96.2</b>	<b>84.9 (<math>\pm 13.1</math>)</b>
					MS-COCO5K				
LMH1	-	58.3	84.7	91.8	78.3 ( $\pm 14.4$ )	42.5	72.7	83.0	66.1 ( $\pm 17.2$ )
LSEH	BERT	<b>59.1</b>	85.3	92.6	<b>79.0 (<math>\pm 14.4</math>)</b>	42.0	72.7	83.1	65.9 ( $\pm 17.4$ )
LSEH	SVD	58.8	<b>85.4</b>	<b>92.6</b>	78.9 ( $\pm 14.5$ )	<b>43.1</b>	<b>73.2</b>	<b>83.2</b>	<b>66.5 (<math>\pm 17.0</math>)</b>
					IAPR TC-12				
LMH1	-	81.5	97.3	<b>99.0</b>	92.6 ( $\pm 7.9$ )	79.1	96.4	<b>98.8</b>	91.4 ( $\pm 8.8$ )
LSEH	BERT	82.9	97.4	98.5	92.9 ( $\pm 7.1$ )	80.6	96.5	97.9	91.7 ( $\pm 7.8$ )
LSEH	SVD	<b>83.7</b>	<b>97.7</b>	98.8	<b>93.4 (<math>\pm 6.9</math>)</b>	<b>81.4</b>	<b>96.9</b>	98.5	<b>92.3 (<math>\pm 7.7</math>)</b>



**Fig. 7.** Comparison of training efficiency when using LSEH-SVD and LSEH-BERT to train VSE $_{\infty}$  on the Flickr30K, MS-COCO, and IAPR TC-12 datasets.

retrieval in the Flickr30K, MS-COCO, and IAPR TC-12 benchmark datasets. The experiments employ the latest VSE network (i.e. VSE $_{\infty}$ ) to test LSEH-SVD and LSEH-BERT. VSE $_{\infty}$  using LSEH-BERT employs the pre-trained BERT-base model [26] and follows the hyperparameter settings of VSE $_{\infty}$  using LSEH-SVD (shown in Table 2).

**Training Efficiency.** Fig. 7 compares the training efficiency of VSE $_{\infty}$  using LSEH-SVD and LSEH-BERT, where the validation step for VSE $_{\infty}$  is at every 500 mini-batches. Fig. 7 shows that VSE $_{\infty}$  with LSEH-SVD reached a larger M-Recall than with LSEH-BERT at fewer epochs on the Flickr30K and IAPR TC-12 datasets. Observing the curves of VSE $_{\infty}$  using LSEH-SVD and LSEH-BERT on the MS-COCO dataset, their lines are close to overlapping and hence they are similar with regards to M-Recall and training efficiency.

**Cross-modal Information Retrieval Performance.** As shown in Table 10, compared to the baseline (LMH1), both LSEH-SVD and LSEH-BERT have outperformed LMH1 for cross-modal information retrieval on all datasets. The performance of LSEH-BERT is lower than that of LSEH-SVD. LSEH-BERT outperformed LSEH-SVD only for the task of image-to-text retrieval on the dataset of MS-COCO5K, where LSEH-BERT reached a mean average Recall value of 79.0% and outperformed LSEH-SVD by 0.1%. However, LSEH-SVD outperformed LSEH-BERT on all the rest of the tasks and datasets in Table 10.

#### 4.8. Reasoning with the results of the proposed LSEH using visualisation

For reasoning with how the proposed LSEH improved the training of VSE networks, VSRN [13] was adopted because it is a state-

of-the-art image graph reasoning network that provides the means of visualising its results.

**Reasoning with Training Efficiency.** When the queries are looped in a mini-batch during training, LMH1 and LSEH select one irrelevant embedding from a set of irrelevant embeddings as the hard negative sample for the query embedding, because of the max operation in LMH1 and LSEH (as shown in Eq. (2) and Eq. (8)). However, some irrelevant embeddings may be repeatedly selected as the hard negative samples and this leads to low efficiency when utilising embeddings (see Section 3). In this experiment, to illustrate and compare the training efficiency of VSRN when using LMH1 and the proposed LSEH loss function, the number of unique image- and description- embeddings (denoted as #ImEmbs and #DesEmbs respectively), that are used as the hard negative samples, are counted within each mini-batch and across a range of training epochs.

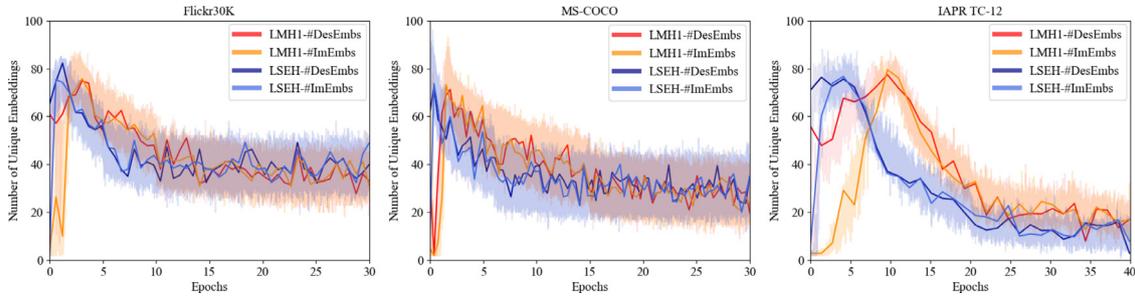
The results of each dataset are shown in Fig. 8 and described as follows.

(1) **Flickr30K.** LSEH reached its largest number of unique #DesEmbs and #ImEmbs (i.e. both 85) at 1.1 epochs, whereas LMH1 needed 2.9 epochs (i.e. #DesEmbs and #ImEmbs were both 87).

(2) **MS-COCO.** LSEH needed 0.2 epochs to reach its largest number of #DesEmbs (91) and #ImEmbs (i.e. 96), whereas LMH1 needed 1.7 epochs to obtain its largest number of unique #DesEmbs and #ImEmbs with values 89 and 93 respectively.

(3) **IAPR TC-12.** LSEH reached its largest number of unique #DesEmbs (i.e. 88) and #ImEmbs (i.e. 87) in 2.2 epochs, whereas LMH1 needed 10 epochs to obtain its largest number of unique #DesEmbs (i.e. 87) and #ImEmbs (i.e. 85). The proposed LSEH outperformed LMH1 on the efficiency of utilising unique embeddings, which visually reasons the improvement in training efficiency by LSEH.

**Illustrating one Hard Negative Sample Selected by LMH1 and LSEH.** For visually comparing LMH1 and LSEH when selecting hard negative samples, the irrelevant images to each textual query found in a mini-batch are ranked by LMH1 and LSEH respectively according to the descending order of their gradients. The top eight irrelevant images to a sample textual query are presented in Fig. 9. Images selected by LMH1 in Fig. 9 (a) contain no semantic similarity relevant to the query description. However, in comparison to the set of results retrieved by LSEH and shown in set (b) of Fig. 9, the selected images are more semantically relevant to the query, e.g. more images were retrieved of ‘a man’ or which relate to the clothes described in the query. The ranked top eight irrelevant images selected by LSEH have a higher probability of being incorrectly retrieved than those irrelevant images selected by LMH1,

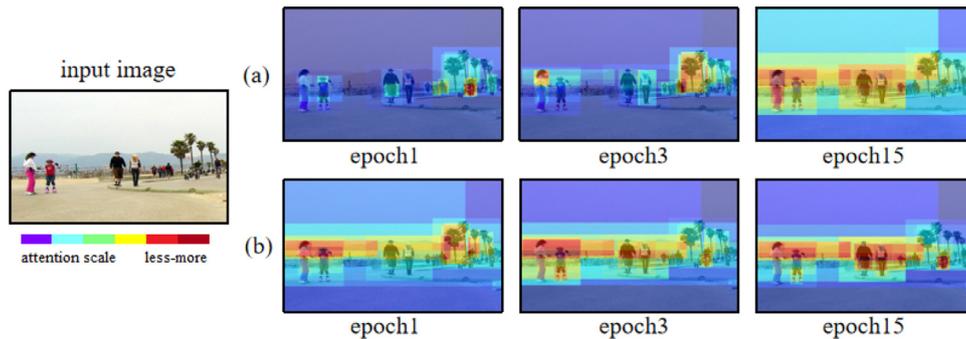


**Fig. 8.** Results of reasoning with training efficiency of VSRN using LMH1 and LSEH on various datasets. The Y-axis presents the numbers of unique image embeddings (#ImEmbs) and description embeddings (#DesEmbs) that are used as hard negative samples in each mini-batch, and the X-axis is the number of training epochs.

Query: A man wearing camouflage shorts and a white t-shirt shines the black dress shoes of a man wearing black dress slacks.



**Fig. 9.** Examples (a) and (b) are the top eight irrelevant images in a mini-batch to the query description that is selected by LMH1 and LSEH respectively according to the descending order of their gradients. The images in (b) are semantically closer to the query's textual description than the images in (a), which are hard negative samples.



**Fig. 10.** Visual comparison of LMH1 (a) and LSEH (b) on the various attention regions of an image processed by VSRN during training. The degree of attention is indicated on the attention scale.

hence why it is more efficient to use LSEH's hard negative samples during the training.

**Visualisation of Learning with LMH1 and LSEH.** For visually comparing the learning progress of VSRN when using LMH1 and LSEH, the various attention regions of an image processed by VSRN during the training are presented, where the method for generating the attention region heating maps follows that of Li et al. [13]. VSRN using LMH1, as shown in Fig. 10 (a) only focused some attention on the main objects of the image (e.g. 'people' and 'trees') on epochs one and three. On the contrary, as shown in Fig. 10 (b), the VSRN using the proposed LSEH has paid much attention to the main objects of the image on epoch one.

## 5. Conclusion

Existing Visual Semantic Embedding (VSE) networks are trained by a hard negatives loss function that learns an objective margin between the similarity of the relevant and irrelevant image-description embedding pairs and ignores the semantic differences between the irrelevant pairs. This paper proposes a novel Semantically-Enhanced Hard negatives Loss function (LSEH) for

Cross-modal Information Retrieval that considers the semantic differences between irrelevant training pairs, to dynamically adjust the learning objectives of VSE networks to make their learning flexible and efficient. Extensive experiments were carried out by integrating the proposed methods into five state-of-the-art VSE networks (VSE++, VSRN, SGRAF-SAF, SGRAF-SGR, and VSE $\infty$ ) that were applied to the Flickr30K, MS-COCO, and IAPR TC-12 datasets. The experiments revealed that the proposed LSEH function, when integrated into VSE networks, improves their training efficiency and cross-modal information retrieval performance. With regards to training time, LSEH reduced the training epochs of LMH on average by 53.2% on the Flickr30K dataset, 43.0% on the MS-COCO dataset, and 48.5% on the IAPR TC-12 dataset. In terms of retrieval performance using the mean average Recall evaluation metric, LSEH outperformed LMH1 by (1) 2.3% for image-to-text and 2.0% for text-to-image retrieval on the Flickr30K dataset; (2) 0.2% for image-to-text and 0.2% for text-to-image retrieval on the MS-COCO1K dataset; (3) 0.7% for image-to-text and 0.3% for text-to-image retrieval on the MS-COCO5K dataset; and (4) 3.5% for image-to-text and 2.8% for text-to-image retrieval on the IAPR TC-12 dataset. Section 4.7 describes our experimental results using SVD

and BERT and the findings revealed that SVD outperformed BERT when these are embedded in the proposed loss function. From this perspective, it is more efficient to use an unsupervised approach because such approaches do not need labels or computationally expensive training processes. The focus of the paper has been on the application of VSE networks for cross-modal information retrieval. Future work includes applying the proposed loss function to other similar cross-modal tasks such as hashing-based networks for cross-modal retrieval and text-to-video retrieval. Future work will also focus on extracting image description semantics using other supervised and unsupervised feature extraction methods and comparing those to the methods (i.e. SVD and BERT) that have already been utilised for the task. Finally, future work also includes embedding the proposed methods into a next generation cross-modal search engine and evaluating its capabilities with real-world datasets.

### Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Data availability

The experiments described in this paper use publicly available datasets. For instructions on how to download and prepare the datasets please visit the paper's Github repository (<https://github.com/yangong23/VSEnetworksLSEH>).

### References

- [1] Y. Gong, G. Cosma, H. Fang, On the limitations of visual-semantic embedding networks for image-to-text information retrieval, *Journal of Imaging* 7 (8) (2021) 125.
- [2] X. Shu, G. Zhao, Scalable multi-label canonical correlation analysis for cross-modal retrieval, *Pattern Recognit* 115 (2021) 107905.
- [3] F. Faghri, D.J. Fleet, J.R. Kiros, S. Fidler, VSE++: improving visual-semantic embeddings with hard negatives, in: *Proceedings of the British Machine Vision Conference*, 2018, p. 12.
- [4] H. Diao, Y. Zhang, L. Ma, H. Lu, Similarity reasoning and filtration for image-text matching, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35, 2021, pp. 1218–1226.
- [5] P. Hu, X. Peng, H. Zhu, J. Lin, L. Zhen, W. Wang, D. Peng, Cross-modal discriminant adversarial network, *Pattern Recognit* 112 (2021) 107734.
- [6] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C.L. Zitnick, Microsoft COCO: common objects in context, in: *Proceedings of the European Conference on Computer Vision*, 2014, pp. 740–755.
- [7] P. Young, A. Lai, M. Hodosh, J. Hockenmaier, From image descriptions to visual denotations: new similarity metrics for semantic inference over event descriptions, *Transactions of the Association for Computational Linguistics* 2 (2014) 67–78.
- [8] J. Chen, H. Hu, H. Wu, Y. Jiang, C. Wang, Learning the best pooling strategy for visual semantic embedding, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15789–15798.
- [9] G. Song, X. Tan, J. Zhao, M. Yang, Deep robust multilevel semantic hashing for multi-label cross-modal retrieval, *Pattern Recognit* 120 (2021) 108084.
- [10] J. Wei, Y. Yang, X. Xu, X. Zhu, H.T. Shen, Universal weighting metric learning for cross-modal retrieval, *IEEE Trans Pattern Anal Mach Intell* (2021).
- [11] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, L. Zhang, Bottom-up and top-down attention for image captioning and visual question answering, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6077–6086.
- [12] K. Cho, B.V. Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using RNN encoder-decoder for statistical machine translation, in: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 2014, pp. 1724–1734.
- [13] K. Li, Y. Zhang, K. Li, Y. Li, Y. Fu, Image-text embedding learning via visual and textual semantic reasoning, *IEEE Trans Pattern Anal Mach Intell* (2022).
- [14] S. Zhang, H. Tong, J. Xu, R. Maciejewski, Graph convolutional networks: a comprehensive review, *Computational Social Networks* 6 (1) (2019) 1–23.
- [15] C. Liu, Z. Mao, T. Zhang, H. Xie, B. Wang, Y. Zhang, Graph structured network for image-text matching, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 10921–10930.
- [16] G. Li, N. Duan, Y. Fang, M. Gong, D. Jiang, Unicoder-VL: a universal encoder for vision and language by cross-modal pre-training, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34, 2020, pp. 11336–11344.
- [17] Y. Duan, N. Chen, P. Zhang, N. Kumar, L. Chang, W. Wen, MS<sup>2</sup>GAH: multi-label semantic supervised graph attention hashing for robust cross-modal retrieval, *Pattern Recognit* 128 (2022) 108676.
- [18] X. Liu, Z. Hu, H. Ling, Y.M. Cheung, MTFH: a matrix tri-factorization hashing framework for efficient cross-modal retrieval, *IEEE Trans Pattern Anal Mach Intell* 43 (3) (2019) 964–981.
- [19] A. Karpathy, L. Fei-Fei, Deep visual-semantic alignments for generating image descriptions, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3128–3137.
- [20] Y. Liu, L. Liu, Y. Guo, M.S. Lew, Learning visual and textual representations for multimodal matching and classification, *Pattern Recognit* 84 (2018) 51–67.
- [21] L. Zhang, X. Wu, Multi-task framework based on feature separation and reconstruction for cross-modal retrieval, *Pattern Recognit* 122 (2022) 108217.
- [22] Y. Liu, Y. Guo, L. Liu, E.M. Bakker, M.S. Lew, CycleMatch: a cycle-consistent embedding network for image-text matching, *Pattern Recognit* 93 (2019) 365–379.
- [23] Y. Song, M. Soleymani, Polysemous visual-semantic embedding for cross-modal retrieval, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1979–1988.
- [24] K.-H. Lee, X. Chen, G. Hua, H. Hu, X. He, Stacked cross attention for image-text matching, in: *Proceedings of the European Conference on Computer Vision*, 2018, pp. 201–216.
- [25] H. Wang, Z. Ji, Z. Lin, Y. Pang, X. Li, Stacked squeeze-and-excitation recurrent residual network for visual-semantic matching, *Pattern Recognit* 105 (2020) 107359.
- [26] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, in: *Proceedings of Conference of the North American Chapter of the Association for Computational Linguistics*, Vol. 1: Long and Short Papers, 2019, pp. 4171–4186.
- [27] S. Lipovetsky, PCA and SVD with nonnegative loadings, *Pattern Recognit* 42 (1) (2009) 68–76.
- [28] G.W. Furnas, S. Deerwester, S.T. Durnais, T.K. Landauer, R.A. Harshman, L.A. Streeter, K.E. Lochbaum, Information retrieval using a singular value decomposition model of latent semantic structure, in: *ACM SIGIR Forum*, Vol. 51, ACM New York, NY, USA, 2017, pp. 90–105.
- [29] L. Maltoudoglou, A. Paisios, L. Lenc, J. Martínek, P. Král, H. Papadopoulos, Well-calibrated confidence measures for multi-label text classification with a large number of labels, *Pattern Recognit* 122 (2022) 108271.
- [30] E. Haddi, X. Liu, Y. Shi, The role of text pre-processing in sentiment analysis, *Procedia Comput Sci* 17 (2013) 26–32.
- [31] G. Cosma, T.M. McGinnity, Feature extraction and classification using leading eigenvectors: applications to biomedical and multi-modal mhealth data, *IEEE Access* 7 (2019) 107400–107412.
- [32] M. Grubinger, P. Clough, H. Müller, T. Deselaers, The IAPR TC-12 benchmark: a new evaluation resource for visual information systems, in: *International workshop ontolmage*, Vol. 2, 2006, pp. 13–23.
- [33] J. Martineau, T. Finin, Delta TFIDF: an improved feature space for sentiment analysis, in: *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 3, 2009, pp. 258–261.
- [34] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, *Adv Neural Inf Process Syst* 25 (2012) 1097–1105.
- [35] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D.A. Shamma, et al., Visual genome: connecting language and vision using crowdsourced dense image annotations, *Int J Comput Vis* 123 (1) (2017) 32–73.



**Yan Gong** is currently pursuing the doctor's degree in Computer Science at the Department of Computer Science at Loughborough University of UK. He received master's degree with distinction from Newcastle University of UK in 2012. His major research interests reside in cross- and multi- modal information retrieval, Nature Language Processing, Computer Vision, Artificial Intelligence, and Data Science.



**Georgina Cosma** is a Senior Lecturer in Natural Language Processing and AI at the Department of Computer Science at Loughborough University. Dr Cosma received a first class B.Sc. (Hons.) degree in computer science from Coventry University, U.K., in 2003, and the Ph.D. degree in computer science from the University of Warwick, U.K., in 2008. Dr. Cosma's research interests reside in cross- and multi- modal information retrieval, Artificial Intelligence, and Data Science.