



HAL
open science

Image search using multilingual texts: a cross-modal learning approach between image and text

Maxime Portaz, Hicham Randrianarivo, Adrien Nivaggioli, Estelle Maudet,
Christophe Servan, Sylvain Peyronnet

► To cite this version:

Maxime Portaz, Hicham Randrianarivo, Adrien Nivaggioli, Estelle Maudet, Christophe Servan, et al.. Image search using multilingual texts: a cross-modal learning approach between image and text. [Research Report] qwant research. 2019. hal-02077556

HAL Id: hal-02077556

<https://hal.science/hal-02077556v1>

Submitted on 25 Mar 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Image search using multilingual texts: a cross-modal learning approach between image and text

Maxime Portaz
Qwant Research
m.portaz@qwant.com

Hicham Randrianarivo
Qwant Research
h.randrianarivo@qwant.com

Adrien Nivaggioli
Qwant Research
a.nivaggioli@qwant.com

Estelle Maudet
Qwant Research
e.maudet@qwant.com

Christophe Servan
Qwant Research
c.servan@qwant.com

Sylvain Peyronnet
Qwant Research
s.peyronnet@qwant.com

Abstract

Multilingual (or cross-lingual) embeddings represent several languages in a unique vector space. Using a common embedding space enables for a shared semantic between words from different languages. In this paper, we propose to embed images and texts into a unique distributional vector space, enabling to search images by using text queries expressing information needs related to the (visual) content of images, as well as using image similarity. Our framework forces the representation of an image to be similar to the representation of the text that describes it. Moreover, by using multilingual embeddings we ensure that words from two different languages have close descriptors and thus are attached to similar images. We provide experimental evidence of the efficiency of our approach by experimenting it on two datasets: Common Objects in COntext (COCO) [19] and Multi30K [7].

1. Introduction

Neural networks can embed data into features vectors that were used primarily for information retrieval of texts. Evolution of these embeddings (namely multilingual embeddings) made it possible to solve multilingual tasks such as cross-lingual classification of texts. Another important task of information retrieval is to be able to deal with text and image queries.

Cross-modal networks are using recurrent and convolutional networks together in order to embed texts and images in a common vector space. There are plenty of literature on these topics, we present a selection of the relevant papers in section 2.

While many approaches embed visual and semantic information together, they are most of the time limited to only one

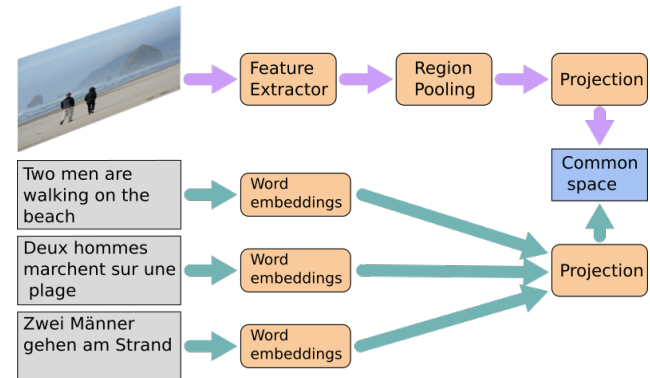


Figure 1: Overall pipeline of our method. The method is decomposed into two path. The visual path that extracts a representation from the image and the text path that extracts a representation from texts. The text path project texts of different languages in the same vector space.

language. In this paper, we propose a novel approach for multilingual text and image embeddings (section 3). Our method uses Convolutional Neural Networks (CNN) to extract Image information and aligned multilingual Word embeddings and Recurrent Neural Networks (RNN) to produce a text representation. With this framework, we provide image and text embeddings that can be trained to produce comparable features. We can thus retrieve images from text, and texts from image, with a multilingual representation.

More precisely, we propose two approaches. The first uses Bilingual Word Representations (BIVector) embeddings (see [20]) and improves on the state-of-the-art for English when trained on another language. The second uses Multilingual Unsupervised or Supervised word Embeddings (MUSE) (see [5]). It enables recognition for several languages with only one model, with a slight performance drawbacks. More precisely, by using MUSE aligned embeddings

in 30 different languages, we can retrieve images with languages never seen by the model.

The major lesson learned with respect to our method is that it provides close to the state-of-the-art performance in English-only context, but enables the use of multilingual datasets to improve results. We show in section 4 that adding a new language during the training phase improves image retrieval for other languages. The experiments show that the approach based on BIVVEC embeddings gives a 3.35% increase in performance on the COCO dataset, and a 15.15% increase on Multi30K. By using MUSE, we are able to encode more languages in the same model. The experiments for that approach show that adding other languages gives a small decrease of performance for English, but increases the recall for a multilingual environment. Indeed, we obtain 49.38% recall@10 on Multi30K dataset for image retrieval from captions in 4 languages.

2. Related Work

2.1. Text embeddings

The use of word representation is an important step when it comes to search information from text documents. In order to perform this task, we want to be able to extract meaningful embeddings from words. One useful property of word embeddings is that words with a similar meaning must have representations with a close distance. A lot of people work on this task but the most popular methods are Word2Vec (W2V) [21] and FastText [3]. These methods are simplified methods of the neural language model proposed by Bengio et al. [1] with several tricks to boost performance.

2.2. Multilingual word embeddings.

Word embeddings can be used in multilingual tasks (e.g. machine translation or cross-lingual document classification) by training a model independently for each language. However, the representations will be in distinct vector spaces, which means that the same words in different languages will most likely have different representations.

There are several methods to solve this problem. One consists in training both models independently and then to learn a mapping from one representation to the other [10, 21]. It is also possible to constrain the training to keep the representations of similar words close to each other [13, 22]. Finally, the training can be performed jointly using parallel corpora [17].

In the latter category, the BIVVEC approach [20] tries to predict words based on the inner context of the sentence like W2V does, but also uses words in the source sentence to predict words in the target sentence (and conversely). Thus, for each update in W2V, BIVVEC performs 4 updates: source to source, source to target, target to target and target to source. This leads to a common representation for the two

languages.

Recently, MUSE [5] proposes to learn a mapping from several word embeddings trained independently. This approach enables a mapping between word embeddings from different languages.

2.3. Cross-modal representation

In order to provide queries as sentences or as images, the image embeddings and the text embeddings must be comparable, i.e. in the same representation space. Recent works have shown the possibility to learn text and image representation simultaneously [15, 9, 8]. They rely on cross-modal networks, that are able to extract information from images and read the caption describing it. Those networks use word embeddings followed by an RNN to encode sentence embeddings in the same space as the image embeddings, extracted with a CNN.

CNN methods provide ways to encode images in meaning full embeddings. Prior works considered image similarity based on the categories [12, 26]. Recent approaches [9, 8] use ResNet [27] as CNN image features extractor. For the text part on the network, they use W2V or Skip-Thought [16, 8] word embeddings, followed by a multi-layer RNN to encode the sentence.

Those methods are multi-branches networks. The loss function has to be a comparative loss, with a similarity function. The similarity function is generally estimated using the euclidean distance, or the cosine similarity. To train network with this type of loss function, we use Siamese network [4, 2, 28]. Siamese networks have been extended to triplet networks with three branches in order to give better results [14, 25]. Triplet are composed of an anchor image, an example of a similar image (positive image) and a dissimilar example (negative image).

Those methods enable to learn complex image representation with few examples, as it is possible to select the best triplet example for the training [26, 11, 23]. Our model is based on triplet networks, with each branch based on text or on images, interchangeably.

3. Multilingual Joint Text Image Embedding

We present a model for multilingual and cross-modal (joint learning of text and image representation) embeddings. We propose to embed images and sentences from different languages in the same space $[-1, 1]^d$, where the distance between two elements (image or sentence) is inversely proportional to their similarity.

To do this, we train a triplet network [14, 25] that compares an image to two sentences: one that describes the image, another that does not (See figure 2).

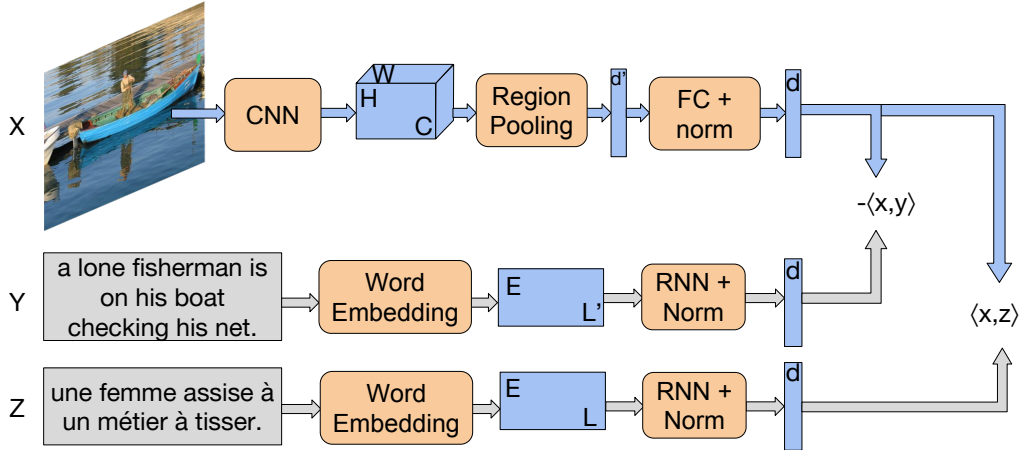


Figure 2: Proposed multilingual text image embedding architecture. This pipeline shows how during the training phase how the model learn how to match an image with a sentence. X and Y are an image and a sentence that match each other. Z is an unrelated sentence. For the loss computation we want the dissimilarity between X and Y to be as big as possible and the dissimilarity between X and Z as small as possible.

3.1. Overview

Our framework enables to take advantage of the availability of multilingual corpora in order to learn a cross-modal representation between texts and images. We present a pipeline that learns a common representation between texts in different languages and images. This common representation enables to compare image and texts using similarity measure for fast information retrieval.

The figure 1 illustrate the different step of the method to extract a representation from texts or images. We show that we can learn a model on several languages and obtain state-of-the-art results on several retrieval task. We also show that our method can improve state-of-the-art methods and generalize on languages the model have not seen during the training but which are model by MUSE. Our method is composed of two paths.

One path extract information from an image using a feature extractor like ResNet. Then a feature pooling method [6] enables to extract a signature from the features. We use the Weldon pooling method which automatically select areas with highest and lowest activation and compute the image signature using these areas.

The second path compute a embedding of the input text using MUSE. This method enables to compute close embeddings for words with close meanings in different languages. A RNN is then use to extract a representation from the set of embeddings.

One interesting property of the Weldon pooling is that it produce a mapping between the highest response in the image path and the most significant words in the text path. This property enables to extract an accurate representation

between text and images, enabling to search precisely images with text and vice versa.

3.2. Multilingual Sentence Embeddings

To embed the sentence, our model first relies on individual word embeddings. As we want to embed every sentence from every language in a unique vector space, we use word embeddings aligned in different languages. The usual method consists in using BIVVEC [20]. BIVVEC aligns two languages in the same space by learning the word embeddings on the two languages simultaneously.

There are 4 languages in our captions: English, French, German, Czech. On the one hand, we propose to use pre-trained embeddings from MUSE. MUSE is a multilingual extension of FastText [5] that embeds and aligns 30 languages in a single vector space. On the other hand, we propose to jointly use BIVVEC and MUSE approaches in order to enhance our multilingual representation.

The main idea is to train independently several bilingual word embeddings, in which, one of the language is English. Then, we learn a mapping between the different English representations (from the several bilingual word embeddings) to maximize the link between the bilingual representations. For instance, one can train two bilingual word embedding models like English-French and English-Czech representations, apply the MUSE approach on the two English parts. From the two bilingual representation (English-French and English-Czech), we obtain a third one: French-Czech.

As shown on figure 3, we combine sentences from different languages by using word embedding models from different languages according to the approach described previously. Those word embeddings being in the same space we

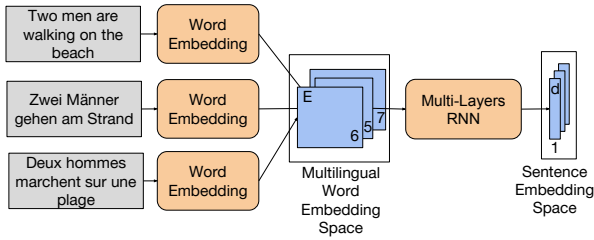


Figure 3: Sentences from different languages are mapped to a common Sentence embedding space. The word embedding method project words from different languages into a common space. For each word of the input sentence an embedding is computed and a RNN is used to extract the embedding of the whole sentence.

can use a multi-layer RNN to learn a sentence embedding . This network is composed of 4 layers of Simple Recurrent Unit (SRU) [18], with a dropout after each layer. The goal of this RNN is to encode a vector of word embeddings of size E into a Sentence Embedding Space \mathbb{R}^d . Lastly, we normalize the output of the RNN to obtain an embedding of the sentence in $[-1, 1]^d$.

3.3. Joint embedding

The visual path of our network is similar to the one used by Engilberge et al. [8]. It starts with a ResNet152 [27], on which we replaced its last fully connected layer (usually used for classification) with a Weldon pooling layer [6]. This layer pulls the regions with the maximum activation in the network, i.e. the regions of interest, and gives us an embedding of the image, a vector of dimension d' . Finally, this vector goes through a fully connected layer that normalizes it, which aims to obtain an embedding of the image in $[-1, 1]^d$.

Both pipelines are learned simultaneously, each image being paired with two sentences, one that describes the image, the other that doesn't. The architecture of the model is shown on figure 2. The two outputs are compared using a cosine similarity, which is equivalent to the inner product as both embeddings are normalized.

We use a triplet loss [26, 25, 11] to converge correctly and increase our performances. This loss enables us to compare the relative similarity between the image and both sentences: the corresponding sentence should be closer to the image than the unrelated one.

The Figure 2 presents the model with a triplet of one image and two captions. The sentence Y describes the image, and the caption Z is an unrelated caption. The triplet loss is shown in equation 1, with x , y , and z being respectively the embeddings of X , Y , and Z . α is minimum margin be-

tween the similarity of the correct caption and the unrelated caption. During the training, it was set to 0.2.

$$loss(x, y, z) = \max(0, \alpha - x \cdot y + x \cdot z) \quad (1)$$

3.4. Training

We train the ResNet on a classification task. This enables the CNN to learn the extraction of interesting image features. We used a ResNet-152 pretrained on the ImageNet [24] dataset, which provides a large collection of images, over 1 million, for 1000 categories. The last layer of the ResNet, that was used for classification, is removed and replaced by a Weldon pooling, followed by a randomly initialized fully connected layer with a dropout regularization.

For the text pipeline, we use pre-trained W2V, FastText, BIVVEC and MUSE word embeddings. We then freeze the CNN and region pooling of the network and train the RNN and the fully connected layer. This enables to project the embeddings of both sentence and image into a common space. Finally, we fine-tune the entire network.

As shown by Schroff et al. [25], the triplets used to train the model have to be carefully selected. Indeed, by using “easy” triplets, i.e. triplets on which the network performs already well, the network learns almost nothing. As proposed previously in [9, 8], we aim to focus on the “hardest” triplets only, i.e. the hardest ones to differentiate by the network. Instead of looking for the best triplet throughout the entire dataset at each iteration, we stay in the current batch. For each image and its corresponding sentence (i, s) , we select the closest non-similar image to s and the closest non-similar sentence to i .

The following equation shows the loss computation over the batch B composed of couple of image and caption (i, s) :

$$\sum_{i,s \in B} \left(\max_{z \in U_i} loss(i, s, z) + \max_{z \in D_i} loss(s, i, z) \right) \quad (2)$$

Where D_i represents every image in the batch B that differs from i . U_i represents every sentences unrelated to the image i , inside the batch B , in every languages. Each batch can contain different captions in different language corresponding to the same image. This enables the selection of the best example inside each training batch.

3.5. Visualization

To enhance the recall evaluation made previously, we propose some visual evaluation. The figure 4 shows the five closest images for the same sentence in French and German. These images comes from the Google Semantic Caption dataset, which contains 3 Millions images.

We show, in Figure 5, the maximum of activation of the network, given different words.



Figure 4: Closest images for the same sentence in different languages from the Google Semantic Caption dataset. Although the sentence is the same in 3 different languages, we can see that the results are slightly different. This is explained by the words embeddings which are close between languages but not the same.

We observe the activation zone of the CNN depending on the word used for the RNN. The network responds mostly in the same way for words from different languages. Some differences appear with less common words like "Mantel" in German, with a noisier activation than for the French "Man-teau" or the English "Jacket".

4. Experiments

In this section we present our experimental protocol: hardware/software setup, datasets used and numerical results. We evaluate and compare our method with state-of-the-art approaches by using classic metrics.

4.1. Experimental setup

All the experiments were done on an NVIDIA DGX-1¹. Training the network on COCO and Multi30k take roughly two days on 4 V100 GPUs with 16GB of RAM, for each experiment. We used Facebook implementation of FastText² to compute word embeddings. Finally, we rely on Pytorch³ for deep learning implementation.

4.2. Datasets

To train and evaluate our model, we used three datasets of images with their corresponding captions. The first dataset is COCO [19]. It contains 123 287 images with 5 English

captions per image. We used the *val split* from Karpathy et al. [15] (113 287 train, 5000 validation and 5000 test images) to train and evaluate our model on English sentences.

To train our model on other languages, we used the Multi30K [7] dataset, containing 31 014 images with captions in French, German, and Czech. 29 000 are kept for training, 1014 for validation and 1000 for testing. Lastly, for evaluation purposes we used Google's Conceptual Captions⁴ dataset, containing 3 154 240 captioned images.

4.3. Evaluation method

We evaluate the quality of our results using *recall@k*, which is the proportion of relevant images found in the top-*k* returned images for a given query. We evaluate caption retrieval with images as queries using recall at 1, 5 and 10. This means that we verify if the sentence corresponding to an image is in the first, fifth, or tenth closest results. For image retrieval, each caption is evaluated in the same way. The presence of the image corresponding to a sentence is verified in the first, fifth, or tenth closest results.

The caption retrieval test is made in batches of 1000 images and caption pairs, using the COCO dataset. On the Multi30K dataset, each image has a caption in each language. The recall is computed across languages.

¹<https://www.nvidia.com/>

²<https://fasttext.cc/>

³<https://pytorch.org/>

⁴<https://ai.google.com/research/ConceptualCaptions>




Figure 5: Activation maps of the network with different input words in different languages. The first column shows the activation map for English, the second for French, the third for German and the last for Czech. These activations show the ability of the network at recognize the important areas of the image according to the input word. This confirms that the method is able to match an object in the image with an associated word.

4.4. Results and analysis

We perform three experiments. The aim of the first experiment is to verify the performance of our model with English captions, depending of the word embeddings used.

The second experiment is similar to the first one, but measures image retrieval instead of caption retrieval. Experiment 1 and 2 measure each model performances in English only.

Finally, in experiment 3 we evaluate the model, with image recall, on the Multi30K dataset, with captions in different languages.

Experiment 1. The models are trained on the COCO dataset for English and on a Multi30K dataset for French, German and Czech. We use the COCO dataset for evaluation.

The table 1 shows the caption retrieval recall on COCO dataset. The first two lines show the state-of-the-art results.

Table 1: Experiment 1: Caption retrieval on the COCO dataset. We compare the different reminders of the different methods first on English and then by adding new languages. We also evaluate variations of DSVE method with different word embedding.

Embedding	lang.	r@1	r@5	r@10
VSE++ [9]	en	64.60	∅	95.70
DSVE [8]	en	69.8	91.9	96.60
DSVE w/ W2V	en	63.48	89.48	95.64
DSVE w/ FastText	en	66.08	90.70	96.20
Ours w/ BIVEC	en	65.58	90.52	96.10
	en+fr	67.78	91.58	96.92
Ours w/ MUSE	en	63.10	89.58	95.56
	en+fr	63.88	89.20	95.24
	en+fr+de	62.40	89.18	95.16
	all	63.28	88.30	94.60

The second pair of lines presents the results of our model, with W2V and FastText embeddings used as baseline. We can see that our model is close to the Deep Semantic-Visual Embedding (DSVE) method [8] while the W2V method is slightly worst, as the representation power of the word embedding is reduced.

The BIVEC English-French method is used on English and on both languages simultaneously. If trained only on English, i.e. only on the COCO dataset like the two previous methods, it shows performance similar to the one of the the state-of-the-art. This means training using BIVEC does not weaken the English representation. When trained on English and French together, the recall is increased by 3.35 %, going from 65.58 % to 67.78 %. We can also see an improvement for recall@5 and recall@10, with respectively 1.17 % and 0.85 % of increase. This imply that the similarity learning with French captions increases the English recognition when using BIVEC.

To verify if we can generalize this result to a larger number of languages, we used MUSE aligned for 30 languages. Using the Multi30K dataset, we can also train the model on German (de) and Czech (cs).

First of all, when training with MUSE for English only, we can see a sharp decrease of performance, with a recall going from 66.08 % to 63.10 %. By comparing the model trained with W2V, we obtain similar results. This could come from the fact that both MUSE and W2V embeddings do not have representation for out of vocabulary words like the FastText ones. Moreover, rare words have much more chance to be wrongly projected because of the space transformation. When we train the model with additional languages, we can see a slight decrease of performance in English. The maximum decrease is 1.01 % for recall@10, but it is counter-balanced by an increase of 0.29 % for the recall@1.

Table 2: Experiment 2: Image retrieval on the COCO dataset. The methods are the same as in table 1.

Embedding	lang.	r@1	r@5	r@10
VSE++ [9]	en	52.00	∅	92.00
DSVE [8]	en	55.90	86.90	94.00
DSVE w/ W2V	en	51.87	84.31	92.48
DSVE w/ FastText	en	54.12	85.74	92.93
Ours w/ BIVEC	en	55.57	86.92	93.86
	en+fr	56.09	87.22	94.03
Ours w/ MUSE	en	51.81	84.70	92.82
	en+fr	52.25	84.72	92.74
	en+fr+de	51.17	84.09	92.22
	all	50.44	83.39	91.80

Experiment 2. Given a sentence, in any language, we evaluate the rank of the corresponding image. The evaluation is again made by batches of 1000. The results are presented in table 2.

The first two lines of the table present the state-of-the-art results, with W2V and FastText embeddings. We can see similar results as in the previous experiment. With BIVEC, we have results close to the FastText embeddings when training only in English. This time, the recall is better with an increase of 2.68 % for recall@1. When trained with English and French, the recall@1 is increased by 3.65 %. This implies, again, that we can improve performance by learning on an additional language.

Our model is able to use the multi-language representing power of MUSE embeddings. We train the model with English, and different combinations of French, German and Czech. On English only, we have similar results to the W2V approach. When adding new languages, we can see a decrease in performance for English. We obtain a maximum decrease of 2.62 % for recall@1 when the models saw English, French, German and Czech.

Experiment 3. The model is trained with English only, then with English and French (en+fr), with English, French and German (en+fr+de) and with English, French, German and Czech (all). We can see a decrease in performance when adding French that is not present with other languages. Otherwise, every time we add a new language the recall for this language logically increase. The best performance is achieved with English+French+German+Czech, with an increase of 6.42 % for multilingual retrieval.

Experiment 4: With BIVEC embeddings, we learn two languages at the same time, and test retrieval on one or two of these languages. Results are shown in table 4. Trained on English alone, the model gives worse performance than

Table 3: Image Recall@10 on the Multi30k dataset with different languages with MUSE.

train. lang.	en	fr	de	cs	all
en	56.60	46.05	44.18	38.75	46.40
en+fr	50.93	43.69	41.61	34.02	42.43
en+fr+de	54.63	46.94	45.07	38.26	46.22
all	55.32	49.30	46.84	46.06	49.38

Table 4: Image Recall@10 on Multi30k dataset with different languages with BIVEC Embeddings.

train. lang.	en	fr	de	en+fr	en+de
en	53.35	26.13	22.96	39.74	34.57
en+fr	59.76	55.22	∅	57.50	∅
en+de	61.44	∅	43.59	∅	52.51

MUSE for languages not seen previously. For example, with English-German BIVEC and a model trained only in English, and test on German, we obtain only 22.96 % recall@10, where MUSE embeddings obtain 44.18 %. But when train on English and French, we obtain 55.22 % recall, an increase of 26.39 % compared to MUSE. With German and English training, we have an increase of 15.16 % on English only recall, with a recall of 61.44 %. Meaning that, once again, learning a new language with BIVEC enables better results in English, as same kind of results are visible with French as well.

5. Conclusion

We presented a novel approach for multilingual text and image embeddings. While the method provides close to the state-of-the-art performance in English-only context, its main advantage is that it enables to use multilingual dataset in order to improve the performance. We showed that using a new language during the training process improve image retrieval for other languages.

Our method uses a CNN to extract image information. It also uses aligned multilingual word embeddings and a RNN to produce text representations. This way, it provides image and text embeddings that can be trained to produce comparable features. We demonstrate that we can use this network to retrieve image from text and text from image, with a multilingual representation.

We evaluated our method on the COCO dataset for English-only results, and shown that using BIVEC embeddings enables the use of another language in order to improve the performance. The obtained improvement is a 3.35 % increase in performance on the COCO dataset, and a 15.15 % increase on the Multi30K dataset. By using MUSE embeddings, we are able to embed more languages in the

same model. We showed that adding other languages decrease performance for English, but increase the recall in a multilingual environment. For image retrieval from caption in 4 languages, we obtain a 49.38 % recall@10 on the Multi30K dataset.

References

- [1] Y. Bengio, R. Ducharme, P. Vincent, and C. Janvin. A Neural Probabilistic Language Model. *The Journal of Machine Learning Research*, 2003. [2](#)
- [2] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. S. Torr. Fully-Convolutional Siamese Networks for Object Tracking. In *European Conference on Computer Vision*, 2016. [2](#)
- [3] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics*, 7 2016. [2](#)
- [4] J. Bromley, I. Guyon, Y. LeCun, E. Sackinger, R. Shah, and C. Moore. Signature Verification using a "Siamese" Time Delay Neural Network. In *Advances in Neural Information Processing Systems*, 8 1993. [2](#)
- [5] A. Conneau, G. Lample, M. Ranzato, L. Denoyer, and H. Jégou. Word Translation Without Parallel Data. In *International Conference on Learning Representations*, 2018. [1](#), [2](#), [3](#)
- [6] T. Durand, N. Thome, and M. Cord. WELDON: Weakly supervised learning of deep convolutional neural networks. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2016. [3](#), [4](#)
- [7] D. Elliott, S. Frank, K. Sima'an, and L. Specia. Multi30K: Multilingual English-German Image Descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, Stroudsburg, PA, USA, 2016. Association for Computational Linguistics. [1](#), [5](#)
- [8] M. Engilberge, L. Chevallier, P. Pérez, and M. Cord. Finding beans in burgers: Deep semantic-visual embedding with localization. In *Conference on Computer Vision and Pattern Recognition*. IEEE, 6 2018. [2](#), [4](#), [6](#), [7](#)
- [9] F. Faghri, D. J. Fleet, J. R. Kiros, and S. Fidler. VSE++: Improving Visual-Semantic Embeddings with Hard Negatives. In *British Machine Vision Conference*, 2018. [2](#), [4](#), [6](#), [7](#)
- [10] M. Faruqui and C. Dyer. Improving Vector Space Word Representations Using Multilingual Correlation. In *European Chapter of ACL*, 2014. [2](#)
- [11] A. Gordo, J. Almazán, J. Revaud, and D. Larlus. End-to-End Learning of Deep Visual Representations for Image Retrieval. *International Journal of Computer Vision*, 2017. [2](#), [4](#)
- [12] M. Guillaumin, T. Mensink, J. Verbeek, and C. Schmid. Tag-Prop: Discriminative metric learning in nearest neighbor models for image auto-annotation. In *International Conference on Computer Vision*. IEEE, 9 2009. [2](#)
- [13] K. M. Hermann. Multilingual Distributed Representations without Word Alignment. In *International Conference on Learning Representations*, 2014. [2](#)
- [14] E. Hoffer and N. Ailon. Deep Metric Learning Using Triplet Network. In *International Conference on Learning Representations*, 2015. [2](#)
- [15] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. [2](#), [5](#)
- [16] R. Kiros, Y. Zhu, R. R. Salakhutdinov, R. Zemel, R. Urtasun, A. Torralba, and S. Fidler. Skip-thought vectors. In *Advances in neural information processing systems*, 2015. [2](#)
- [17] A. Klementiev, I. Titov, and B. Bhattacharai. Inducing Crosslingual Distributed Representations of Words. In *COLING*, 2012. [2](#)
- [18] T. Lei, Y. Zhang, S. I. Wang, H. Dai, and Y. Artzi. Simple Recurrent Units for Highly Parallelizable Recurrence. In *Conference on Empirical Methods in Natural Language Processing*, 9 2017. [4](#)
- [19] T. Lin, M. Maire, S. Belongie, and J. Hays. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision*, pages 1–14, 2014. [1](#), [5](#)
- [20] T. Luong, H. Pham, and C. D. Manning. Bilingual Word Representations with Monolingual Quality in Mind. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, Stroudsburg, PA, USA, 2015. Association for Computational Linguistics. [1](#), [2](#), [3](#)
- [21] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems*, 10 2013. [2](#)
- [22] S. C. A. P, S. Lauly, H. Larochelle, M. M. Khapra, B. Ravindran, V. Raykar, and A. Saha. An Autoencoder Approach to Learning Bilingual Word Representations. In *Advances in Neural Information Processing System*, 2014. [2](#)
- [23] M. Portaz, M. Kohl, G. Quenot, and J. P. Chevallet. Fully Convolutional Network and Region Proposal for Instance Identification with Egocentric Vision. In *IEEE International Conference on Computer Vision Workshops*, 2017. [2](#)
- [24] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 9 2015. [4](#)
- [25] F. Schroff, D. Kalenichenko, and J. Philbin. FaceNet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. [2](#), [4](#)
- [26] J. Wang, Y. Song, T. Leung, C. Rosenberg, J. Wang, J. Philbin, B. Chen, and Y. Wu. Learning Fine-Grained Image Similarity with Deep Ranking. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 6 2014. [2](#), [4](#)
- [27] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [2](#), [4](#)
- [28] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Deep Metric Learning for Person Re-identification. In *International Conference on Pattern Recognition*. IEEE, 8 2014. [2](#)