

Lessons learned in multilingual grounded language learning

Ákos Kádár
Tilburg University
a.kadar@uvt.nl

Desmond Elliott*
University of Copenhagen
de@di.ku.dk

Marc-Alexandre Côté
Microsoft Research Montreal
macote@microsoft.com

Grzegorz Chrupała
Tilburg University
g.chrupala@uvt.nl

Afra Alishahi
Tilburg University
a.alishahi@uvt.nl

Abstract

Recent work has shown how to learn better visual-semantic embeddings by leveraging image descriptions in more than one language. Here, we investigate in detail which conditions affect the performance of this type of grounded language learning model. We show that multilingual training improves over bilingual training, and that low-resource languages benefit from training with higher-resource languages. We demonstrate that a multilingual model can be trained equally well on either translations or comparable sentence pairs, and that annotating the same set of images in multiple language enables further improvements via an additional caption-caption ranking objective.

1 Introduction

Multimodal representation learning is largely motivated by evidence of perceptual grounding in human concept acquisition and representation (Barsalou et al., 2003). It has been shown that visually grounded word and sentence-representations (Kielbaso et al., 2014; Baroni, 2016; Elliott and Kádár, 2017; Kielbaso et al., 2017; Yoo et al., 2017) improve performance on the downstream tasks of paraphrase identification, semantic entailment, and multimodal machine translation (Dolan et al., 2004; Marelli et al., 2014; Specia et al., 2016). Multilingual sentence representations have also been successfully applied to many-languages-to-one character-level machine translation (Chung et al., 2016) and multilingual dependency parsing (Ammar et al., 2016).

Recently, Gella et al. (2017) proposed to learn both bilingual and multimodal sentence representations using images paired with captions independently collected in English and German. Their results show that bilingual training improves image-sentence ranking performance over a monolingual

baseline, and it improves performance on semantic textual similarity benchmarks (Agirre et al., 2014, 2015). These findings suggest that it may be beneficial to consider another language as another *modality* in a monolingual grounded language learning model. In the grounded learning scenario, descriptions of an image in multiple languages can be considered as multiple views of the same or closely related data. These additional views can help overcome the problems of data sparsity, and have practical implications for efficiently collecting image-text datasets in different languages. In real-life applications, many tasks and domains can involve code switching (Barman et al., 2014), which is easier to deal with using a multilingual model. Furthermore, it is more convenient to maintain a single multilingual system than one system for each considered language. However, there is a need for a systematic exploration of the conditions under which it is useful to add additional views of the data. We investigate the impact of the following conditions on the performance of a multilingual grounded language learning model in sentence and image retrieval tasks:

Additional languages. Multilingual models have not been explored yet in a multimodal setting. We investigate the contribution of adding more than one language by performing bilingual experiments on English and German (Section 5) as well as adding French and Czech captioned images (Section 6).

Data alignment: We assess the performance of a multilingual models trained using either captions that are translations of each other, or captions that are independently collected in different languages for the same set of images. The two scenarios are illustrated in Figure 1. Additionally we consider the setup when non-overlapping sets of images and their

*Work carried out at the University of Edinburgh.



En: A group of people are eating noodles.
De: Eine Gruppe von Leuten isst Nudeln.
Fr: Un groupe de gens mangent des nouilles.
Cs: Skupina lidí jedí nudle.

(a) A translation tuple

En: Several asian people eating around a table.
De: Drei Männer und zwei Frauen südostasiatischen Aussehens sitzen, aus Schälchen essend, an einem schwarzen, Tisch, auf dem sich u.a. auch Pappbecher und eine Tasche befinden, im Hintergrund sind weitere Personen und Tische.¹

(b) A comparable pair

Figure 1: An example taken from the *Translation* and *Comparable* portions of the Multi30K dataset. The translation portion (a) contains professional translations of the English captions into German, French, and Czech. The comparable portion (b) consists of five independently crowdsourced English and German descriptions, given only the image. Note that the sentences in (b) convey different information from the English–German translation pair in (a).

captions are collected in different languages. Such disjoint settings have been explored in pivot-based multimodal representation learning (Funaki and Nakayama, 2015; Rajendran et al., 2015) or zero-shot multi-modal machine translation (Nakayama and Nishida, 2017). We compare translated vs. independently collected captions in Sections 5.2 and 6.1, and overlapping vs. disjoint images in Section 5.3.

High-to-low resource transfer: In Section 6.2 we investigate whether low-resource languages benefit from jointly training on larger data sets from higher-resource languages. This type of transfer has previously been shown to be effective in machine translation (e.g., Zoph et al., 2016).

Training objective: In addition to learning to map images to sentences, we study the effect of also learning relationships between captions of the same image in different languages Gella et al. (2017). We assess the contribution of such a caption–caption ranking objective throughout our experiments.

Our results show that multilingual joint train-

ing improves upon bilingual joint training, and that grounded sentence representations for a low-resource language can be substantially improved with data from different high-resource languages. Our results suggest that independently-collected captions are more useful than translated captions, for the task of learning multilingual multimodal sentence embeddings. Finally, we recommend to collect captions for the same set of images in multiple languages, due to the benefits of the additional caption–caption ranking objective function.

2 Related work

Learning visually grounded word-representations has been an active area of research in the fields of multi-modal semantics and cross-situational word-learning. Such perceptually-grounded word representations have been shown to lead to higher correlation with human judgements on word-similarity benchmarks such as WordSim353 (Finkelstein et al., 2001) or SimLex999 (Hill et al., 2015) compared to uni-modal representations (Kádár et al.,

¹Gloss: Three men and two women with a South-East Asian appearance eat out of bowls at a black table, on which there are, among other things, paper cups and a bag; in the background there are other people and tables.

2015; Bruni et al., 2014; Kiela and Bottou, 2014).

Grounded representations of sentences that are learned from image–caption data sets also improve performance on a number of sentence-level tasks (Kiela et al., 2017; Yoo et al., 2017) when used as additional features to skip-thought vectors (Kiros et al., 2015). The model architectures used for these studies have the same overall structure as our model and coincide with image–sentence retrieval systems (Kiros et al., 2014; Karpathy and Fei-Fei, 2015): a pre-trained CNN is fixed or fine-tuned as image feature extractor, followed by a learned transformation, while sentence representations are learned by a randomly initialized recurrent neural network. These models are trained to push the true image–caption pairs closer together, and the false image–caption pairs further from each other, in a joint embedding space.

In addition to learning grounded representations for image-sentence ranking, joint vision and language systems have been proposed to solve a wide range of tasks across modalities such as image captioning (Mao et al., 2014; Vinyals et al., 2015; Xu et al., 2015), visual question answering (Antol et al., 2015; Fukui et al., 2016; Jabri et al., 2016), text-to-image synthesis (Reed et al., 2016) and multimodal machine translation (Libovicky and Helcl, 2017; Elliott and Kádár, 2017).

Our work is also closely related to multilingual joint representation learning. In this scenario, a single model is trained to solve a task across multiple languages. Ammar et al. (2016) train a multilingual dependency parser on the Universal Dependencies treebank (Nivre et al., 2015) and show that on average the single multilingual model outperforms the monolingual baselines. Johnson et al. (2016) present a zero-shot neural machine translation model that is jointly trained on language pairs $A \leftrightarrow B$ and $B \leftrightarrow C$ and show that the model is capable of performing well on the unseen language pair $A \leftrightarrow C$. Lee et al. (2017) find that jointly training a many-languages-to-one translation model on unsegmented character sequences improves BLEU scores compared to monolingual training. They also show evidence that the model can handle intra-sentence code-switching. Peters et al. (2017) train a multilingual sequence-to-sequence translation architecture on grapheme-to-phoneme conversion using more than 300 languages. They report better performance when adding multiple languages, even those which are not present in the test data. Finally,

Require: p : task switching probability.

\mathcal{D}_{c2i} : datasets $D_1 \dots D_k$ of image-caption pairs
 $\langle c, i \rangle$ for all k languages.

\mathcal{D}_{c2c} : data set of all possible caption pairs
 $\langle c_a, c_b \rangle$ for all k languages.

$\phi(c, \theta_\phi)$: caption encoder

$\psi(i, \theta_\psi)$: image encoder

while not stopping criterion **do**

$T \sim \text{Bern}(p)$

if $T = 1$ **then**

$D_n \sim \mathcal{D}_{c2i}$

$\langle c, i \rangle \sim D_n$

$\mathbf{a} \leftarrow \phi(c, \theta_\phi)$

$\mathbf{b} \leftarrow \psi(i, \theta_\psi)$

else

$\langle c_a, c_b \rangle \sim \mathcal{D}_{c2c}$

$\mathbf{a} \leftarrow \phi(c_a, \theta_\phi)$

$\mathbf{b} \leftarrow \phi(c_b, \theta_\phi)$

end if

$[\theta_\phi; \theta_\psi] \leftarrow \text{SGD}(\nabla_{[\theta_\phi; \theta_\psi]} \mathcal{J}(\mathbf{a}, \mathbf{b}))$

end while

Figure 2: Pseudo-code of the training procedure used to train our multilingual multi-task model.

massively multilingual language representations trained on over 900 languages have been shown to resemble language families (Östling and Tiedemann, 2016) and can successfully predict linguistic typology features (Malaviya et al., 2017).

In the vision and language domain, multilingual-multimodal sentence representation learning has been limited so far to two languages. The joint training of models on English and German data has been shown to outperform monolingual baselines on image-sentence ranking and semantic textual similarity tasks (Gella et al., 2017; Calixto et al., 2017). Recently Harwath et al. (2018) also showed the benefit of joint bilingual training in the domain of speech-to-image and image-to-speech retrieval using English and Hindi data.

3 Multilingual grounded learning

We train a standard model of grounded language learning which projects images and their textual descriptions into the same space (Kiros et al., 2014; Karpathy and Fei-Fei, 2015). The training procedure is illustrated by the pseudo-code in Figure 2. Images i are encoded by a fixed pre-trained CNN followed by a learned affine transformation

$\psi(i, \theta_\psi)$, and captions c are encoded by a randomly initialized RNN $\phi(c, \theta_\phi)$. The model learns to minimize the distance between pairs $\langle a, b \rangle$ using a max-of-hinges ranking loss (Faghri et al., 2017):

$$\mathcal{J}(a, b) = \max_{\langle \hat{a}, \hat{b} \rangle} [\max(0, \alpha - s(a, b) + s(\hat{a}, \hat{b}))] + \max_{\langle a, \hat{b} \rangle} [\max(0, \alpha - s(a, b) + s(a, \hat{b}))]$$

where $\langle a, b \rangle$ are the true pairs, and $\langle a, \hat{b} \rangle$ and $\langle \hat{a}, \hat{b} \rangle$ are all possible contrastive pairs in the mini-batch. The pairs either consists of image-caption pairs $\langle i, c \rangle$, where the model solves a caption-image **c2i** ranking task, or pairs of captions in multiple languages belonging to the same image $\langle c_a, c_b \rangle$, where the model solves a caption-caption **c2c** ranking task (Gella et al., 2017). Our monolingual models are trained to minimize the caption-image ranking objective **c2i** on the training set. The multilingual models are trained to minimize the ranking loss for the set of all languages \mathcal{L} in the collection: at each iteration the model is either updated for the **c2i** objective or the caption-caption **c2c** objective given either $\langle c^l, i \rangle$ or a $\langle c_a^k, c_b^m \rangle$ pair in languages $l, k, m, \dots \in \mathcal{L}$. All models are trained by first selecting a task, either **c2i** or **c2c**. In the **c2i** case, a language is sampled at random followed by sampling a random batch; in the **c2c** case, all possible $\langle c_a, c_b \rangle$ pairs across all languages are treated as a single data set. All of the model parameters are shared across all tasks and languages.

Implementation. We build our model on the PyTorch implementation of the VSE++ model (Faghri et al., 2017). Images are represented by the 2048D average-pool features extracted from the ResNet50 architecture (He et al., 2016) trained on ImageNet (Deng et al., 2009); this is followed by a trained linear layer $\mathbf{W}_I \in \mathbb{R}^{2048 \times 1024}$. Other implementation details follow (Faghri et al., 2017): sentences are represented as the final hidden state of a GRU (Chung et al., 2014) with 1024 units and 300 dimensional word-embeddings trained from scratch. We use a single word embedding matrix containing the union of all words in all considered languages. The similarity function s in the ranking loss is cosine similarity. We ℓ_2 normalize both the caption and image representations. The model is trained with the Adam optimizer (Kingma and Ba, 2014) using default parameters and learning-rate of $2e-4$. We train the model with an early stopping crite-

	En	De	Fr	Cz
En	1.0	0.04	0.06	0.02
De	-	1.0	0.03	0.01
Fr	-	-	1.0	0.01
Cz	-	-	-	1.0

Table 1: Vocabulary overlap as measured by the Jaccard coefficient between the different languages on the translation portion of the Multi30K dataset.

rion, which is to maximise the sum of the image-sentence recall scores $R@1$, $R@5$, $R@10$ on the validation set with patience of 10 evaluations. In the monolingual setting the stopping criterion is evaluated at the end of each epoch, whereas in the multilingual setup it is evaluated every 500 iterations. The probability of switching between the **c2i** and **c2c** tasks is set to 0.5. Batches from all data sets are sampled by shuffling the full dataset, going through each batch and re-shuffling when exhausted. The sentence-pair dataset used to train the **c2c** ranking model for ℓ languages is generated as follows. For a given image i , a set of languages $1 \dots \ell$, and a set of captions C_1^i, \dots, C_ℓ^i associated with an image i , we generate the set of all possible combinations of size 2 from caption sets C^i and add the Cartesian product between all resulting pairs $C_m^i \times C_n^i$ in C^i to the training set.

4 Experimental setup²

Datasets. We train and evaluate our models on the *translation* and *comparable* portions of the Multi30K dataset (Elliott et al., 2016, 2017). The translation portion (a low-resource dataset) contains 29K images, each described in one English caption with German, French, and Czech translations. The comparable portion (a higher-resource dataset) contains the same 29K images paired with five English and five German descriptions collected independently. Figure 1 presents an example of the translation and comparable portions of the data. We used the preprocessed version of the dataset, in which the text is lowercased, punctuation is normalized, and the text is tokenized³. To reduce the vocabulary size of the joint models, we replace all words occurring fewer than four times with a spe-

²Code to reproduce our results is available at <https://github.com/kadarakos/muliserla>.

³<https://github.com/multi30k/dataset>

cial “UNK” symbol. Table 1 shows the overlap between the vocabularies of the *translation* portion of the Multi30K dataset. The total number of tokens across all four languages is 17,571, and taking the union of the tokens in these four languages results in vocabulary of 16,553 tokens – a 6% reduction in vocabulary size. On the *comparable* portion of the dataset, the total vocabulary between English and German contains 18,337 tokens, with a union of 17,667, which is a 4% reduction in vocabulary size.

Evaluation. We evaluate our models on the 1K images of the 2016 test set of Multi30K either using the 5K captions from the comparable data or the 1K translation pairs. We evaluate on image-to-text ($I \rightarrow T$) and text-to-image ($T \rightarrow I$) retrieval tasks. For most experiments we report Recall at 1 (R@1), 5 (R@5) and 10 (R@10) scores averaged over 10 randomly initialised models. However, in Section 6 we only report R@10 due to space limitations and because it has less variance than R@1 or R@5.

5 Bilingual Experiments

5.1 Reproducing Gella et al. (2017)

We start by attempting to reproduce the findings of Gella et al. (2017). In these experiments we train our multi-task learning model on the *comparable* portion of Multi30K. Our models reimplement their setups used for VSE (Monolingual) and bilingual models `Pivot-Sym` (Bilingual) and `Parallel-Sym` (Bilingual + c2c). The `OE`, `Pivot-Asym` and `Parallel-Asym` models are trained using the asymmetric similarity measure introduced for the order-embeddings (Vendrov et al., 2015). The main differences between our models and Gella et al. (2017) is that they use VGG-19 image features, whereas we use ResNet50 features, and we use the max-of-hinges loss instead of the more common sum-of-hinges loss.

Table 2 shows the results on the English comparable 2016 test set. Overall our scores are higher than Gella et al. (2017), which is most likely due to the different image features (Faghri et al. (2017) also report a large performance gain when they use the ResNet instead of the VGG image features). Nevertheless, our results show a similar trend to the symmetric cosine similarity models from Gella et al. (2017): our best results are achieved with bilingual joint training with the added c2c objective. Their models trained with an asymmetric similarity

measure show a different trend: the monolingual model is stronger than the bilingual model, and the c2c loss provides no clear improvement.

Table 3 presents the German results. Once again, our implementation outperforms Gella et al. (2017), and this is likely due to the different visual features and max-of-hinges loss. However, our Bilingual model with the additional c2c objective performs the best for German, whereas Gella et al. (2017) reports the overall best results for the monolingual baseline VSE. Their models that use the asymmetric similarity function are clearly better than the Monolingual OE model. In general, the results from Gella et al. (2017) indicate the benefits of bilingual joint training, however, they do not find a clear pattern between the model configurations across languages. In our implementation, we only focused on the symmetric cosine similarity function and found a systematic pattern across both languages: bilingual training improves results on all performance metrics for both languages, and the additional c2c objective always provides further improvements.

5.2 Translations vs. independent captions

We now study whether the model can be trained on either translation pairs or independently collected bilingual captions. Gella et al. (2017) only conducted experiments on independently collected captions. However, it is known that humans have equally strong preference for translated or independently collected captions of images (Frank et al., 2018), which has implications for the difficulty and cost of collecting training data. Our baseline is a Monolingual model trained on 29K single-captioned images in the *translation* portion of Multi30K. The Bi-translation model is trained on both German and English, with shared parameters. Table 4 shows that there is a substantial improvement in performance for both languages in the bilingual setting. However, the additional c2c loss degrades performance here. This could be because we only have one caption per image in each language and it is easier to find a relationship between these views of the translation pairs.

In the Bi-comparable setting, we randomly select an English and a German sentence for each image in the *comparable* portion of Multi30K. We only find a minor difference in performance between the Bi-translation and Bi-comparable models for English, but the German results are improved. Cru-

		I→T			T→I		
		R@1	R@5	R@10	R@1	R@5	R@10
Symmetric	VSE	31.6	60.4	72.7	23.3	53.6	65.8
	Pivot-Sym	31.6	61.2	73.8	23.5	53.4	65.8
	Parallel-Sym	31.7	62.4	74.1	24.7	53.9	65.7
Asymmetric	OE	34.8	63.7	74.8	25.8	56.5	67.8
	Pivot-Asym	33.8	62.8	75.2	26.2	56.4	68.4
	Parallel-Asym	31.5	61.4	74.7	27.1	56.2	66.9
Monolingual		42.4	69.9	79.8	30.5	57.8	67.9
Bilingual		42.7	70.7	80.1	30.6	58.1	68.3
+ c2c		43.8	71.8	81.4	32.3	59.9	70.2

Table 2: English Image-to-text (I→T) and text-to-image (T→I) retrieval results on the *comparable* part of Multi30K, measured by Recall at 1, 5 at 10. Typewriter font shows performance of two sets of symmetric and asymmetric models from Gella et al. (2017).

		I→T			T→I		
		R@1	R@5	R@10	R@1	R@5	R@10
Symmetric	VSE	29.3	58.1	71.8	20.3	47.2	60.1
	Pivot-Sym	26.9	56.6	70.0	20.3	46.4	59.2
	Parallel-Sym	28.2	57.7	71.3	20.9	46.9	59.3
Asymmetric	OE	26.8	57.5	70.9	21.0	48.5	60.4
	Pivot-Asym	28.2	61.9	73.4	22.5	49.3	61.7
	Parallel-Asym	30.2	60.4	72.8	21.8	50.5	62.3
Monolingual		34.2	63.0	74.0	23.9	49.5	60.5
Bilingual		35.2	64.3	75.3	24.6	50.8	62.0
+ c2c		37.9	66.1	76.8	26.6	53.0	64.0

Table 3: German Image-to-text (I→T) and text-to-image (T→I) retrieval results on the *comparable* part of Multi30K, measured by Recall at 1, 5 at 10. Typewriter font shows performance of two sets of symmetric and asymmetric models from Gella et al. (2017).

cially, it is still better than training on monolingual data. In the Bi-comparable setting, the c2c loss does not have a detrimental effect on model performance, unlike in the Bi-translation experiment. Overall we find that the *comparable* data leads to larger improvements in retrieval performance.

5.3 Overlapping vs. non-overlapping images

In a bilingual setting, we can improve an image-sentence ranking model by collecting more data in

a second language. This can be achieved in two ways: by collecting captions in a new language for the same overlapping set of images, or by using a disjoint set of images and captions in a new language. We compare these two settings here.

In the Bi-overlap condition, we collect captions for the existing images in a new language, i.e. we use all of the English and German captions paired with a random selection of 50% of the images in *comparable* Multi30K. This results in a training

	English		German	
	I→T	T→I	I→T	T→I
Monolingual	56.3	40.1	39.5	20.9
Bi-translation	67.4	55.1	58.3	44.6
+ c2c	58.2	47.7	51.0	39.6
Bi-comparable	67.9	55.7	62.0	48.1
+ c2c	67.6	56.0	61.9	49.1

Table 4: R@10 retrieval results on the *comparable* part of Multi30K. Bi-translation is trained on 29K *translation pair* data; bi-comparable is trained by downsampling the *comparable* data to 29K.

dataset of 14.5K images with 145K bilingual captions. In the Bi-disjoint condition, we collect captions for new images in a new language, i.e. we use all of the English captions from a random selection of 50% of the images, and all of the German captions for the remaining 50% of the images. This results in a training dataset on 29K images with a total of 145K bilingual captions.

Table 5 shows the results of this experiment. The upper-bound is to train a Monolingual model on the full *comparable* corpus. For the lower bound, we train Half Monolingual models by randomly sampling half of the 29K images and their associated captions, giving 72.5K captions over 14.5K images. Unsurprisingly, the Half Monolingual models perform worse than the Full Monolingual models. In the Bi-overlap experiment, the German model is improved by collecting captions for the existing images in English. There is no difference in the performance of the English model, echoing the results from Section 5.1. The Bi-overlap model also benefits from the added c2c objective. Finally, the Bi-disjoint model performs as well as the Bi-overlap model without the c2c objective. (It was not possible to train the Bi-disjoint model with the additional c2c objective because there are no caption pairs for the same image.)

Overall, these results suggest that it is best to collect additional captions in the original language, but when adding a second language, it is better to collect extra captions for existing images and exploit the additional c2c ranking objective.

	English		German	
	I→T	T→I	I→T	T→I
Full Monolingual	79.8	67.9	74.0	60.5
Half Monolingual	73.7	61.6	66.4	53.9
Bi-overlap	73.6	62.2	67.6	54.9
+ c2c	76.0	65.9	71.2	59.1
Bi-disjoint	73.1	62.1	67.9	54.9

Table 5: R@10 retrieval results on the *comparable* part of Multi30K. Full model trained on the 29K images of the *comparable* part, Half model on 14.5K images using random downsampling. For Bi-overlap, both English and German captions are used for 14.5K images. For Bi-disjoint, 14.5K images are used for English and the remaining 14.5K images for German.

6 Multilingual experiments

We now turn our attention to multilingual learning using the English, German, French and Czech annotations in the *translation* portion of Multi30K. We only report the text-to-image (T→I) R@10 results due to space limitations.

We did not repeat the overlapping vs. non-overlapping experiments from Section 5.3 in a multilingual setting because this would introduce too much data sparsity. In order to conduct this experiment, we would have to downsample the already low-resource French and Czech captions by 50%, or even further for multi-way experiments.

6.1 Translation vs. independent captions

Table 6 shows the results of repeating the translations vs. comparable captions experiment from Section 5.2 with data in four languages. The Multi-translation models are trained on 29K images paired with a single caption in each language. These models perform better than their Monolingual counterparts, and the German, French, and Czech models are further improved with the c2c objective. The Multi-comparable models are trained by randomly sampling one English and one German caption from the *comparable* dataset, alongside the French and Czech translation pairs. These models perform as well as the Multi-translation models, and the c2c objective brings further improvements for all languages in this setting.

	En	De	Fr	Cz
Monolingual	50.4	39.5	47.0	42.0
Multi-translation	58.7	51.2	57.0	51.0
+ c2c	56.3	52.2	55.0	51.6
Multi-comparable	59.2	49.6	57.2	50.8
+ c2c	61.8	52.7	59.2	55.2

Table 6: The Monolingual and joint Multi-translation models trained on *translation pairs*, and the Multi-comparable trained on the downsampled *comparable* set with one caption per image.

These results clearly demonstrate the advantage of jointly training on more than two languages. Text-to-image retrieval performance increases by more than 11 R@10 points for each of the four languages in our experiment.

6.2 High-to-low resource transfer

We now examine whether the lower-resource French and Czech models benefit from training with the full complement of the higher-resource English and German comparable data. Therefore we train a joint model on the *translation* as well as *comparable* portions of Multi30K, and examine the performance on French and Czech.

Table 7 shows the results of this experiment. We find that the French and Czech models improve by 8.8 and 5.5 R@10 points respectively when they are only trained on the multilingual translation pairs (compared to the monolingual version), and by another 2.2 and 2.8 points if trained on the extra 155K English and German *comparable* descriptions. We also find that the additional c2c objective improves the Czech model by a further 4.8 R@10 points (this improvement is likely caused by training the model on 46 possible caption pairs). Our results show the impact of jointly training with the larger English and German resources, which demonstrates the benefits of high-to-low resource transfer.

6.3 Bilingual vs. multilingual

Finally, we investigate how useful it is to train on four languages instead of two. Figure 3 presents the image-to-text and text-to-image retrieval results of training Monolingual, Bilingual, or Multilingual models. The Monolingual and Bilingual models are trained on a random single-caption-image subsam-

	French	Czech
Monolingual	47.0	42.0
Multilingual	56.3	51.3
+ Comparable	58.9	52.4
+ c2c	61.6	57.2

Table 7: Multilingual is trained on all *translation pairs*, + Comparable adds the *comparable* data set.

ple of the *comparable* dataset with the additional c2c objective, as this configuration provided the overall best results in Sections 5.2 and 6.1. The Multilingual models are trained with the additional French and Czech *translation* data. As can be seen in Figure 3, the performance on both tasks and for both languages improves as we move from using data from one to two to four languages.

7 Conclusions

We learn multilingual multimodal sentence embeddings and show that multilingual joint training improves over bilingual joint training. We also demonstrate that low-resource languages can benefit from the additional data found in high-resource languages. Our experiments suggest that either translation pairs or independently-collected captions improve the performance of a multilingual model, and that the latter data setting provides further improvements through a caption-caption ranking objective. We also show that when collecting data in an additional language, it is better to collect captions for the existing images because we can exploit the caption-caption objective. Our results lead to several directions for future work. We would like to pin down the mechanism via which multilingual training contributes to improved performance for image-sentence ranking. Additionally, we only consider four languages and show the gain of multilingual over bilingual training only for the English-German language pair. In future work we will incorporate more languages from data sets such as the Chinese Flickr8K (Li et al., 2016) or Japanese COCO (Miyazaki and Shimizu, 2016).

Acknowledgements

Desmond Elliott was supported by an Amazon Research Award.




Figure 3: Comparing models from the Monolingual, Bilingual and Multilingual settings. The Monolingual and Bilingual models are trained on the downsampled English and German *comparable* sets with additional c2c objective. The Multilingual model uses the French and Czech *translation pairs* as additional data. The results are reported on the full 2016 test set of the *comparable* portion of Multi30K.

References

- Agirre, E., Banea, C., Cardie, C., Cer, D., Diab, M., Gonzalez-Agirre, A., Guo, W., Lopez-Gazpio, I., Maritxalar, M., Mihalcea, R., et al. (2015). Semeval-2015 task 2: Semantic textual similarity, english, spanish and pilot on interpretability. In *Proceedings of the 9th international workshop on semantic evaluation (SemEval 2015)*, pages 252–263.
- Agirre, E., Banea, C., Cardie, C., Cer, D., Diab, M., Gonzalez-Agirre, A., Guo, W., Mihalcea, R., Rigau, G., and Wiebe, J. (2014). Semeval-2014 task 10: Multilingual semantic textual similarity. In *Proceedings of the 8th international workshop on semantic evaluation (SemEval 2014)*, pages 81–91.
- Ammar, W., Mulcaire, G., Ballesteros, M., Dyer, C., and Smith, N. A. (2016). Many languages, one parser. *arXiv preprint arXiv:1602.01595*.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., and Parikh, D. (2015). Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433.
- Barman, U., Das, A., Wagner, J., and Foster, J. (2014). Code mixing: A challenge for language identification in the language of social media. In *Proceedings of the first workshop on computational approaches to code switching*, pages 13–23.
- Baroni, M. (2016). Grounding distributional semantics in the visual world. *Language and Linguistics Compass*, 10(1):3–13.
- Barsalou, L. W., Simmons, W. K., Barbey, A. K., and Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in cognitive sciences*, 7(2):84–91.
- Bruni, E., Tran, N.-K., and Baroni, M. (2014). Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49:1–47.
- Calixto, I., Liu, Q., and Campbell, N. (2017). Multilingual multi-modal embeddings for natural language processing. *arXiv preprint arXiv:1702.01101*.
- Chung, J., Cho, K., and Bengio, Y. (2016). A character-level decoder without explicit segmentation for neural machine translation. *arXiv preprint arXiv:1603.06147*.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE.
- Dolan, B., Quirk, C., and Brockett, C. (2004). Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources. In *Proceedings of the 20th international conference on Computational Linguistics*, page 350. Association for Computational Linguistics.
- Elliott, D., Frank, S., Barrault, L., Bougares, F., and Specia, L. (2017). Findings of the second shared

- task on multimodal machine translation and multilingual image description. In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*.
- Elliott, D., Frank, S., Sima'an, K., and Specia, L. (2016). Multi30k: Multilingual english-german image descriptions. *arXiv preprint arXiv:1605.00459*.
- Elliott, D. and Kádár, A. (2017). Imagination improves multimodal translation. *arXiv preprint arXiv:1705.04350*.
- Faghri, F., Fleet, D. J., Kiros, J. R., and Fidler, S. (2017). Vse++: Improved visual-semantic embeddings. *arXiv preprint arXiv:1707.05612*.
- Finkelstein, L., Gabilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., and Ruppin, E. (2001). Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414. ACM.
- Frank, S., Elliott, D., and Specia, L. (2018). Assessing multilingual multimodal image description: Studies of native speaker preferences and translator choices. *Natural Language Engineering*, 24(3):393–413.
- Fukui, A., Park, D. H., Yang, D., Rohrbach, A., Darrell, T., and Rohrbach, M. (2016). Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*.
- Funaki, R. and Nakayama, H. (2015). Image-mediated learning for zero-shot cross-lingual document retrieval. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 585–590.
- Gella, S., Sennrich, R., Keller, F., and Lapata, M. (2017). Image pivoting for learning multilingual multimodal representations. *arXiv preprint arXiv:1707.07601*.
- Harwath, D., Chuang, G., and Glass, J. (2018). Vision as an interlingua: Learning multilingual semantic embeddings of untranscribed speech. *arXiv preprint arXiv:1804.03052*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Hill, F., Reichart, R., and Korhonen, A. (2015). Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- Jabri, A., Joulin, A., and van der Maaten, L. (2016). Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer.
- Johnson, M., Schuster, M., Le, Q. V., Krikun, M., Wu, Y., Chen, Z., Thorat, N., Viégas, F., Wattenberg, M., Corrado, G., et al. (2016). Google’s multilingual neural machine translation system: enabling zero-shot translation. *arXiv preprint arXiv:1611.04558*.
- Kádár, Á., Alishahi, A., and Chrupała, G. (2015). Learning word meanings from images of natural scenes. *Traitement Automatique des Langues*, 55(3).
- Karpathy, A. and Fei-Fei, L. (2015). Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137.
- Kiela, D. and Bottou, L. (2014). Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 36–45.
- Kiela, D., Conneau, A., Jabri, A., and Nickel, M. (2017). Learning visually grounded sentence representations. *arXiv preprint arXiv:1707.06320*.
- Kiela, D., Hill, F., Korhonen, A., and Clark, S. (2014). Improving multi-modal representations using image dispersion: Why less is sometimes more. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 835–841.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kiros, R., Salakhutdinov, R., and Zemel, R. S. (2014). Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*.
- Kiros, R., Zhu, Y., Salakhutdinov, R. R., Zemel, R., Urtasun, R., Torralba, A., and Fidler, S. (2015). Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.
- Lee, J., Cho, K., and Hofmann, T. (2017). Fully character-level neural machine translation without explicit segmentation. *Transactions of the Association for Computational Linguistics*, 5:365–378.
- Li, X., Lan, W., Dong, J., and Liu, H. (2016). Adding chinese captions to images. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval*, pages 271–275. ACM.
- Libovicky, J. and Helcl, J. (2017). Attention strategies for multi-source sequence-to-sequence learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 196–202, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Malaviya, C., Neubig, G., and Littell, P. (2017). Learning language representations for typology prediction. *arXiv preprint arXiv:1707.09569*.

- Mao, J., Xu, W., Yang, Y., Wang, J., Huang, Z., and Yuille, A. (2014). Deep captioning with multimodal recurrent neural networks (m-rnn). *arXiv preprint arXiv:1412.6632*.
- Marelli, M., Menini, S., Baroni, M., Bentivogli, L., Bernardi, R., Zamparelli, R., et al. (2014). A sick cure for the evaluation of compositional distributional semantic models. In *LREC*, pages 216–223.
- Miyazaki, T. and Shimizu, N. (2016). Cross-lingual image caption generation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1780–1790. Association for Computational Linguistics.
- Nakayama, H. and Nishida, N. (2017). Zero-resource machine translation by multimodal encoder–decoder network with multimedia pivot. *Machine Translation*, 31(1-2):49–64.
- Nivre, J., Agić, Ž., Aranzabe, M. J., Asahara, M., Atutxa, A., Ballesteros, M., Bauer, J., Bengoetxea, K., Bhat, R. A., Bosco, C., et al. (2015). Universal dependencies 1.2.
- Östling, R. and Tiedemann, J. (2016). Continuous multilinguality with language vectors. *arXiv preprint arXiv:1612.07486*.
- Peters, B., Dehdari, J., and van Genabith, J. (2017). Massively multilingual neural grapheme-to-phoneme conversion. In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, pages 19–26.
- Rajendran, J., Khapra, M. M., Chandar, S., and Ravindran, B. (2015). Bridge correlational neural networks for multilingual multimodal representation learning. *arXiv preprint arXiv:1510.03519*.
- Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., and Lee, H. (2016). Generative adversarial text to image synthesis. *arXiv preprint arXiv:1605.05396*.
- Specia, L., Frank, S., Sima'an, K., and Elliott, D. (2016). A shared task on multimodal machine translation and crosslingual image description. In *Proceedings of the First Conference on Machine Translation*, pages 543–553, Berlin, Germany. Association for Computational Linguistics.
- Vendrov, I., Kiros, R., Fidler, S., and Urtasun, R. (2015). Order-embeddings of images and language. *arXiv preprint arXiv:1511.06361*.
- Vinyals, O., Toshev, A., Bengio, S., and Erhan, D. (2015). Show and tell: A neural image caption generator. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*, pages 3156–3164. IEEE.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., and Bengio, Y. (2015). Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, pages 2048–2057.
- Yoo, K. M., Shin, Y., and Lee, S.-g. (2017). Improving visually grounded sentence representations with self-attention. *arXiv preprint arXiv:1712.00609*.
- Zoph, B., Yuret, D., May, J., and Knight, K. (2016). Transfer learning for low-resource neural machine translation. *arXiv preprint arXiv:1604.02201*.