

Multimodal Machine Translation with Embedding Prediction

Tosho Hirasawa and Hayahide Yamagishi and Yukio Matsumura and Mamoru Komachi

Tokyo Metropolitan University

{tosho.hirasawa@, yamagishi-hayahide@ed., matsumura-yukio@ed.,
komachi@}tmu.ac.jp

Abstract

Multimodal machine translation is an attractive application of neural machine translation (NMT). It helps computers to deeply understand visual objects and their relations with natural languages. However, multimodal NMT systems suffer from a shortage of available training data, resulting in poor performance for translating rare words. In NMT, pretrained word embeddings have been shown to improve NMT of low-resource domains, and a search-based approach is proposed to address the rare word problem. In this study, we effectively combine these two approaches in the context of multimodal NMT and explore how we can take full advantage of pretrained word embeddings to better translate rare words. We report overall performance improvements of 1.24 METEOR and 2.49 BLEU and achieve an improvement of 7.67 F-score for rare word translation.

1 Introduction

In **multimodal machine translation**, a target sentence is translated from a source sentence together with related nonlinguistic information such as visual information. Recently, neural machine translation (NMT) has superseded traditional statistical machine translation owing to the introduction of the attentional encoder-decoder model, in which machine translation is treated as a sequence-to-sequence learning problem and is trained to pay attention to the source sentence while decoding (Bahdanau et al., 2015).

Most previous studies on multimodal machine translation are classified into two categories: visual feature adaptation and data augmentation. In visual feature adaptation, multitask learning (Elliott and Kádár, 2017) and feature integration architecture (Caglayan et al., 2017a; Calixto et al., 2017) are proposed to improve neural network

models. Data augmentation aims to deal with the fact that the size of available datasets for multimodal translation is quite small. To alleviate this problem, parallel corpora without a visual source (Elliott and Kádár, 2017; Grönroos et al., 2018) and pseudo-parallel corpora obtained using back-translation (Helcl et al., 2018) are used as additional learning resources.

Due to the availability of parallel corpora for NMT, Qi et al. (2018) suggested that initializing the encoder with pretrained word embedding improves the translation performance in low-resource language pairs. Recently, Kumar and Tsvetkov (2019) proposed an NMT model that predicts the embedding of output words and searches for the output word instead of calculating the probability using the softmax function. This model performed as well as conventional NMT, and it significantly improved the translation accuracy for rare words.

In this study, we introduce an NMT model with embedding prediction for multimodal machine translation that fully uses pretrained embeddings to improve the translation accuracy for rare words.

The main contributions of this study are as follows:

1. We propose a novel multimodal machine translation model with embedding prediction and explore various settings to take full advantage of word embeddings.
2. We show that pretrained word embeddings improve the model performance, especially when translating rare words.

2 Multimodal Machine Translation with Embedding Prediction

We integrate an embedding prediction framework (Kumar and Tsvetkov, 2019) with the multimodal

machine translation model and take advantage of pretrained word embeddings. To highlight the effect of pretrained word embeddings and embedding prediction architecture, we adopt IMAGINATION (Elliott and Kádár, 2017) as a simple multi-modal baseline.

IMAGINATION jointly learns machine translation and visual latent space models. It is based on a conventional NMT model for a machine translation task. In latent space learning, a source sentence and the paired image are mapped closely in the latent space. We use the latent space learning model as it is, except for the preprocessing of images. The models for each task share the same textual encoder in a multitask scenario.

The loss function for multitask learning is the linear interpolation of loss functions for each task.

$$J = \lambda J_T(\theta, \phi_T) + (1 - \lambda) J_V(\theta, \phi_V) \quad (1)$$

where θ is the parameter of the shared encoder; ϕ_T and ϕ_V are parameters of the machine translation model and latent space model, respectively; and λ is the interpolation coefficient¹.

2.1 Neural Machine Translation with Embedding Prediction

The machine translation part in our proposed model is an extension of Bahdanau et al. (2015). However, instead of using the probability of each word in the decoder, it searches for output words based on their similarity with word embeddings. Once the model predicts a word embedding, its nearest neighbor in the pretrained word embeddings is selected as the system output.

$$\hat{e}_j = \tanh(\mathbf{W}_o s_j + \mathbf{b}_o) \quad (2)$$

$$\hat{y}_j = \operatorname{argmin}_{w \in \mathcal{V}} \{d(\hat{e}_j, \mathbf{e}(w))\} \quad (3)$$

where s_j , \hat{e}_j , and \hat{y}_j are the hidden state of the decoder, predicted embedding, and system output, respectively, for each timestep j in the decoding process. $\mathbf{e}(w)$ is the pretrained word embedding for a target word w . d is a distance function that is used to calculate the word similarity. \mathbf{W}_o and \mathbf{b}_o are parameters of the output layer.

We adopt margin-based ranking loss (Lazaridou et al., 2015) as the loss function

of the machine translation model.

$$J_T(\theta, \phi_T) = \sum_j^M \max\{0, \gamma + d(\hat{e}_j, \mathbf{e}(w_j^-)) - d(\hat{e}_j, \mathbf{e}(y_j))\} \quad (4)$$

$$w_j^- = \operatorname{argmax}_{w \in \mathcal{V}} \{d(\hat{e}_j, \mathbf{e}(w)) - d(\hat{e}_j, \mathbf{e}(y_j))\} \quad (5)$$

where M is the length of a target sentence and γ is the margin². w_j^- is a negative sample that is close to the predicted embedding and far from the gold embedding as measuring using d .

Pretrained word embeddings are also used to initialize the embedding layers of the encoder and decoder, and the output layer of the decoder. The embedding layer of the encoder is updated during training, and the embedding layer of the decoder is fixed to the initial value.

2.2 Visual Latent Space Learning

The decoder of this model calculates the average vector over the hidden states \mathbf{h}_i in the encoder and maps it to the final vector $\hat{\mathbf{v}}$ in the latent space.

$$\hat{\mathbf{v}} = \tanh(\mathbf{W}_v \cdot \frac{1}{N} \sum_i^N \mathbf{h}_i) \quad (6)$$

where N is the length of an input sentence and $\mathbf{W}_v \in \mathbb{R}^{N \times M}$ is learned parameter of the model.

We use max margin loss as the loss function; it learns to make corresponding latent vectors of a source sentence and the paired image closer.

$$J_V(\theta, \phi_V) = \sum_{v' \neq v} \max\{0, \alpha + d(\hat{\mathbf{v}}, \mathbf{v}') - d(\hat{\mathbf{v}}, \mathbf{v})\} \quad (7)$$

where \mathbf{v} is the latent vector of the paired image; \mathbf{v}' , the image vector for other examples; and α , the margin that adjusts the sparseness of each vector in the latent space³.

3 Experiment

3.1 Dataset

We train, validate, and test our model with the Multi30k (Elliott et al., 2016) dataset published in the WMT17 Shared Task.

We choose French as the source language and English as the target one. The vocabulary size of both the source and the target languages is 10,000.

²We use $\gamma = 0.5$ in the experiment.

³We use $\alpha = 0.1$ in our experiment.

¹We use $\lambda = 0.01$ in the experiment.

Following Kumar and Tsvetkov (2019), byte pair encoding (Sennrich et al., 2016) is not applied. The source and target sentences are preprocessed with lower-casing, tokenizing and normalizing the punctuation.

Visual features are extracted using pretrained ResNet (He et al., 2016). Specifically, we encode all images in Multi30k with ResNet-50 and pick out the hidden state in the pool5 layer as a 2,048-dimension visual feature. We calculate the centroid of visual features in the training dataset as the bias vector and subtract the bias vector from all visual features in the training, validation and test datasets.

3.2 Model

The model is implemented using nmtpytorch toolkit v3.0.0⁴ (Caglayan et al., 2017b).

The shared encoder has 256 hidden dimensions, and therefore the bidirectional GRU has 512 dimensions. The decoder in NMT model has 256 hidden dimension. The input word embedding size and output vector size is 300 each. The latent space vector size is 2,048.

We used the Adam optimizer with learning rate of 0.0004. The gradient norm is clipped to 1.0. The dropout rate is 0.3.

BLEU (Papineni et al., 2002) and METEOR (Denkowski and Lavie, 2014) are used as performance metrics. We also evaluated the models using the F-score of each word; this shows how accurately each word is translated into target sentences, as was proposed in Kumar and Tsvetkov (2019). The F-score is calculated as the harmonic mean of the precision (fraction of produced sentences with a word that is in the references sentences) and the recall (fraction of reference sentences with a word that is in model outputs). We ran the experiment three times with different random seeds and obtained the mean and variance for each model.

To clarify the effect of pretrained embeddings on machine translation, we also initialized the encoder and decoder of our models with random values instead of pretrained embeddings, and investigated the effect of fixing decoder embeddings.

3.3 Word Embedding

We use publicly available pretrained FastText (Bojanowski et al., 2017) embeddings (Grave et al., 2018). These word embeddings are

⁴<https://github.com/toshohirasawa/nmtpytorch-emb-pred>

Model	val	test	
	BLEU	BLEU	METEOR
NMT	50.83	51.00±.37	42.65±.12
+ pretrained	52.05	52.33±.66	43.42±.13
IMAG+	51.03	51.18±.16	42.80±.19
+ pretrained	52.40	52.75±.25	43.56±.04
Ours	53.14	53.49±.20	43.89±.14

Table 1: Results on Multi30k validation and test dataset. NMT denotes the text-only conventional NMT model (Bahdanau et al., 2015) and IMAG+ denotes our reimplementation of the IMAGINATION (Elliott and Kádár, 2017) model. “+ pretrained” models are initialized with pretrained embeddings.

trained on Wikipedia and Common Crawl using the CBOW algorithm, and the dimension is 300.

The embedding for unknown words is calculated as the average embedding over words that are a part of pretrained embeddings but are not included in the vocabularies. Both the target and the source embeddings are preprocessed according to Mu and Viswanath (2018), in which all embeddings are debiased to make the average embedding into a zero vector and the top five principal components are subtracted for each embedding.

4 Results

Table 1 shows the overall performance of the proposed and baseline models. Compared with randomly initialized models, our model outperforms the text-only baseline by +2.49 BLEU and +1.24 METEOR, and the multimodal baseline by +2.31 BLEU and +1.09 METEOR, respectively. While pretrained embeddings improve NMT/IMAGINATION models as well, the improved models are still beyond our model.

Table 2 shows the results of ablation experiments of the initialization and fine-tuning methods. The pretrained embedding models outperform other models by up to +2.77 BLEU and +1.37 METEOR.

5 Discussion

Rare Words Our model shows a great improvement for low-frequency words. Figure 1 shows a variety of F-score according to the word frequency in the training corpus. Whereas IMAGINATION improves the translation accuracy uniformly, our model shows substantial improvement

Encoder	Decoder	Fixed	BLEU	METEOR
fasttext	fasttext	Yes	53.49	43.89
random	fasttext	Yes	53.22	43.83
fasttext	random	No	51.53	43.07
random	random	No	51.42	42.77
fasttext	fasttext	No	51.42	42.88
random	fasttext	No	50.72	42.52

Table 2: Results on test dataset with variations of model initialization and fine-tuning in decoder.




Figure 1: F-score of word prediction per frequency breakdown in training corpus.

for rare words.

Word Embeddings Furthermore, we found that decoder embeddings must be fixed to improve multimodal machine translation with embedding prediction. When we allow fine-tuning on the embedding layer, the performance drops below the baseline. It seems that fine-tuning embeddings in NMT with embedding prediction makes the model search for common words more than expected, thus preventing it from predicting rare words.

More interestingly, using pretrained FastText embeddings on the decoder rather than the encoder improves performance. This finding is different from Qi et al. (2018), in which only the encoder benefits from pretrained embeddings. Compared with the model initialized with a random value, initializing the decoder with the embedding results in an increase of +1.80 BLEU; in contrast, initializing the encoder results in an increase of only +0.11 BLEU. This is caused by the multitask learning model that trains the encoder with images and takes it away from what the embedding prediction model wants to learn from the sentences.

Model	val		test
	BLEU	BLEU	METEOR
Ours	53.14	53.49	43.89
– Debias	52.65	53.27	43.91
– Images	52.97	53.25	43.91

Table 3: Ablation experiments of visual features. “– Debias” denotes the result without subtracting the bias vector. “– Images” shows the result of text-only NMT with embedding prediction.

Visual Feature We also investigated the effect of images and its preprocessing in NMT with embedding prediction (Table 3). The interesting result is that multitask learning with raw images would not help the predictive model. Debiasing images is an essential preprocessing for NMT with embedding prediction to use images effectively in multitask learning scenario.

Translation Examples In Table 4, we show French-English translations generated by different models. In the left example, our proposed model correctly translates “voûte” into “archway” (occurs five times in the training set), Although the baseline model translates it to its synonym having higher frequency (nine times for “arch” and 12 times for “monument”). At the same time, our outputs tend to be less fluent for long sentences. The right example shows that our model translates some words (“patterned” and “carpet”) more concisely; however, it generates a less fluent sentence than the baseline.

6 Related Works

Most studies on multimodal machine translation are divided into two categories: visual feature adaptation and data augmentation.

First, in visual feature adaptation, visual features are extracted using image processing techniques and then integrated into a machine translation model. In contrast, most multitask learning models use latent space learning as their auxiliary task. Elliott and Kádár (2017) proposed the IMAGINATION model that learns to construct the corresponding visual feature from the textual hidden states of a source sentence. The visual model shares its encoder with the machine translation model; this helps in improving the textual encoder.

Second, in data augmentation, parallel corpora without images are widely used as additional train-


Image		
Source	un homme en vélo pédale devant une voûte .	quatre hommes , dont trois portent des kippas , sont assis sur un tapis à motifs bleu et vert olive .
Reference	a man on a bicycle pedals through an archway .	four men , three of whom are wearing prayer caps , are sitting on a blue and olive green patterned mat .
NMT	a man on a bicycle pedal past an arch .	four men , three of whom are wearing aprons , are sitting on a blue and green speedo carpet .
IMAG+	a man on a bicycle pedals outside a monument .	four men , three of them are wearing alaska , are sitting on a blue patterned carpet and green green seating .
Ours	a man on a bicycle pedals in front of a archway .	four men , three are wearing these are wearing these are sitting on a blue and green patterned mat .

Table 4: French to English translation examples in the Multi30k test set.

ing data. Grönroos et al. (2018) trained their multimodal model with parallel corpora and achieved state-of-the-art performance in the WMT 2018. However, the use of monolingual corpora has seldom been studied in multimodal machine translation. Our study proposes using word embeddings that are pretrained on monolingual corpora.

7 Conclusion

We have proposed a multimodal machine translation model with embedding prediction and showed that pretrained word embeddings improve the performance in multimodal translation tasks, especially when translating rare words.

In the future, we will tailor the training corpora for embedding learning, especially for handling the embedding for unknown words in the context of multimodal machine translation. We will also incorporate visual features into contextualized word embeddings.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. In *TACL*, volume 5, pages 135–146.
- Ozan Caglayan, Walid Aransa, Adrien Bardet, Mercedes García-Martínez, Fethi Bougares, Loïc Barrault, Marc Masana, Luis Herranz, and Joost van de Weijer. 2017a. LIUM-CVC submissions for WMT17 multimodal translation task. In *WMT*, pages 432–439.
- Ozan Caglayan, Mercedes García-Martínez, Adrien Bardet, Walid Aransa, Fethi Bougares, and Loïc Barrault. 2017b. NMT-PY: A flexible toolkit for advanced neural machine translation systems. *Prague Bull. Math. Linguistics*, 109:15–28.
- Iacer Calixto, Qun Liu, and Nick Campbell. 2017. Doubly-attentive decoder for multi-modal neural machine translation. In *ACL*, pages 1913–1924.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *WMT*, pages 376–380.
- Desmond Elliott, Stella Frank, Khalil Sima’an, and Lucia Specia. 2016. Multi30k: Multilingual English-German image descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74.
- Desmond Elliott and Àkos Kádár. 2017. Imagination improves multimodal translation. In *IJCNLP*, volume 1, pages 130–141.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning

- word vectors for 157 languages. In *LREC*, pages 3483–3487.
- Stig-Arne Grönroos, Benoit Huet, Mikko Kurimo, Jorma Laaksonen, Bernard Merialdo, Phu Pham, Mats Sjöberg, Umut Sulubacak, Jörg Tiedemann, Raphael Troncy, and Raúl Vázquez. 2018. The MeMAD submission to the WMT18 multimodal translation task. In *WMT*, pages 603–611.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*, pages 770–778.
- Jindřich Helcl, Jindřich Libovický, and Dusan Varis. 2018. CUNI system for the WMT18 multimodal translation task. In *WMT*, pages 616–623.
- Sachin Kumar and Yulia Tsvetkov. 2019. Von Mises-Fisher loss for training sequence to sequence models with continuous outputs. In *ICLR*.
- Angeliki Lazaridou, Georgiana Dinu, and Marco Baroni. 2015. Hubness and pollution: Delving into cross-space mapping for zero-shot learning. In *ACL-IJCNLP*, pages 270–280.
- Jiaqi Mu and Pramod Viswanath. 2018. All-but-the-Top: Simple and effective postprocessing for word representations. In *ICLR*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318.
- Ye Qi, Devendra Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. When and why are pre-trained word embeddings useful for neural machine translation? In *NAACL*, pages 529–535.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *ACL*, pages 1715–1725.