

Incorporating Global Visual Features into Attention-Based Neural Machine Translation

Iacer Calixto

ADAPT Centre
Dublin City University
School of Computing
Glasnevin, Dublin 9

iacer.calixto@adaptcentre.ie

Qun Liu

ADAPT Centre
Dublin City University
School of Computing
Glasnevin, Dublin 9

qun.liu@adaptcentre.ie

Abstract

We introduce multi-modal, attention-based Neural Machine Translation (NMT) models which incorporate visual features into different parts of both the encoder and the decoder. Global image features are extracted using a pre-trained convolutional neural network and are incorporated (i) as words in the source sentence, (ii) to initialise the encoder hidden state, and (iii) as additional data to initialise the decoder hidden state. In our experiments, we evaluate translations into English and German, how different strategies to incorporate global image features compare and which ones perform best. We also study the impact that adding synthetic multi-modal, multilingual data brings and find that the additional data have a positive impact on multi-modal models. We report new state-of-the-art results and our best models also significantly improve on a comparable Phrase-Based Statistical MT (PBSMT) model trained on the Multi30k data set according to all metrics evaluated. To the best of our knowledge, it is the first time a purely neural model significantly improves over a PBSMT model on all metrics evaluated on this data set.

1 Introduction

Neural Machine Translation (NMT) has recently been proposed as an instantiation of the *sequence to sequence* (*seq2seq*) learning problem (Kalchbrenner and Blunsom, 2013; Cho et al., 2014b; Sutskever et al., 2014). In this problem, each training example consists of one source and one target variable-length sequence, with no prior information regarding the alignments between the two.

A model is trained to *translate* sequences in the source language into corresponding sequences in the target. This framework has been successfully used in many different tasks, such as handwritten text generation (Graves, 2013), image description generation (Hodosh et al., 2013; Kiros et al., 2014; Mao et al., 2014; Elliott et al., 2015; Karpathy and Fei-Fei, 2015; Vinyals et al., 2015), machine translation (Cho et al., 2014b; Sutskever et al., 2014) and video description generation (Donahue et al., 2015; Venugopalan et al., 2015).

Recently, there has been an increase in the number of natural language generation models that explicitly use *attention-based decoders*, i.e. decoders that model an *intra-sequential* mapping between source and target representations. For instance, Xu et al. (2015) proposed an attention-based model for the task of Image Description Generation (IDG) where the model learns to *attend* to specific parts of an image (the source) as it generates its description (the target). In MT, one can intuitively interpret this attention mechanism as inducing an *alignment* between source and target sentences, as first proposed by Bahdanau et al. (2015). The common idea is to explicitly frame a learning task in which the decoder learns to attend to the relevant parts of the source sequence when generating each part of the target sequence.

We are inspired by recent successes in using attention-based models in both IDG and NMT. Our main goal in this work is to propose end-to-end multi-modal NMT models which effectively incorporate visual features in different parts of the attention-based NMT framework. The main contributions of our work are:

- We propose novel attention-based multi-modal NMT models which incorporate visual features into the encoder and the decoder.
- We discuss the impact that adding synthetic

multi-modal and multilingual data brings to multi-modal NMT.

- We show that images bring useful information to an NMT model and report state-of-the-art results.

One additional contribution of our work is that we corroborate previous findings by Vinyals et al. (2015) that suggested that using image features directly as additional context to update the hidden state of the decoder (at each time step) prevents learning.

The remainder of this paper is structured as follows. In §1.1 we briefly discuss relevant previous related work. We then revise the attention-based NMT framework and further expand it into different multi-modal NMT models (§2). In §3 we introduce the data sets we use in our experiments. In §4 we detail the hyperparameters, parameter initialisation and other relevant details of our models. Finally, in §6 we draw conclusions and provide some avenues for future work.

1.1 Related work

Attention-based encoder-decoder models for MT have been actively investigated in recent years. Some researchers have studied how to improve attention mechanisms (Luong et al., 2015; Tu et al., 2016) and how to train attention-based models to translate between many languages (Dong et al., 2015; Firat et al., 2016).

However, multi-modal MT has only recently been addressed by the MT community in the form of a shared task (Specia et al., 2016). We note that in the official results of this first shared task no submissions based on a purely neural architecture could improve on the Phrase-Based SMT (PBSMT) baseline. Nevertheless, researchers have proposed to include global visual features in re-ranking n -best lists generated by a PBSMT system or directly in a purely NMT framework with some success (Caglayan et al., 2016; Calixto et al., 2016; Libovický et al., 2016; Shah et al., 2016). The best results achieved by a purely NMT model in this shared task are those of Huang et al. (2016), who proposed to use global and regional image features extracted with the VGG19 (Simonyan and Zisserman, 2014) and the RCNN (Girshick et al., 2014) convolutional neural networks (CNNs).

Similarly to one of the three models we pro-

pose,¹ Huang et al. (2016) extract global features for an image, project these features into the vector space of the source words and then add it as a word in the input sequence. Their best model improves over a strong NMT baseline and is comparable to results obtained with a PBSMT model trained on the same data, although not significantly better. For that reason, their models are used as baselines in our experiments. Next, we point out some key differences between the work of Huang et al. (2016) and ours.

Architecture Their implementation is based on the attention-based model of Luong et al. (2015), which has some differences to that of Bahdanau et al. (2015), used in our work (§2.1). Their encoder is a single-layer unidirectional LSTM and they use the last hidden state of the encoder to initialise the decoder’s hidden state, therefore indirectly using the image features to do so. We use a bi-directional recurrent neural network (RNN) with GRU (Cho et al., 2014a) as our encoder, better encoding the semantics of the source sentence.

Image features We include image features separately either as a word in the source sentence (§2.2.1) or *directly* for encoder (§2.2.2) or decoder initialisation (§2.2.3), whereas Huang et al. (2016) only use it as a word. We also show it is better to include an image exclusively for the encoder *or* the decoder initialisation (Tables 1 and 2).

Data Huang et al. (2016) use object detections obtained with the RCNN of Girshick et al. (2014) as additional data, whereas we study the impact that additional back-translated data brings.

Performance All our models outperform Huang et al. (2016)’s according to all metrics evaluated, even when they use additional object detections. If we use additional back-translated data, the difference becomes even larger.

2 Attention-based NMT

In this section, we briefly revise the attention-based NMT framework (§2.1) and expand it into a multi-modal NMT framework (§2.2).

2.1 Text-only attention-based NMT

We follow the notation of Bahdanau et al. (2015) and Firat et al. (2016) throughout this section.

¹This idea has been developed independently by both research groups.

Given a source sequence $X = (x_1, x_2, \dots, x_N)$ and its translation $Y = (y_1, y_2, \dots, y_M)$, an NMT model aims at building a single neural network that translates X into Y by directly learning to model $p(Y|X)$. Each x_i is a row index in a source lookup matrix $\mathbf{W}_x \in \mathbb{R}^{|V_x| \times d_x}$ (the *source word embeddings matrix*) and each y_j is an index in a target lookup matrix $\mathbf{W}_y \in \mathbb{R}^{|V_y| \times d_y}$ (the *target word embeddings matrix*). V_x and V_y are source and target vocabularies and d_x and d_y are source and target word embeddings dimensionalities, respectively.

A bidirectional RNN with GRU is used as the encoder. A forward RNN $\overrightarrow{\Phi}_{\text{enc}}$ reads X word by word, from left to right, and generates a sequence of *forward annotation vectors* $(\overrightarrow{\mathbf{h}}_1, \overrightarrow{\mathbf{h}}_2, \dots, \overrightarrow{\mathbf{h}}_N)$ at each encoder time step $i \in [1, N]$. Similarly, a backward RNN $\overleftarrow{\Phi}_{\text{enc}}$ reads X from right to left, word by word, and generates a sequence of *backward annotation vectors* $(\overleftarrow{\mathbf{h}}_1, \overleftarrow{\mathbf{h}}_2, \dots, \overleftarrow{\mathbf{h}}_N)$, as in (1):

$$\begin{aligned} \overrightarrow{\mathbf{h}}_i &= \overrightarrow{\Phi}_{\text{enc}}(\mathbf{W}_x[x_i], \overrightarrow{\mathbf{h}}_{i-1}), \\ \overleftarrow{\mathbf{h}}_i &= \overleftarrow{\Phi}_{\text{enc}}(\mathbf{W}_x[x_i], \overleftarrow{\mathbf{h}}_{i+1}). \end{aligned} \quad (1)$$

The final annotation vector for a given time step i is the concatenation of forward and backward vectors $\mathbf{h}_i = [\overrightarrow{\mathbf{h}}_i; \overleftarrow{\mathbf{h}}_i]$.

In other words, each source sequence X is encoded into a sequence of annotation vectors $h = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N)$, which are in turn used by the decoder: essentially a neural language model (LM) (Bengio et al., 2003) conditioned on the previously emitted words and the source sentence via an attention mechanism.

At each time step t of the decoder, we compute a *time-dependent* context vector \mathbf{c}_t based on the annotation vectors h , the decoder’s previous hidden state \mathbf{s}_{t-1} and the target word \tilde{y}_{t-1} emitted by the decoder in the previous time step.²

We follow Bahdanau et al. (2015) and use a single-layer feed-forward network to compute an *expected alignment* $e_{t,i}$ between each source annotation vector \mathbf{h}_i and the target word to be emitted at the current time step t , as in (2):

$$e_{t,i} = \mathbf{v}_a^T \tanh(\mathbf{U}_a \mathbf{s}_{t-1} + \mathbf{W}_a \mathbf{h}_i). \quad (2)$$

In Equation (3), these expected alignments are

normalised and converted into probabilities:

$$\alpha_{t,i} = \frac{\exp(e_{t,i})}{\sum_{j=1}^N \exp(e_{t,j})}, \quad (3)$$

where $\alpha_{t,i}$ are called the model’s *attention weights*, which are in turn used in computing the time-dependent context vector $\mathbf{c}_t = \sum_{i=1}^N \alpha_{t,i} \mathbf{h}_i$. Finally, the context vector \mathbf{c}_t is used in computing the decoder’s hidden state \mathbf{s}_t for the current time step t , as shown in Equation (4):

$$\mathbf{s}_t = \Phi_{\text{dec}}(\mathbf{s}_{t-1}, \mathbf{W}_y[\tilde{y}_{t-1}], \mathbf{c}_t), \quad (4)$$

where \mathbf{s}_{t-1} is the decoder’s previous hidden state, $\mathbf{W}_y[\tilde{y}_{t-1}]$ is the embedding of the word emitted in the previous time step, and \mathbf{c}_t is the updated time-dependent context vector.

We use a single-layer feed-forward neural network to initialise the decoder’s hidden state \mathbf{s}_0 at time step $t = 0$ and feed it the concatenation of the last hidden states of the encoder’s forward RNN ($\overrightarrow{\Phi}_{\text{enc}}$) and backward RNN ($\overleftarrow{\Phi}_{\text{enc}}$), as in (5):

$$\mathbf{s}_0 = \tanh(\mathbf{W}_{di}[\overrightarrow{\mathbf{h}}_1; \overleftarrow{\mathbf{h}}_N] + \mathbf{b}_{di}), \quad (5)$$

where \mathbf{W}_{di} and \mathbf{b}_{di} are model parameters. Since RNNs normally better store information about recent inputs in comparison to more distant ones (Hochreiter and Schmidhuber, 1997; Bahdanau et al., 2015), we expect to initialise the decoder’s hidden state with a strong source sentence representation, i.e. a representation with a strong focus on both the first and the last tokens in the source sentence.

2.2 Multi-modal NMT (MNMT)

Our models can be seen as expansions of the attention-based NMT framework described in §2 with the addition of a *visual component* to incorporate image features.

Simonyan and Zisserman (2014) trained and evaluated an extensive set of deep Convolutional Neural Networks (CNNs) for classifying images into one out of the 1000 classes in ImageNet (Russakovsky et al., 2015). We use their 19-layer VGG network (VGG19) to extract image feature vectors for all images in our dataset. We feed an image to the pre-trained VGG19 network and use the 4096D activations of the penultimate fully-connected layer FC7³ as our *image feature vector*, henceforth referred to as \mathbf{q} .

We propose three different methods to incorporate images into the attentive NMT framework:

²At training time, the correct previous target word y_{t-1} is known and therefore used instead of \tilde{y}_{t-1} . At test or inference time, y_{t-1} is not known and \tilde{y}_{t-1} is used instead. Bengio et al. (2015) discussed problems that may arise from this difference between training and inference distributions.

³We use the activations of the FC7 layer, which encode information about the entire image, of the VGG19 network (configuration E) in Simonyan and Zisserman (2014)’s paper.

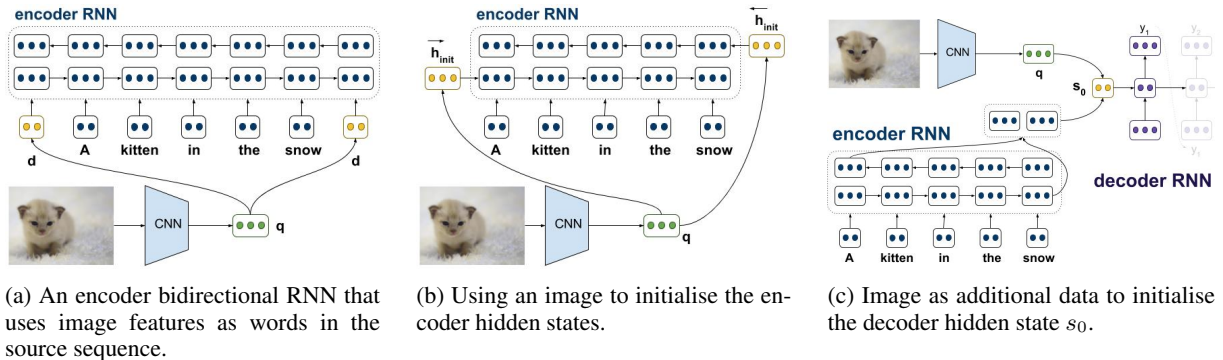


Figure 1: Multi-modal neural machine translation models IMG_W , IMG_E , and IMG_D .

using an image as words in the source sentence (§2.2.1), using an image to initialise the source language encoder (§2.2.2) and the target language decoder (§2.2.3).

We also evaluated a fourth mechanism to incorporate images into NMT, namely to use an image as one of the different contexts available to the decoder at each time step of the decoding process. We added the image features directly as an additional context, in addition to $\mathbf{W}_y[\hat{y}_{t-1}]$, \mathbf{s}_{t-1} and \mathbf{c}_t , to compute the hidden state \mathbf{s}_t of the decoder at a given time step t . We corroborate previous findings by Vinyals et al. (2015) in that adding the image features as such prevents the model from learning.⁴

2.2.1 Images as source words: IMG_W

One way we propose to incorporate images into the encoder is to project an image feature vector into the space of the words of the source sentence. We use the projected image as the first and/or last word of the source sentence and let the attention model learn when to attend to the image representation. Specifically, given the global image feature vector $\mathbf{q} \in \mathbb{R}^{4096}$, we compute (6):

$$\mathbf{d} = \mathbf{W}_I^2 \cdot (\mathbf{W}_I^1 \cdot \mathbf{q} + \mathbf{b}_I^1) + \mathbf{b}_I^2, \quad (6)$$

where $\mathbf{W}_I^1 \in \mathbb{R}^{4096 \times 4096}$ and $\mathbf{W}_I^2 \in \mathbb{R}^{4096 \times d_x}$ are image transformation matrices, $\mathbf{b}_I^1 \in \mathbb{R}^{4096}$ and $\mathbf{b}_I^2 \in \mathbb{R}^{d_x}$ are bias vectors, and d_x is the source words vector space dimensionality, all trained with the model. We then directly use \mathbf{d} as words in the source words vector space: as the first word only (model IMG_{1W}), and as the first and last words of

the source sentence (model IMG_{2W}).

An illustration of this idea is given in Figure 1a, where a source sentence that originally contained N tokens, after including the image as source words will contain $N + 1$ tokens (model IMG_{1W}) or $N + 2$ tokens (model IMG_{2W}). In model IMG_{1W} , the image is projected as the first source word only (solid line in Figure 1a); in model IMG_{2W} , it is projected into the source words space as both first and last words (both solid and dashed lines in Figure 1a).

Given a sequence $X = (x_1, x_2, \dots, x_N)$ in the source language, we concatenate the transformed image vector \mathbf{d} to $\mathbf{W}_x[X]$ and apply the forward and backward encoder RNN passes, generating hidden vectors as in Figure 1a. When computing the context vector \mathbf{c}_t (Equations (2) and (3)), we effectively make use of the transformed image vector, i.e. the $\alpha_{t,i}$ attention weight parameters will use this information to attend or not to the image features.

By including images into the encoder in models IMG_{1W} and IMG_{2W} , our intuition is that (i) by including the image as the *first word*, we propagate image features into the source sentence vector representations when applying the forward RNN $\vec{\Phi}_{\text{enc}}$ (vectors $\vec{\mathbf{h}}_i$), and (ii) by including the image as the *last word*, we propagate image features into the source sentence vector representations when applying the backward RNN $\overleftarrow{\Phi}_{\text{enc}}$ (vectors $\overleftarrow{\mathbf{h}}_i$).

2.2.2 Images for encoder initialisation: IMG_E

In the original attention-based NMT model described in §2, the hidden state of the encoder is initialised with the zero vector $\vec{\mathbf{0}}$. Instead, we propose to use two new single-layer feed-forward neural networks to compute the initial states of the forward RNN $\vec{\Phi}_{\text{enc}}$ and the backward RNN $\overleftarrow{\Phi}_{\text{enc}}$,

⁴Outputs would typically consist of sets of 2-5 words repeated many times, usually without any syntax. For comparison, translations for the translated Multi30k test set (described in §3) achieve just 3.8 BLEU (Papineni et al., 2002), 15.5 METEOR (Denkowski and Lavie, 2014) and 93.0 TER (Snover et al., 2006).

respectively, as illustrated in Figure 1b.

Similarly to §2.2.1, given a global image feature vector $\mathbf{q} \in \mathbb{R}^{4096}$, we compute a vector \mathbf{d} using Equation (6), only this time the parameters \mathbf{W}_f^2 and \mathbf{b}_f^2 project the image features into the same dimensionality as the textual encoder hidden states.

The feed-forward networks used to initialise the encoder hidden state are computed as in (7):

$$\begin{aligned}\overleftarrow{\mathbf{h}}_{\text{init}} &= \tanh(\mathbf{W}_f \mathbf{d} + \mathbf{b}_f), \\ \overrightarrow{\mathbf{h}}_{\text{init}} &= \tanh(\mathbf{W}_b \mathbf{d} + \mathbf{b}_b),\end{aligned}\quad (7)$$

where \mathbf{W}_f and \mathbf{W}_b are multi-modal projection matrices that project the image features \mathbf{d} into the encoder forward and backward hidden states dimensionality, respectively, and \mathbf{b}_f and \mathbf{b}_b are bias vectors.

2.2.3 Images for decoder initialisation: IMG_D

To incorporate an image into the decoder, we introduce a new single-layer feed-forward neural network to be used instead of the one described in Equation 5. Originally, the decoder’s initial hidden state was computed using the concatenation of the last hidden states of the encoder forward RNN ($\overrightarrow{\Phi}_{\text{enc}}$) and backward RNN ($\overleftarrow{\Phi}_{\text{enc}}$), respectively $\overrightarrow{\mathbf{h}}_N$ and $\overleftarrow{\mathbf{h}}_1$.

Our proposal is that we include the image features as additional input to initialise the decoder hidden state at time step $t = 0$, as in (8):

$$\mathbf{s}_0 = \tanh(\mathbf{W}_{di}[\overleftarrow{\mathbf{h}}_1; \overrightarrow{\mathbf{h}}_N] + \mathbf{W}_m \mathbf{d} + \mathbf{b}_{di}), \quad (8)$$

where \mathbf{W}_m is a multi-modal projection matrix that projects the image features \mathbf{d} into the decoder hidden state dimensionality and \mathbf{W}_{di} and \mathbf{b}_{di} are the same as in Equation (5).

Once again we compute \mathbf{d} by applying Equation (6) onto a global image feature vector $\mathbf{q} \in \mathbb{R}^{4096}$, only this time the parameters \mathbf{W}_f^2 and \mathbf{b}_f^2 project the image features into the same dimensionality as the decoder hidden states. We illustrate this idea in Figure 1c.

3 Data set

Our multi-modal NMT models need bilingual sentences accompanied by one or more images as training data. The original Flickr30k data set contains 30K images and 5 English sentence descriptions for each image (Young et al., 2014). We use the translated and the comparable Multi30k datasets (Elliott et al., 2016), henceforth referred to as M30k_T and M30k_C, respectively, which are multilingual expansions of the original Flickr30k.

For each of the 30K images in the Flickr30k, the M30k_T has one of its English descriptions manually translated into German by a professional translator. Training, validation and test sets contain 29K, 1014 and 1K images, respectively, each accompanied by one sentence pair (the original English sentence and its German translation). For each of the 30K images in the Flickr30k, the M30k_C has five descriptions in German collected independently of the English descriptions. Training, validation and test sets contain 29K, 1014 and 1K images, respectively, each accompanied by 5 English and 5 German sentences.

We use the scripts in Moses (Koehn et al., 2007) to normalise, truecase and tokenize English and German descriptions and we also convert space-separated tokens into subwords (Sennrich et al., 2016b). All models use a common vocabulary of $\sim 83\text{K}$ English and $\sim 91\text{K}$ German subword tokens. If sentences in English or German are longer than 80 tokens, they are discarded.

We use the entire M30k_T training set for training, its validation set for model selection with BLEU, and its test set to evaluate our models. In order to study the impact that additional training data brings to the models, we use the baseline model described in §2 trained on the textual part of the M30k_T data set (German→English and English→German) without the images to build back-translation models (Sennrich et al., 2016a). We back-translate the 145K German (English) descriptions in the M30k_C into English (German) and include the triples (synthetic English description, German description, image) as additional training data when translating into German, and (synthetic German description, English description, image) when translating into English.

We train models to translate from English into German and from German into English, and report evaluation of cased, tokenized sentences with punctuation.

4 Experimental setup

Our encoder is a bidirectional RNN with GRU (one 1024D single-layer forward RNN and one 1024D single-layer backward RNN). Source and target word embeddings are 620D each and both are trained jointly with our model. All non-recurrent matrices are initialised by sampling from a Gaussian ($\mu = 0, \sigma = 0.01$), recurrent matrices are orthogonal and bias vectors are all initialised

to zero. Our decoder RNN also uses GRU and is a neural LM (Bengio et al., 2003) conditioned on its previous emissions and the source sentence by means of the source attention mechanism.

Image features are obtained by feeding images to the pre-trained VGG19 network of Simonyan and Zisserman (2014) and using the activations of the penultimate fully-connected layer FC7. We apply dropout with a probability of 0.2 in both source and target word embeddings and with a probability of 0.5 in the image features (in all MNMT models), in the encoder and decoder RNNs inputs and recurrent connections, and before the readout operation in the decoder RNN. We follow Gal and Ghahramani (2016) and apply dropout to the encoder bidirectional RNN and decoder RNN using the same mask in all time steps.

Our models are trained using stochastic gradient descent with Adadelta (Zeiler, 2012) and minibatches of size 40 for improved generalisation (Keskar et al., 2017), where each training instance consists of one English sentence, one German sentence and one image. We apply early stopping for model selection based on BLEU scores, so that if a model does not improve on the validation set for more than 20 epochs, training is halted.

We evaluate our models’ translation quality quantitatively in terms of BLEU4 (Papineni et al., 2002), METEOR (Denkowski and Lavie, 2014), TER (Snover et al., 2006), and chrF3 scores⁵ (Popović, 2015) and we report statistical significance for the three first metrics using approximate randomisation computed with MultEval (Clark et al., 2011).

As our main baseline we train an attention-based NMT model (§2) in which only the textual part of M30k_T is used for training. We also train a PBSMT model built with Moses on the same English→German (German→English) data, where the LM is a 5-gram LM with modified Kneser-Ney smoothing (Kneser and Ney, 1995) trained on the German (English) of the M30k_T dataset. We use minimum error rate training (Och, 2003) for tuning the model parameters for BLEU scores. Our third baseline (English→German), is the best comparable multi-modal model by Huang et al. (2016) and also their best model with additional object detections: respectively models *m1* (image at head) and *m3* in the authors’ paper. Finally, our fourth baseline (German→English) is

⁵We specifically compute character 6-gram F3 scores.

	BLEU4↑	METEOR↑	TER↓	chrF3↑
English→German				
PBSMT	32.9	<u>54.1</u>	<u>45.1</u>	<u>67.4</u>
NMT	<u>33.7</u>	52.3	46.7	64.5
Huang + RCNN	35.1	52.2	—	—
IMG _{1W}	37.1 ^{†‡} († 3.4)	54.5 [‡] († 0.4)	42.7 ^{†‡} (↓ 2.4)	66.9 (↓ 0.5)
IMG _{2W}	36.9 ^{†‡} († 3.2)	54.3 [‡] († 0.2)	41.9^{†‡} (↓ 3.2)	66.8 (↓ 0.6)
IMG _E	37.1 ^{†‡} († 3.4)	55.0 ^{†‡} († 0.9)	43.1 ^{†‡} (↓ 2.0)	67.6 († 0.2)
IMG _D	37.3^{†‡} († 3.6)	55.1^{†‡} († 1.0)	42.8 ^{†‡} (↓ 2.3)	67.7 († 0.3)
IMG _{2W+D}	35.7 ^{†‡} († 2.0)	53.6 [‡] (↓ 0.5)	43.3 ^{†‡} (↓ 1.8)	66.2 (↓ 1.2)
IMG _{E+D}	37.0 ^{†‡} († 3.3)	54.7 [‡] († 0.6)	42.6 ^{†‡} (↓ 2.5)	67.2 (↓ 0.2)
German→English				
PBSMT	32.8	34.8	43.9	61.8
NMT	<u>38.2</u>	<u>35.8</u>	<u>40.2</u>	<u>62.8</u>
IMG _{2W}	39.5 [‡] († 1.3)	37.1 ^{†‡} († 1.3)	37.1 ^{†‡} (↓ 3.1)	63.8 († 1.0)
IMG _E	41.1 ^{†‡} († 2.9)	37.7 ^{†‡} († 1.9)	37.9 ^{†‡} (↓ 2.3)	65.7 († 2.9)
IMG _D	41.3 ^{†‡} († 3.1)	37.8 ^{†‡} († 2.0)	37.9 ^{†‡} (↓ 2.3)	65.7 († 2.9)
IMG _{2W+D}	39.9 ^{†‡} († 1.7)	37.2 ^{†‡} († 1.4)	37.0^{†‡} (↓ 3.2)	64.4 († 1.6)
IMG _{E+D}	41.9^{†‡} († 3.7)	37.9^{†‡} († 2.1)	37.1 ^{†‡} (↓ 3.1)	66.0 († 3.2)

Table 1: BLEU4, METEOR, chrF3 (higher is better) and TER scores (lower is better) on the M30k_T test set for the two text-only baselines PBSMT and NMT, the two multi-modal NMT models by Huang et al. (2016) (English→German only) and our MNMT models that: (i) use images as words in the source sentence (IMG_{1W}, IMG_{2W}), (ii) use images to initialise the encoder (IMG_E), and (iii) use images as additional data to initialise the decoder (IMG_D). Best text-only baselines are underlined and best overall results appear in bold. We highlight in parentheses the improvements brought by our models compared to the best corresponding text-only baseline score. Results differ significantly from PBSMT baseline (†) or NMT baseline (‡) with $p = 0.05$.

the best-performing model in the WMT’16 multi-modal MT shared task (Shah et al., 2016), henceforth PBSMT⁺. It uses image features as well as additional data from WordNet (Miller, 1995) to re-rank n -best lists.

4.1 Results

The Multi30K dataset contains images and bilingual descriptions. Overall, it is a small dataset with a small vocabulary whose sentences have simple syntactic structures and not much ambiguity (Elliott et al., 2016). This is reflected in the fact that even the simplest baselines perform fairly well on it, i.e. the smallest BLEU scores of 32.9 for translating into German, which are still reasonably good results.

	BLEU4 \uparrow	METEOR \uparrow	TER \downarrow	chrF3 \uparrow
English\rightarrowGerman				
original training data				
IMG _{2W}	36.9	54.3	41.9	66.8
IMG _E	37.1	55.0	43.1	67.6
IMG _D	37.3	55.1	42.8	67.7
+ back-translated training data				
PBSMT	34.0	<u>55.0</u>	44.7	<u>68.0</u>
NMT	<u>35.5</u>	53.4	<u>43.3</u>	65.3
IMG _{2W}	36.7 \ddagger (\uparrow 1.2)	54.6 \ddagger (\downarrow 0.4)	42.0 \ddagger (\downarrow 1.3)	66.8 (\downarrow 1.2)
IMG _E	38.5\ddagger (\uparrow 3.0)	55.7 \ddagger (\uparrow 0.9)	41.4\ddagger (\downarrow 1.9)	68.3 (\uparrow 0.3)
IMG _D	38.5\ddagger (\uparrow 3.0)	55.9\ddagger (\uparrow 1.1)	41.6 \ddagger (\downarrow 1.7)	68.4 (\uparrow 0.4)
German\rightarrowEnglish				
PBSMT ⁺	42.5	<u>39.5</u>	<u>35.6</u>	68.7
original training data				
IMG _{2W}	39.5	37.1	37.1	63.8
IMG _E	41.1	37.7	37.9	65.7
IMG _D	41.3	37.8	37.9	65.7
+ back-translated training data				
NMT	<u>42.6</u>	38.9	36.1	67.6
IMG _{2W}	42.4 \ddagger (\downarrow 0.2)	39.0 \ddagger (\uparrow 0.1)	34.7\ddagger (\downarrow 1.4)	67.6 (\uparrow 0.0)
IMG _E	43.9\ddagger (\uparrow 1.3)	39.7 \ddagger (\uparrow 0.8)	34.8 \ddagger (\downarrow 1.3)	68.7 (\uparrow 1.1)
IMG _D	43.4 \ddagger (\uparrow 0.8)	39.3 \ddagger (\uparrow 0.4)	35.2 \ddagger (\downarrow 0.9)	67.8 (\uparrow 0.2)
Improvements (original vs. + back-translated)				
English \rightarrow German / German \rightarrow English				
IMG _{2W}	\downarrow 0.2 / \uparrow 2.9	\uparrow 0.1 / \uparrow 1.9	\uparrow 0.1 / \downarrow 2.4	\uparrow 0.0 / \uparrow 3.8
IMG _E	\uparrow 1.4 / \uparrow 2.8	\uparrow 0.7 / \uparrow 2.0	\downarrow 1.8 / \downarrow 3.1	\uparrow 0.7 / \uparrow 2.9
IMG _D	\uparrow 1.2 / \uparrow 2.1	\uparrow 0.8 / \uparrow 1.5	\downarrow 1.2 / \downarrow 2.7	\uparrow 0.7 / \uparrow 2.1

Table 2: BLEU4, METEOR, TER and chrF3 scores on the M30k_T test set for models trained on original and additional back-translated data. Best text-only baselines are underscored and best overall results in bold. We highlight in parentheses the improvements brought by our models compared to the best baseline score. Results differ significantly from PBSMT baseline (\dagger) or NMT baseline (\ddagger) with $p = 0.05$. We also show the improvements each model yields in each metric when only trained on the original M30k_T training set vs. also including additional back-translated data. PBSMT⁺ is the best model in the multi-modal MT shared task (Specia et al., 2016).

Multi30k In Table 1, we show results for translating from English \rightarrow German and German \rightarrow English. When translating into German, our multi-modal models perform well, with models IMG_E and IMG_D improving on both baselines according to all metrics analysed. We also note that all models but IMG_{2W+D} perform consistently better than the strong multi-modal NMT baseline of Huang et al. (2016), even when this model has access to more data (+RCNN features).⁶ Combining image features in the

⁶In fact, model IMG_{2W+D} still improves on the multi-modal baseline of Huang et al. (2016) when trained on the

encoder and the decoder at the same time does not seem to improve results compared to using the image features in only the encoder or the decoder when translating into German. To the best of our knowledge, it is the first time a purely neural model significantly improves over a PBSMT model in all metrics on this data set.

When translating into English, all multi-modal models significantly improve over the NMT baseline, with the only exception being model IMG_{2W}'s BLEU scores. In this scenario, model IMG_{E+D} is the best performing one according to all but one metric. However, differences between multi-modal models are not statistically significant, i.e. all multi-modal models but IMG_{2W} and IMG_{2W+D} perform comparably.

Additional back-translated data Arguably, the main downside of applying multi-modal NMT in a real-world scenario is the small amount of publicly available training data (\sim 30K entries). For that reason, we back-translated the German and English sentences in the M30k_C and created two sets of 145K synthetic triples, one for each translation direction, as described in §3.

In Table 2, we present results for some of the models evaluated in Table 1 but when also trained on the additional data. In order to add more data to our PBSMT baseline, we simply added the German sentences in the M30k_C to train the LM.⁷ We also include results for PBSMT⁺, which uses image features as well as additional features extracted using WordNet (Shah et al., 2016). When translating into German, both our models IMG_E and IMG_D that use global image features to initialise the encoder and the decoder, respectively, significantly improve according to BLEU, METEOR and TER with the additional back-translated data, and also achieved better chrF3 scores. When translating into English, IMG_E is the only model to significantly improve over both baselines according to all metrics with the additional back-translated data, also improving chrF3 scores. We highlight that our best-performing model IMG_E significantly outperforms PBSMT⁺ according to BLEU and TER, and all our other multi-modal models perform comparably to it. This is a noteworthy finding, since

same data.

⁷Adding the synthetic sentence pairs to train the baseline PBSMT model, as we did with all neural MT models, deteriorated the results.

Multi30k test set (English→German)					
	Ensemble?	BLEU4 ↑	METEOR ↑	TER ↓	chrF3 ↑
IMG _D	×	37.3	55.1	42.8	67.7
IMG _D + IMG _E	✓	40.1 (↑ 2.8)	58.5 (↑ 3.4)	40.7 (↓ 2.1)	68.1 (↑ 0.4)
IMG _D + IMG _E + IMG _{2W}	✓	41.0 (↑ 3.7)	58.9 (↑ 3.8)	39.7 (↓ 3.1)	68.3 (↑ 0.6)
IMG _D + IMG _E + IMG _{2W} + IMG _D	✓	41.3 (↑ 4.0)	59.2 (↑ 4.1)	39.5 (↓ 3.3)	68.5 (↑ 0.8)

Table 3: Results for different combinations of multi-modal models, all trained on the original M30k_T training data only, evaluated on the M30k_T test set.

PBSMT⁺ uses image features as well as additional data from WordNet and, to the best of our knowledge, is the best published model in this language pair and data set to date.

Ensemble decoding We now report on how can ensemble decoding be used to improve translations obtained with our multi-modal NMT models. In order to do that, we use different combinations of models trained on the original M30k_T training set to translate from English into German. We built ensembles of different models by starting with our best performing multi-modal model on this language pair and data set, IMG_D, and by adding new models to the ensemble one by one, until we reach a maximum of four independent models, all of which are trained separately and on the original M30k_T training data only. In Table 3, we show results when translating the M30k_T’s test set. These models were also evaluated in our recent participation in the WMT 2017 multi-modal MT shared task (Calixto et al., 2017a).

We first note that to add more models to the ensemble seems to always improve translations, and by a considerable margin (~ 3 BLEU/METEOR points). Adding model IMG_{2W} to the ensemble already consisting of models IMG_E and IMG_D improves translations according to all metrics evaluated. This is an interesting result, since compared to these other two multi-modal models, model IMG_{2W} performs poorly according to BLEU, METEOR and chrF3. Regardless of that fact, our best results are obtained with an ensemble of four different multi-modal models, including model IMG_{2W}.

By using an ensemble of four different multi-modal NMT models trained on the translated Multi30k training data, we were able to obtain translations comparable to or even better than those obtained with the strong multi-modal NMT model of Calixto et al. (2017b), which is pre-trained on large amounts of English–German data

and uses local image features. Finally, we have recently participated in the WMT 2017 multi-modal MT shared task, and our system submissions ranked among the best performing systems under the constrained data regime (Calixto et al., 2017a). We note that our models performed particularly well on the ambiguous MSCOCO test set (Elliott et al., 2017), which indicate its ability to use the image information in disambiguating difficult source sentences into their correct translations.

5 Error Analysis

In Table 4 we show translations into German generated by different models for one entry of the M30k test set. In this example, the last three multi-modal models extrapolate the reference+image and describe “ceremony” as a “wedding ceremony” (IMG_{2W}) and as an “Olympics ceremony” (IMG_E and IMG_D). This could be due to the fact that the training set is small, depicts a small variation of different scenes and contains different forms of biases (van Miltenburg, 2015).

In Table 5, we draw attention to an example where some models generate what seems to be *novel visual terms*. Neither the source German sentence nor the English reference translation contained the translated units “having fun” or “Mexican restaurant”, although both could have been inferred at least partially from the image. In this example, the visual term “having fun” is also generated by the baseline NMT model, making it clear that at times what seems like a translation extracted exclusively from the image may have been learnt from the training text data. However, none of the two baselines translated “Mexikanischen Setting” as “Mexican restaurant”, but four out of the five multi-modal models did. The multi-modal models also had problems translating the German “trinkt Shots” (drinking shots). We observe translations such as “having drinks” (IMG_{2W}), which



src.	a woman with long hair is at a graduation ceremony .
ref.	eine Frau mit langen Haaren bei einer Abschluss Feier.
NMT	eine Frau mit langen Haaren ist an einer Sta Zeremonie.
PBSMT	eine Frau mit langen Haaren steht an einem Abschluss
IMG _{1W}	eine Frau mit langen Haaren ist an einer warmen Zeremonie teil
IMG _{2W}	eine Frau mit langen Haaren steht bei einer Hochzeit Feier .
IMG _E	eine lang haarige Frau bei einer olympischen Zeremonie.
IMG _D	eine lang haarige Frau bei einer olympischen Zeremonie.

Table 4: Translations for the 668th example in the M30k test set.



src.	eine Gruppe junger Menschen trinkt Shots in einem Mexikanischen Setting .
ref.	a group of young people take shots in a Mexican setting .
NMT	a group of young people are having fun in an auditorium .
PBSMT	a group of young people drinking at a Shots Mexikanischen Setting .
IMG _{2W}	a group of young people having drinks in a Mexican restaurant .
IMG _E	a group of young people drinking apples in a Mexican restaurant .
IMG _D	a group of young people drinking food in a Mexican restaurant .
IMG _{2W+D}	a group of young people having fun in a Mexican room .
IMG _{E+D}	a group of young people drinking dishes in a Mexican restaurant .

Table 5: Translations for 300th example in the M30k test set.

although not a novel translation is still correct, but also “drinking apples” (IMG_E), “drinking food” (IMG_D), and “drinking dishes” (IMG_{E+D}), which are clearly incorrect.

6 Conclusions and future work

In this work, we introduced models that incorporate images into state-of-the-art attention-based NMT, by using images as words in the source sentence, to initialise the encoder’s hidden state and as additional data in the initialisation of the decoder’s hidden state. The intuition behind our effort is to use images to visually *ground* translations, and consequently increase translation quality. We demonstrate with extensive experiments that adding global image features into NMT significantly improves the translations of image descriptions compared to text-only NMT and PBSMT. It also improves significantly on the previous state-of-the-art model of Huang et al. (2016) (English→German), and performs comparably to the best published results of Shah et al. (2016) (German→English). Overall, we note that using images as words in the source sequence (IMG_{1W}, IMG_{2W}), an idea similarly entertained by Huang et al. (2016), does not fare as well as to directly incorporate the image either in the encoder or the decoder (IMG_E and IMG_D), independently of the target language. The fact that multi-modal NMT models can benefit from back-translated data is

also an interesting finding.

In future work, we will conduct a more systematic study on the impact that synthetic back-translated data brings to multi-modal NMT, and run an error analysis to identify what particular types of errors our models make (and prevent).

Acknowledgments

This project has received funding from Science Foundation Ireland in the ADAPT Centre for Digital Content Technology (www.adaptcentre.ie) at Dublin City University funded under the SFI Research Centres Programme (Grant 13/RC/2106) co-funded under the European Regional Development Fund and the European Union Horizon 2020 research and innovation programme under grant agreement 645452 (QT21).

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *International Conference on Learning Representations, ICLR 2015*. San Diego, California.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam M. Shazeer. 2015. *Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks*. In *Advances in Neural Information Processing Systems, NIPS*. <http://arxiv.org/abs/1506.03099>.

- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. *A Neural Probabilistic Language Model*. *J. Mach. Learn. Res.* 3:1137–1155. <http://dl.acm.org/citation.cfm?id=944919.944966>.
- Ozan Caglayan, Walid Aransa, Yaxing Wang, Marc Masana, Mercedes García-Martínez, Fethi Bougares, Loïc Barrault, and Joost van de Weijer. 2016. *Does Multimodality Help Human and Machine for Translation and Image Captioning?* In *Proceedings of the First Conference on Machine Translation*. Berlin, Germany, pages 627–633. <http://www.aclweb.org/anthology/W/W16/W16-2358>.
- Iacer Calixto, Koel Dutta Chowdhury, and Qun Liu. 2017a. DCU System Report on the WMT 2017 Multi-modal Machine Translation Task. In *Proceedings of the Second Conference on Machine Translation*. Copenhagen, Denmark.
- Iacer Calixto, Desmond Elliott, and Stella Frank. 2016. *DCU-UvA Multimodal MT System Report*. In *Proceedings of the First Conference on Machine Translation*. Berlin, Germany, pages 634–638. <http://www.aclweb.org/anthology/W/W16/W16-2359>.
- Iacer Calixto, Qun Liu, and Nick Campbell. 2017b. *Doubly-Attentive Decoder for Multi-modal Neural Machine Translation*. In *Proceedings of the 55th Conference of the Association for Computational Linguistics: Volume 1, Long Papers*. Vancouver, Canada (Paper Accepted).
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014a. On the properties of neural machine translation: Encoder–decoder approaches. *Syntax, Semantics and Structure in Statistical Translation* page 103.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014b. *Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation*. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar, pages 1724–1734. <http://www.aclweb.org/anthology/D14-1179>.
- Jonathan H. Clark, Chris Dyer, Alon Lavie, and Noah A. Smith. 2011. *Better Hypothesis Testing for Statistical Machine Translation: Controlling for Optimizer Instability*. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers - Volume 2*. Portland, Oregon, HLT '11, pages 176–181. <http://dl.acm.org/citation.cfm?id=2002736.2002774>.
- Michael Denkowski and Alon Lavie. 2014. *Meteor Universal: Language Specific Translation Evaluation for Any Target Language*. In *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*.
- Jeff Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Trevor Darrell, and Kate Saenko. 2015. *Long-term Recurrent Convolutional Networks for Visual Recognition and Description*. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*. Boston, US, pages 2625–2634.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. *Multi-Task Learning for Multiple Language Translation*. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Beijing, China, pages 1723–1732. <http://www.aclweb.org/anthology/P15-1166>.
- Desmond Elliott, Stella Frank, Loïc Barrault, Fethi Bougares, and Lucia Specia. 2017. *Findings of the Second Shared Task on Multimodal Machine Translation and Multilingual Image Description*. In *Proceedings of the Second Conference on Machine Translation*. Copenhagen, Denmark.
- Desmond Elliott, Stella Frank, and Eva Hasler. 2015. *Multi-language image description with neural sequence models*. *CoRR* abs/1510.04709. <http://arxiv.org/abs/1510.04709>.
- 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, VL@ACL 2016*. Berlin, Germany. <http://aclweb.org/anthology/W/W16/W16-3210.pdf>.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. *Multi-Way, Multilingual Neural Machine Translation with a Shared Attention Mechanism*. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California, pages 866–875. <http://www.aclweb.org/anthology/N16-1101>.
- Yarin Gal and Zoubin Ghahramani. 2016. *A Theoretically Grounded Application of Dropout in Recurrent Neural Networks*. In *Advances in Neural Information Processing Systems, NIPS*, Barcelona, Spain, pages 1019–1027. <http://papers.nips.cc/paper/6241-a-theoretically-grounded-application-of-dropout-in-recurrent-neural-networks.pdf>.
- Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. *Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation*. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition*. Washington, DC, USA, CVPR '14, pages 580–587. <https://doi.org/10.1109/CVPR.2014.81>.
- Alex Graves. 2013. *Generating Sequences With Recurrent Neural Networks*. *CoRR* abs/1308.0850. <http://arxiv.org/abs/1308.0850>.

- Sepp Hochreiter and Jürgen Schmidhuber. 1997. **Long Short-Term Memory**. *Neural Comput.* 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013. **Framing Image Description As a Ranking Task: Data, Models and Evaluation Metrics**. *J. Artif. Int. Res.* 47(1):853–899. <http://dl.acm.org/citation.cfm?id=2566972.2566993>.
- Po-Yao Huang, Frederick Liu, Sz-Rung Shiang, Jean Oh, and Chris Dyer. 2016. **Attention-based Multimodal Neural Machine Translation**. In *Proceedings of the First Conference on Machine Translation*. Berlin, Germany, pages 639–645. <http://www.aclweb.org/anthology/W/W16/W16-2360>.
- Nal Kalchbrenner and Phil Blunsom. 2013. **Recurrent Continuous Translation Models**. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013*. Seattle, USA, pages 1700–1709.
- Andrej Karpathy and Li Fei-Fei. 2015. **Deep visual-semantic alignments for generating image descriptions**. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015*. Boston, Massachusetts, pages 3128–3137.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2017. **On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima**. In *International Conference on Learning Representations, ICLR 2017*. Toulon, France.
- Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel. 2014. **Unifying visual-semantic embeddings with multimodal neural language models**. *CoRR* abs/1411.2539. <http://arxiv.org/abs/1411.2539>.
- Reinhard Kneser and Hermann Ney. 1995. **Improved backing-off for m-gram language modeling**. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. Detroit, Michigan, volume I, pages 181–184.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. **Moses: Open Source Toolkit for Statistical Machine Translation**. In *Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions*. Association for Computational Linguistics, Prague, Czech Republic, ACL '07, pages 177–180. <http://dl.acm.org/citation.cfm?id=1557769.1557821>.
- Jindřich Libovický, Jindřich Helcl, Marek Tlustý, Ondřej Bojar, and Pavel Pecina. 2016. **CUNI System for WMT16 Automatic Post-Editing and Multimodal Translation Tasks**. In *Proceedings of the First Conference on Machine Translation*. Berlin, Germany, pages 646–654. <http://www.aclweb.org/anthology/W/W16/W16-2361>.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. **Effective Approaches to Attention-based Neural Machine Translation**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Lisbon, Portugal, pages 1412–1421.
- Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, and Alan L. Yuille. 2014. **Explain Images with Multimodal Recurrent Neural Networks**. <http://arxiv.org/abs/1410.1090>.
- George A. Miller. 1995. **Wordnet: A lexical database for english**. *Commun. ACM* 38(11):39–41. <https://doi.org/10.1145/219717.219748>.
- Franz Josef Och. 2003. **Minimum Error Rate Training in Statistical Machine Translation**. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics - Volume 1*. Sapporo, Japan, ACL '03, pages 160–167. <https://doi.org/10.3115/1075096.1075117>.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **BLEU: A Method for Automatic Evaluation of Machine Translation**. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*. Philadelphia, Pennsylvania, ACL '02, pages 311–318. <https://doi.org/10.3115/1073083.1073135>.
- Maja Popović. 2015. **chrF: character n-gram f-score for automatic mt evaluation**. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*. Lisbon, Portugal, pages 392–395. <http://aclweb.org/anthology/W15-3049>.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. **ImageNet Large Scale Visual Recognition Challenge**. *International Journal of Computer Vision (IJCV)* 115(3):211–252. <https://doi.org/10.1007/s11263-015-0816-y>.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. **Improving Neural Machine Translation Models with Monolingual Data**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany, pages 86–96. <http://www.aclweb.org/anthology/P16-1009>.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. **Neural Machine Translation of Rare Words with Subword Units**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany, pages 1715–1725. <http://www.aclweb.org/anthology/P16-1162>.

- Kashif Shah, Josiah Wang, and Lucia Specia. 2016. **SHEF-Multimodal: Grounding Machine Translation on Images**. In *Proceedings of the First Conference on Machine Translation*. Berlin, Germany, pages 660–665. <http://www.aclweb.org/anthology/W/W16/W16-2363>.
- K. Simonyan and A. Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of Association for Machine Translation in the Americas*. Cambridge, MA, pages 223–231.
- Lucia Specia, Stella Frank, Khalil Sima'an, and Desmond Elliott. 2016. **A Shared Task on Multimodal Machine Translation and Crosslingual Image Description**. In *Proceedings of the First Conference on Machine Translation, WMT 2016*. Berlin, Germany, pages 543–553. <http://aclweb.org/anthology/W/W16/W16-2346.pdf>.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to Sequence Learning with Neural Networks. In *Advances in Neural Information Processing Systems*. Montréal, Canada, pages 3104–3112.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. **Modeling Coverage for Neural Machine Translation**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany, pages 76–85. <http://www.aclweb.org/anthology/P16-1008>.
- Emiel van Miltenburg. 2015. Stereotyping and bias in the flickr30k dataset. In *Proceedings of the Workshop on Multimodal Corpora, MMC-2016*. Portorož, Slovenia, pages 1–4.
- Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to Sequence - Video to Text. In *Proceedings of the IEEE International Conference on Computer Vision*. Santiago, Chile, pages 4534–4542.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015*. Boston, Massachusetts, pages 3156–3164.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. **Show, attend and tell: Neural image caption generation with visual attention**. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*. JMLR Workshop and Conference Proceedings, Lille, France, pages 2048–2057. <http://jmlr.org/proceedings/papers/v37/xuc15.pdf>.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics* 2:67–78.
- Matthew D. Zeiler. 2012. **ADADELTA: An Adaptive Learning Rate Method**. *CoRR* abs/1212.5701. <http://arxiv.org/abs/1212.5701>.