

Unsupervised Bilingual Lexicon Induction from Mono-Lingual Multimodal Data

Shizhe Chen,¹ Qin Jin,^{1*} Alexander Hauptmann²

¹School of Information, Renmin University of China, Beijing, China

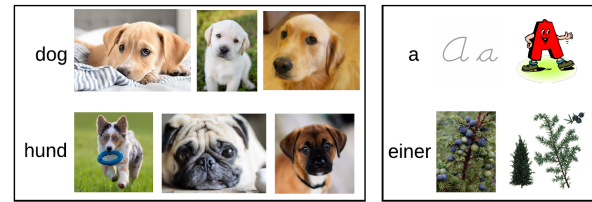
²Language Technology Institute, Carnegie Mellon University, Pittsburgh, USA
{cszhe1, qjin}@ruc.edu.cn, alex@cs.cmu.edu

Abstract

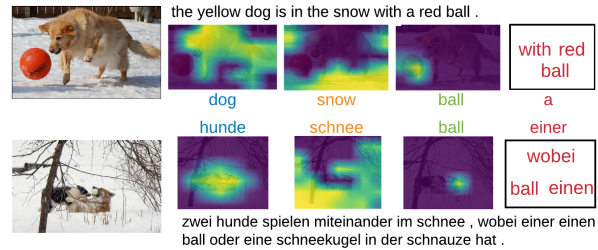
Bilingual lexicon induction, translating words from the source language to the target language, is a long-standing natural language processing task. Recent endeavors prove that it is promising to employ images as pivot to learn the lexicon induction without reliance on parallel corpora. However, these vision-based approaches simply associate words with entire images, which are constrained to translate concrete words and require object-centered images. We humans can understand words better when they are within a sentence with context. Therefore, in this paper, we propose to utilize images and their associated captions to address the limitations of previous approaches. We propose a multi-lingual caption model trained with different mono-lingual multimodal data to map words in different languages into joint spaces. Two types of word representation are induced from the multi-lingual caption model: linguistic features and localized visual features. The linguistic feature is learned from the sentence contexts with visual semantic constraints, which is beneficial to learn translation for words that are less visual-relevant. The localized visual feature is attended to the region in the image that correlates to the word, so that it alleviates the image restriction for salient visual representation. The two types of features are complementary for word translation. Experimental results on multiple language pairs demonstrate the effectiveness of our proposed method, which substantially outperforms previous vision-based approaches without using any parallel sentences or supervision of seed word pairs.

1. Introduction

The bilingual lexicon induction task aims to automatically build word translation dictionaries across different languages, which is beneficial for various natural language processing tasks such as cross-lingual information retrieval (Lavrenko, Choquette, and Croft 2002), multi-lingual sentiment analysis (Denecke 2008), machine translation (Och and Ney 2003) and so on. Although building bilingual lexicon has achieved success with parallel sentences in resource-rich languages (Och and Ney 2003), the parallel data is insufficient or even unavailable especially for resource-scarce languages and it is expensive to collect. On the contrary,



(a) The previous vision-based approach: a word is represented by global features extracted from retrieved images. It requires object-centered images and is unreliable for non-concrete words.



(b) Our proposed approach: the word representation is learned from both sentence contexts and visual localization.

Figure 1: Comparison of previous vision-based approaches and our proposed approach for bilingual lexicon induction. Best viewed in color.

there are abundant multimodal mono-lingual data on the Internet, such as images and their associated tags and descriptions, which motivates researchers to induce bilingual lexicon from these non-parallel data without supervision.

There are mainly two types of mono-lingual approaches to build bilingual dictionaries in recent works. The first is purely text-based, which explores the structure similarity between different linguistic space. The most popular approach among them is to linearly map source word embedding into the target word embedding space (Mikolov, Le, and Sutskever 2013; Conneau et al. 2018). The second type utilizes vision as bridge to connect different languages (Bergsma and Van Durme 2011; Kiela, Vulic, and Clark 2015; Hewitt et al. 2018). It assumes that words correlating to similar images should share similar semantic meanings. So previous vision-based methods search images with

*Qin Jin is the corresponding author.

multi-lingual words and translate words according to similarities of visual features extracted from the corresponding images. It has been proved that the visual-grounded word representation improves the semantic quality of the words (Silberer and Lapata 2012).

However, previous vision-based methods suffer from two limitations for bilingual lexicon induction. Firstly, the accurate translation performance is confined to concrete visual-relevant words such as nouns and adjectives as shown in Figure 1(a). For words without high-quality visual groundings, previous methods would generate poor translations (Hewitt et al. 2018). Secondly, previous works extract visual features from the whole image to represent words and thus require object-centered images in order to obtain reliable visual groundings. However, common images usually contain multiple objects or scenes, and the word might only be grounded to part of the image, therefore the global visual features will be quite noisy to represent the word.

In this paper, we address the two limitations via learning from mono-lingual multimodal data with both sentence and visual context (e.g., image and caption data) to induce bilingual lexicon. Such multimodal data is also easily obtained for different languages on the Internet (Sharma et al. 2018). We propose a multi-lingual image caption model trained on multiple mono-lingual image caption data, which is able to induce two types of word representations for different languages in the joint space. The first is the linguistic feature learned from the sentence context with visual semantic constraints, so that it is able to generate more accurate translations for words that are less visual-relevant. The second is the localized visual feature which attends to the local region of the object or scene in the image for the corresponding word, so that the visual representation of words will be more salient than previous global visual features. The two representations are complementary and can be combined to induce better bilingual word translation.

We carry out experiments on multiple language pairs including German-English, French-English, and Japanese-English. The experimental results show that the proposed multi-lingual caption model not only achieves better caption performance than independent mono-lingual models for data-scarce languages, but also can induce the two types of features, linguistic and visual features, for different languages in joint spaces. Our proposed method consistently outperforms previous state-of-the-art vision-based bilingual word induction approaches on different languages. The contributions of this paper are as follows:

- To the best of our knowledge, we are the first to explore images associated with sentences for bilingual lexicon induction, which mitigates two main limitations of previous vision-based approaches.
- We propose a multi-lingual caption model to induce both the linguistic and localized visual features for multi-lingual words in joint spaces. The linguistic and visual features are complementary to enhance word translation.
- Extensive experiments on different language pairs demonstrate the effectiveness of our approach, which achieves significant improvements over the state-of-the-

art vision-based methods in all part-of-speech classes.

2. Related Work

The early works for bilingual lexicon induction require parallel data in different languages. (Och and Ney 2003) systematically investigates various word alignment methods with parallel texts to induce bilingual lexicon. However, the parallel data is scarce or even unavailable for low-resource languages. Therefore, methods with less dependency on the availability of parallel corpora are highly desired.

There are mainly two types of mono-lingual approaches for bilingual lexicon induction: text-based and vision-based methods. The text-based methods purely exploit the linguistic information to translate words. The initiative works (Fung and Yee 1998; Rapp 1999) utilize word co-occurrences in different languages as clue for word alignment. With the improvement in word representation based on deep learning, (Mikolov, Le, and Sutskever 2013) finds the structure similarity of the deep-learned word embeddings in different languages, and employs a parallel vocabulary to learn a linear mapping from the source to target word embeddings. (Xing et al. 2015) improves the translation performance via adding an orthogonality constraint to the mapping. (Zhang et al. 2017) further introduces a matching mechanism to induce bilingual lexicon with fewer seeds. However, these models require seed lexicon as the start-point to train the bilingual mapping. Recently, (Conneau et al. 2018) proposes an adversarial learning approach to learn the joint bilingual embedding space without any seed lexicon.

The vision-based methods exploit images to connect different languages, which assume that words corresponding to similar images are semantically alike. (Bergsma and Van Durme 2011) collects images with labeled words in different languages to learn word translation with image as pivot. (Kiela, Vulic, and Clark 2015) improves the visual-based word translation performance via using more powerful visual representations: the CNN-based (Krizhevsky, Sutskever, and Hinton 2012) features. The above works mainly focus on the translation of nouns and are limited in the number of collected languages. The recent work (Hewitt et al. 2018) constructs the current largest (with respect to the number of language pairs and types of part-of-speech) multi-modal word translation dataset, MMID. They show that concrete words are easiest for vision-based translation methods while others are much less accurate. In our work, we alleviate the limitations of previous vision-based methods via exploring images and their captions rather than images with unstructured tags to connect different languages.

Image captioning has received more and more research attentions. Most image caption works focus on the English caption generation (Vinyals et al. 2017; Xu et al. 2015), while there are limited works considering generating multi-lingual captions. The recent WMT workshop (Elliott et al. 2017) has proposed a subtask of multi-lingual caption generation, where different strategies such as multi-task captioning and source-to-target translation followed by captioning have been proposed to generate captions in target languages. Our work proposes a multi-lingual image caption model that

shares part of the parameters across different languages in order to benefit each other.

3. Unsupervised Bilingual Lexicon Induction

Our goal is to induce bilingual lexicon without supervision of parallel sentences or seed word pairs, purely based on the mono-lingual image caption data. In the following, we introduce the multi-lingual image caption model whose objectives for bilingual lexicon induction are two folds: 1) explicitly build multi-lingual word embeddings in the joint linguistic space; 2) implicitly extract the localized visual features for each word in the shared visual space. The former encodes linguistic information of words while the latter encodes the visual-grounded information, which are complementary for bilingual lexicon induction.

3.1 Multi-lingual Image Caption Model

Suppose we have mono-lingual image caption datasets $D_{v,x} = \{\langle v^{(i)}, x^{(i)} \rangle\}_{i=1}^M$ in the source language and $D_{v,y} = \{\langle v^{(i)}, y^{(i)} \rangle\}_{i=1}^N$ in the target language. The images v in $D_{v,x}$ and $D_{v,y}$ do not necessarily overlap, but cover overlapped object or scene classes which is the basic assumption of vision-based methods. For notation simplicity, we omit the superscript i for the data sample. Each image caption x and y is composed of word sequences $\{x_1, \dots, x_{T_x}\}$ and $\{y_1, \dots, y_{T_y}\}$ respectively, where T_* is the sentence length.

The proposed multi-lingual image caption model aims to generate sentences in different languages to describe the image content, which connects the vision and multi-lingual sentences. Figure 2 illustrates the framework of the caption model, which consists of three parts: the image encoder, word embedding module and language decoder.

The image encoder encodes the image into the shared visual space. We apply the Resnet152 (He et al. 2016) as our encoder f_e , which produces K vectors corresponding to different spatial locations in the image:

$$a = \{a_1, \dots, a_K\} = f_e(v; \theta_e) \quad (1)$$

where $a_k \in \mathbb{R}^D$. The parameter θ_e of the encoder is shared for different languages in order to encode all the images in the same visual space.

The word embedding module maps the one-hot word representation in each language into low-dimensional distributional embeddings:

$$w_t^x = W_x x_t; \quad w_t^y = W_y y_t \quad (2)$$

where $W_x \in \mathbb{R}^{D \times N_x}$ and $W_y \in \mathbb{R}^{D \times N_y}$ is the word embedding matrix for the source and target languages respectively. N_x and N_y are the vocabulary size of the two languages.

The decoder then generates word step by step conditioning on the encoded image feature and previous generated words. The probability of generating x_t in the source language is as follows:

$$P(x_t | a, x_{<t}) = f_d(h_t^x; \theta_d) = \text{softmax}(h_t^x W_x) \quad (3)$$

where h_t^x is the hidden state of the decoder at step t , which is functioned by LSTM (Hochreiter and Schmidhuber 1997):

$$h_t^x = \text{LSTM}([w_{t-1}^x; c_t^x], h_{t-1}^x) \quad (4)$$

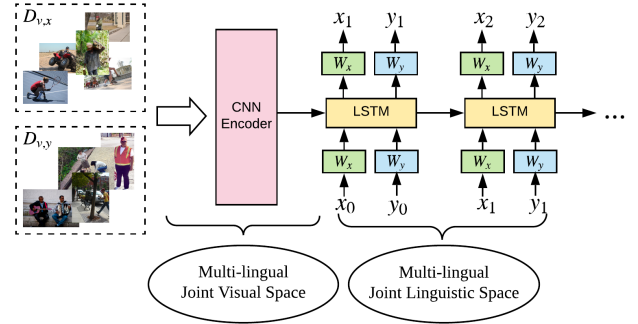


Figure 2: Multi-lingual image caption model. The source and target language caption models share the same image encoder and language decoder, which enforce the word embeddings of different languages to project in the same space.

The c_t^x is the dynamically located contextual image feature to generate word x_t via attention mechanism, which is the weighted sum of a computed by

$$c_t^x = \sum_{k=1}^K \alpha_{t,k} a_k \quad (5)$$

$$\alpha_{t,k} = \frac{\exp(f_a(h_{t-1}^x, a_k))}{\sum_{k=1}^K \exp(f_a(h_{t-1}^x, a_k))} \quad (6)$$

where f_a is a fully connected neural network. The parameter θ_d in the decoder includes all the weights in the LSTM and the attention network f_a .

Similarly, $P(y_t | a, y_{<t})$ is the probability of generating y_t in the target language, which shares θ_d with the source language. By sharing the same parameters across different languages in the encoder and decoder, both the visual features and the learned word embeddings for different languages are enforced to project in a joint semantic space. To be noted, the proposed multi-lingual parameter sharing strategy is not constrained to the presented image captioning model, but can be applied in various image captioning models such as show-tell model (Vinyals et al. 2017) and so on.

We use maximum likelihood as objective function to train the multi-lingual caption model, which maximizes the log-probability of the ground-truth captions:

$$L = \sum_{t=1}^{T_x} \log P(x_t | a, x_{<t}) + \sum_{t=1}^{T_y} \log P(y_t | a, y_{<t}) \quad (7)$$

3.2 Visual-guided Word Representation

The proposed multi-lingual caption model can induce similarities of words in different languages from two aspects: the linguistic similarity and the visual similarity. In the following, we discuss the two types of similarity and then construct the source and target word representations.

The **linguistic similarity** is reflected from the learned word embeddings W_x and W_y in the multi-lingual caption model. As shown in previous works (Mikolov, Yih, and

Zweig 2013), word embeddings learned from the language contexts can capture syntactic and semantic regularities in the language. However, if the word embeddings of different languages are trained independently, they are not in the same linguistic space and we cannot compute similarities directly. In our multi-lingual caption model, since images in $D_{v,x}$ and $D_{v,y}$ share the same visual space, the features of sentence x and y belonging to similar images are bound to be close in the same space with the visual constraints. Meanwhile, the language decoder is also shared, which enforces the word embeddings across languages into the same semantic space in order to generate similar sentence features. Therefore, W_x and W_y not only encode the linguistic information of different languages but also share the embedding space which enables direct cross-lingual similarity comparison. We refer the linguistic features of source and target words x_n and y_n as l_n^x and l_n^y respectively.

For the **visual similarity**, the multi-lingual caption model locates the image region to generate each word base on the spatial attention in Eq (6), which can be used to calculate the localized visual representation of the word. However, since the attention is computed before word generation, the localization performance can be less accurate. It also cannot be generalized to image captioning models without spatial attention. Therefore, inspired by (Zeiler and Fergus 2014), where they occlude over regions of the image to observe the change of classification probabilities, we feed different parts of the image to the caption model and investigate the probability changes for each word in the sentence. Algorithm 1 presents the procedure of word localization and the grounded visual feature generation. Please note that such visual-grounding is learned unsupervisedly from the image caption data. Therefore, every word can be represented as a set of grounded visual features (the set size equals to the word occurrence number in the dataset). We refer the localized visual feature set for source word x_n as $\{i_{n,1}^x, \dots, i_{n,C_n^x}^x\}$, for target word y_n as $\{i_{n,1}^y, \dots, i_{n,C_n^y}^y\}$.

Algorithm 1 Generating localized visual features.

Input: Encoded image features $a = \{a_1, \dots, a_K\}$, sentence $x = \{x_1, \dots, x_T\}$.
Output: Localized visual features for each word
for $t \in [1, T]$ **do**
 for each $a_k \in a$ **do**
 compute $P_{t,k} = P(x_t|a_k, x_{<t})$ according to Eq (3)
 end for
 $\hat{a}_{t,k} = P_{t,k} / (\sum_k P_{t,k})$
 $\hat{c}_t = \sum_k \hat{a}_{t,k} a_k$
end for
return $\{\hat{c}_1, \dots, \hat{c}_T\}$

3.3 Word Translation Prediction

Since the word representations of the source and target language are in the same space, we could directly compute the similarities across languages. We apply l2-normalization on the word representations and measure with the cosine simi-

Table 1: Statistics of image caption datasets.

dataset	lang	#images	#captions	#words
Multi30k	English	14,500	72,500	5,492
	German	14,500	72,500	5,535
	French	14,500	14,500	3,318
COCO + STAIR	English	56,644	283,378	7,491
	Japanese	56,643	283,215	9,331

larity. For linguistic features, the similarity is measured as:

$$s_l(x_n, y_m) = l_n^x \cdot l_m^y \quad (8)$$

However, there are a set of visual features associated with one word, so the visual similarity measurement between two words is required to take two sets of visual features as input. We aggregate the visual features in a single representation and then compute cosine similarity instead of point-wise similarities among two sets:

$$s_i(x_n, y_m) = \left(\frac{1}{C_n^x} \sum_{j=1}^{C_n^x} i_{n,j}^x \right) \cdot \left(\frac{1}{C_m^y} \sum_{j=1}^{C_m^y} i_{m,j}^y \right) \quad (9)$$

The reasons for performing aggregation are two folds. Firstly, the number of visual features is proportional to the word occurrence in our approach instead of fixed numbers as in (Kiela, Vulic, and Clark 2015; Hewitt et al. 2018). So the computation cost for frequent words are much higher. Secondly, the aggregation helps to reduce noise, which is especially important for abstract words. The abstract words such as ‘event’ are more visually diverse, but the overall styles of multiple images can reflect its visual semantics.

Due to the complementary characteristics of the two features, we combine them to predict the word translation. The translated word for x_n is

$$y_* = \operatorname{argmax}_{y_m} (s_l(x_n, y_m) + s_i(x_n, y_m)) \quad (10)$$

4. Experiments

4.1 Datasets

For image captioning, we utilize the multi30k (Elliott et al. 2016), COCO (Chen et al. 2015) and STAIR (Yoshikawa, Shigeto, and Takeuchi 2017) datasets. The multi30k dataset contains 30k images and annotations under two tasks. In task 1, each image is annotated with one English description which is then translated into German and French. In task 2, the image is independently annotated with 5 descriptions in English and German respectively. For German and English languages, we utilize annotations in task 2. For the French language, we can only employ French descriptions in task 1, so the training size for French is less than the other two languages. The COCO and STAIR datasets contain the same image set but are independently annotated in English and Japanese. Since the images in the wild for different languages might not overlap, we randomly split the image set into two disjoint parts of equal size. The images in each part only contain the mono-lingual captions. We use Moses SMT

Table 2: Image captioning performance of different languages on the Multi30k dataset.

	English		German		French	
	BLEU4	METEOR	CIDEr	BLEU4	METEOR	CIDEr
mono-lingual model	23.59	20.62	50.19	16.02	18.79	44.24
multi-lingual model	23.39	20.86	51.04	16.47	18.82	44.75

Toolkit to tokenize sentences and select words occurring more than five times in our vocabulary for each language. Table 1 summarizes the statistics of caption datasets.

For bilingual lexicon induction, we use two visual datasets: BERGSMA and MMID. The BERGSMA dataset (Bergsma and Van Durme 2011) consists of 500 German-English word translation pairs. Each word is associated with no more than 20 images. The words in BERGSMA dataset are all nouns. The MMID dataset (Hewitt et al. 2018) covers a larger variety of words and languages, including 9,808 German-English pairs and 9,887 French-English pairs. The source word can be mapped to multiple target words in their dictionary. Each word is associated with no more than 100 retrieved images. Since both these image datasets do not contain Japanese language, we download the Japanese-to-English dictionary online¹. We select words in each dataset that overlap with our caption vocabulary, which results in 230 German-English pairs in BERGSMA dataset, 1,311 German-English pairs and 1,217 French-English pairs in MMID dataset, and 2,408 Japanese-English pairs.

4.2 Experimental Setup

For the multi-lingual caption model, we set the word embedding size and the hidden size of LSTM as 512. Adam algorithm is applied to optimize the model with learning rate of 0.0001 and batch size of 128. The caption model is trained up to 100 epochs and the best model is selected according to caption performance on the validation set.

We compare our approach with two baseline vision-based methods proposed in (Kiela, Vulic, and Clark 2015; Hewitt et al. 2018), which measure the similarity of two sets of global visual features for bilingual lexicon induction:

1. CNN-mean: taking the similarity score of the averaged feature of the two image sets.
2. CNN-avgmax: taking the average of the maximum similarity scores of two image sets.

We evaluate the word translation performance using MRR (mean-reciprocal rank) as follows:

$$\text{MRR} = \frac{1}{N} \sum_{n=1}^N \max_{y_m \in D(x_n)} \frac{1}{\text{rank}(x_n, y_m)} \quad (11)$$

where $D(x_n)$ is the groundtruth translated words for source word x_n , and $\text{rank}(x_n, y_m)$ denotes the rank of groundtruth word y_m in the rank list of translation candidates. We also measure the precision at K (P@K) score, which is the proportion of source words whose groundtruth translations rank in the top K words. We set K as 1, 5, 10 and 20.

¹<https://github.com/facebookresearch/MUSE#ground-truth-bilingual-dictionaries>

Table 3: Image captioning performance of different languages on the COCO and STAIR dataset.

language	model	BLEU4	METEOR	CIDEr
English	mono	32.29	26.01	102.36
	multi	31.59	26.07	102.41
Japanese	mono	39.85	31.63	98.12
	multi	39.81	31.70	97.81



Figure 3: Visual groundings learned from the caption model.

4.3 Evaluation of Multi-lingual Image Caption

We first evaluate the captioning performance of the proposed multi-lingual caption model, which serves as the foundation stone for our bilingual lexicon induction method.

We compare the proposed multi-lingual caption model with the mono-lingual model, which consists of the same model structure, but is trained separately for each language. Table 2 presents the captioning results on the multi30k dataset, where all the languages are from the Latin family. The multi-lingual caption model achieves comparable performance with mono-lingual model for data sufficient languages such as English and German, and significantly outperforms the mono-lingual model for the data-scarce language French with absolute 3.22 gains on the CIDEr metric. For languages with distinctive grammar structures such as English and Japanese, the multi-lingual model is also on par with the mono-lingual model as shown in Table 3. To be noted, the multi-lingual model contains about twice less of parameters than the independent mono-lingual models, which is more computation efficient.

We visualize the learned visual groundings from the multi-lingual caption model in Figure 3. Though there is certain mistakes such as ‘musicians’ in the bottom image, most of the words are grounded well with correct objects or scenes, and thus can obtain more salient visual features.

Table 4: Performance of German to English word translation.

		BERGSMA					MMID				
		MRR	P@1	P@5	P@10	P@20	MRR	P@1	P@5	P@10	P@20
baselines	CNN-mean	0.650	55.6	75.2	82.7	89.3	0.262	19.9	33.3	37.6	42.4
	CNN-avgmax	0.723	65.0	79.9	84.1	87.9	0.430	38.5	47.2	49.8	52.8
proposed model	linguistic	0.755	67.6	86.2	88.0	91.6	0.467	38.8	55.5	61.6	67.4
	visual	0.762	69.2	84.6	89.7	91.1	0.400	31.5	48.1	54.7	60.3
	linguistic + visual	0.819	76.6	86.9	91.6	93.5	0.529	45.2	62.2	68.2	72.3

Table 5: German-to-English word translation examples. ‘de’ is the source German word and ‘en’ is the groundtruth target English word. The ‘rank’ denotes the position of the groundtruth target word in the candidate ranking list. The ‘top3 translation’ presents the top 3 translated words of the source word by our system.

pos	de	en	rank	top3 translation
noun	telefon	phone	1	phone, pay, telephone
	gebiet	area	10	village, trees, desert
verb	sieht	looks	1	looks, stands, look
	machen	make	4	do, take, performs
adj	roten	red	1	red, orange, pink
	wenigen	few	3	many, several, few,
adv	sehr	very	1	very, large, huge
	knapp	barely	65	pretty, short, ballet
prep	für	for	1	for, as, hold
	oberhalb	above	7	below, near, within
num	einer	a, one	1	a, the, one
	zehn	ten	5	seven, eight, six

4.4 Evaluation of Bilingual Lexicon Induction

We induce the linguistic features and localized visual features from the multi-lingual caption model for word translation from the source to target languages. Table 4 presents the German-to-English word translation performance of the proposed features. In the BERGSMA dataset, the visual features achieve better translation results than the linguistic features while they are inferior to the linguistic features in the MMID dataset. This is because the vocabulary in BERGSMA dataset mainly consists of nouns, but the parts-of-speech is more diverse in the MMID dataset. The visual features contribute most to translate concrete noun words, while the linguistic features are beneficial to other abstract words. The fusion of the two features performs best for word translation, which demonstrates that the two features are complementary with each other.

We also compare our approach with previous state-of-the-art vision-based methods in Table 4. Since our visual feature is the averaged representation, it is fair to compare with the CNN-mean baseline method where the only difference lies in the feature rather than similarity measurement. The localized features perform substantially better than the global

Table 6: Comparison of the image captioning models’ impact on the bilingual lexicon induction. The acronym L is for linguistic and V is for the visual feature.

feat	caption model	MRR	P@1	P@5	P@10	P@20
L	mp	0.437	35.5	52.9	58.8	66.4
	attn	0.467	38.8	55.5	61.6	67.4
V	mp	0.367	28.8	45.0	50.2	57.1
	attn	0.400	31.5	48.1	54.7	60.3
L+V	mp	0.490	40.4	58.6	65.4	71.7
	attn	0.529	45.2	62.2	68.2	72.3

image features which demonstrate the effectiveness of the attention learned from the caption model. The combination of visual and linguistic features also significantly improves the state-of-the-art visual-based CNN-avgmax method with 11.6% and 6.7% absolute gains on P@1 on the BERGSMA and MMID dataset respectively.

In Figure 4, we present the word translation performance for different POS (part-of-speech) labels. We assign the POS label for words in different languages according to their translations in English. We can see that the previous state-of-the-art vision-based approach contributes mostly to noun words which are most visual-relevant, while generates poor translations for other part-of-speech words. Our approach, however, substantially improves the translation performance for all part-of-speech classes. For concrete words such as nouns and adjectives, the localized visual features produce better representation than previous global visual features; and for other part-of-speech words, the linguistic features, which are learned with sentence context, are effective to complement the visual features. The fusion of the linguistic and localized visual features in our approach leads to significant performance improvement over the state-of-the-art baseline method for all types of POS classes.

Some correct and incorrect translation examples for different POS classes are shown in Table 5. The visual-relevant concrete words are easier to translate such as ‘phone’ and ‘red’. But our approach still generates reasonable results for abstract words such as ‘area’ and functional words such as ‘for’ due to the fusion of visual and sentence contexts.

We also evaluate the influence of different image captioning structures on the bilingual lexicon induction. We com-

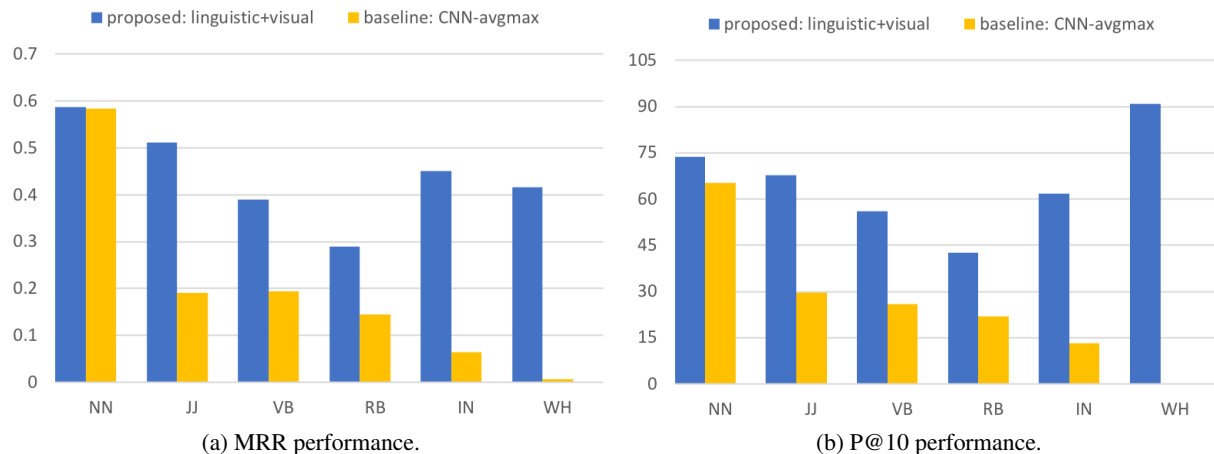


Figure 4: Performance comparison of German-to-English word translation on the MMID dataset. The word translation performance is broken down by part-of-speech labels.

Table 7: Performance of French to English word translation on the MMID dataset.

	MRR	P@1	P@5	P@10	P@20
CNN-mean	0.301	22.8	37.1	43.1	48.6
CNN-avgmax	0.474	41.9	52.6	55.4	59.1
linguistic	0.376	29.3	47.0	52.6	58.9
visual	0.387	31.1	46.7	52.1	58.9
linguistic+visual	0.494	42.0	57.1	62.9	69.4

Table 8: Performance of Japanese to English word translation.

	MRR	P@1	P@5	P@10	P@20
linguistic	0.290	22.1	35.8	42.0	49.5
visual	0.419	34.2	50.4	56.3	61.1
linguistic+visual	0.469	38.3	56.9	62.9	68.1

pare our attention model (‘attn’) with the vanilla show-tell model (‘mp’) (Vinyals et al. 2017), which applies mean pooling over spatial image features to generate captions and achieves inferior caption performance to the attention model. Table 6 shows the word translation performance of the two caption models. The attention model with better caption performance also induces better linguistic and localized visual features for bilingual lexicon induction. Nevertheless, the show-tell model still outperforms the previous vision-based methods in Table 4.

4.5 Generalization to Diverse Language Pairs

Beside German-to-English word translation, we expand our approach to other languages including French and Japanese which is more distant from English.

The French-to-English word translation performance is presented in Table 7. To be noted, the training data of the French captions is five times less than German captions,

which makes French-to-English word translation performance less competitive with German-to-English. But similarly, the fusion of linguistic and visual features achieves the best performance, which has boosted the baseline methods with 4.2% relative gains on the MRR metric and 17.4% relative improvements on the P@20 metric.

Table 8 shows the Japanese-to-English word translation performance. Since the language structures of Japanese and English are quite different, the linguistic features learned from the multi-lingual caption model are less effective but still can benefit the visual features to improve the translation quality. The results on multiple diverse language pairs further demonstrate the generalization of our approach for different languages.

5. Conclusion

In this paper, we address the problem of bilingual lexicon induction without reliance on parallel corpora. Based on the experience that we humans can understand words better when they are within the context and can learn word translations with external world (e.g. images) as pivot, we propose a new vision-based approach to induce bilingual lexicon with images and their associated sentences. We build a multi-lingual caption model from multiple mono-lingual multimodal data to map words in different languages into joint spaces. Two types of word representation, linguistic features and localized visual features, are induced from the caption model. The two types of features are complementary for word translation. Experimental results on multiple language pairs demonstrate the effectiveness of our proposed method, which leads to significant performance improvement over the state-of-the-art vision-based approaches for all types of part-of-speech. In the future, we will further expand the vision-pivot approaches for zero-resource machine translation without parallel sentences.

6. Acknowledgments

This work was supported by National Natural Science Foundation of China under Grant No. 61772535, National Key Research and Development Plan under Grant No. 2016YFB1001202 and Research Foundation of Beijing Municipal Science & Technology Commission under Grant No. Z181100008918002.

References

- Bergsma, S., and Van Durme, B. 2011. Learning bilingual lexicons using the visual similarity of labeled web images. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, 1764.
- Chen, X.; Fang, H.; Lin, T.-Y.; Vedantam, R.; Gupta, S.; Dollár, P.; and Zitnick, C. L. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Conneau, A.; Lample, G.; Ranzato, M.; Denoyer, L.; and Jégou, H. 2018. Word translation without parallel data. *ICLR*.
- Denecke, K. 2008. Using sentiwordnet for multilingual sentiment analysis. In *Data Engineering Workshop, 2008. ICDEW 2008. IEEE 24th International Conference on*, 507–512. IEEE.
- Elliott, D.; Frank, S.; Sima'an, K.; and Specia, L. 2016. Multi30k: Multilingual english-german image descriptions. *arXiv preprint arXiv:1605.00459*.
- 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. *arXiv preprint arXiv:1710.07177*.
- Fung, P., and Yee, L. Y. 1998. An ir approach for translating new words from nonparallel, comparable texts. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics-Volume 1*, 414–420. Association for Computational Linguistics.
- 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*, 770–778.
- Hewitt, J.; Ippolito, D.; Callahan, B.; Kriz, R.; Wijaya, D. T.; and Callison-Burch, C. 2018. Learning translations via images with a massively multilingual image dataset. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 1, 2566–2576.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Kiela, D.; Vulic, I.; and Clark, S. 2015. Visual bilingual lexicon induction with transferred convnet features. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing. ACL*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Lavrenko, V.; Choquette, M.; and Croft, W. B. 2002. Cross-lingual relevance models. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, 175–182. ACM.
- Mikolov, T.; Le, Q. V.; and Sutskever, I. 2013. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Mikolov, T.; Yih, W.-t.; and Zweig, G. 2013. Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 746–751.
- Och, F. J., and Ney, H. 2003. A systematic comparison of various statistical alignment models. *Computational linguistics* 29(1):19–51.
- Rapp, R. 1999. Automatic identification of word translations from unrelated english and german corpora. In *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*, 519–526. Association for Computational Linguistics.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 1, 2556–2565.
- Silberer, C., and Lapata, M. 2012. Grounded models of semantic representation. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 1423–1433. Association for Computational Linguistics.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2017. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence* 39(4):652–663.
- Xing, C.; Wang, D.; Liu, C.; and Lin, Y. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1006–1011.
- 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*, 2048–2057.
- Yoshikawa, Y.; Shigeto, Y.; and Takeuchi, A. 2017. Stair captions: Constructing a large-scale japanese image caption dataset. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, volume 2, 417–421.
- Zeiler, M. D., and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*, 818–833. Springer.
- Zhang, M.; Peng, H.; Liu, Y.; Luan, H.-B.; and Sun, M. 2017. Bilingual lexicon induction from non-parallel data with minimal supervision. In *AAAI*, 3379–3385.