

Doubly-Attentive Decoder for Multi-modal Neural Machine Translation

Iacer Calixto
ADAPT Centre
School of Computing
Dublin City University
Dublin, Ireland

Qun Liu
ADAPT Centre
School of Computing
Dublin City University
Dublin, Ireland

Nick Campbell
ADAPT Centre
Speech Communication Lab
Trinity College Dublin
Dublin 2, Ireland

{iacer.calixto,qun.liu,nick.campbell}@adaptcentre.ie

Abstract

We introduce a Multi-modal Neural Machine Translation model in which a doubly-attentive decoder naturally incorporates *spatial* visual features obtained using pre-trained convolutional neural networks, bridging the gap between image description and translation. Our decoder learns to attend to source-language words and parts of an image independently by means of two *separate attention mechanisms* as it generates words in the target language. We find that our model can efficiently exploit not just back-translated in-domain multi-modal data but also large general-domain text-only MT corpora. We also report state-of-the-art results on the Multi30k data set.

1 Introduction

Neural Machine Translation (NMT) has been successfully tackled as a *sequence to sequence* learning problem (Kalchbrenner and Blunsom, 2013; Cho et al., 2014b; Sutskever et al., 2014) where each training example consists of one source and one target variable-length sequences, with no prior information on the alignment between the two.

In the context of NMT, Bahdanau et al. (2015) first proposed to use an *attention mechanism* in the decoder, which is trained to attend to the relevant source-language words as it generates each word of the target sentence. Similarly, Xu et al. (2015) proposed an attention-based model for the task of image description generation (IDG) where a model learns to attend to specific parts of an image representation (the source) as it generates its description (the target) in natural language.

We are inspired by recent successes in applying attention-based models to NMT and IDG. In this

work, we propose an end-to-end attention-based multi-modal neural machine translation (MNMT) model which effectively incorporates *two independent attention mechanisms*, one over source-language words and the other over different areas of an image.

Our main contributions are:

- We propose a novel attention-based MNMT model which incorporates spatial visual features in a separate visual attention mechanism;
- We use a medium-sized, back-translated multi-modal in-domain data set and large general-domain text-only MT corpora to pre-train our models and show that our MNMT model can efficiently exploit both;
- We show that images bring useful information into an NMT model, e.g. in situations in which sentences describe objects illustrated in the image.

To the best of our knowledge, previous MNMT models in the literature that utilised spatial visual features did not significantly improve over a comparable model that used global visual features or even only textual features (Caglayan et al., 2016a; Calixto et al., 2016; Huang et al., 2016; Libovický et al., 2016; Specia et al., 2016). In this work, we wish to address this issue and propose an MNMT model that uses, in addition to an attention mechanism over the source-language words, an additional visual attention mechanism to incorporate spatial visual features, and still improves on simpler text-only and multi-modal attention-based NMT models.

The remainder of this paper is structured as follows. We first briefly revisit the attention-based NMT framework (§2) and expand it into an MNMT framework (§3). In §4, we introduce the

datasets we use to train and evaluate our models, in §5 we discuss our experimental setup and analyse and discuss our results. Finally, in §6 we discuss relevant related work and in §7 we draw conclusions and provide avenues for future work.

2 Background and Notation

2.1 Attention-based NMT

In this section, we describe the attention-based NMT model introduced by Bahdanau et al. (2015). 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 to build a single neural network that translates X into Y by directly learning to model $p(Y | X)$. The entire network consists of one *encoder* and one *decoder* with one *attention mechanism*, typically implemented as two Recurrent Neural Networks (RNN) and one multilayer perceptron, respectively. Each x_i is a row index in a source lookup or word embedding matrix $E_x \in \mathbb{R}^{|V_x| \times d_x}$, as well as each y_j being an index in a target lookup or word embedding matrix $E_y \in \mathbb{R}^{|V_y| \times d_y}$, V_x and V_y are source and target vocabularies, and d_x and d_y are source and target word embeddings dimensionalities, respectively.

The encoder is a bi-directional RNN with GRU (Cho et al., 2014a), where a forward RNN $\overrightarrow{\Phi}_{\text{enc}}$ reads X word by word, from left to right, and generates a sequence of *forward annotation vectors* $(\overrightarrow{h}_1, \overrightarrow{h}_2, \dots, \overrightarrow{h}_N)$ 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* $(\overleftarrow{h}_N, \overleftarrow{h}_{N-1}, \dots, \overleftarrow{h}_1)$. The final annotation vector is the concatenation of forward and backward vectors $\mathbf{h}_i = [\overrightarrow{h}_i; \overleftarrow{h}_i]$, and $C = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N)$ is the set of source annotation vectors.

These annotation vectors are in turn used by the decoder, which is essentially a neural language model (LM) (Bengio et al., 2003) conditioned on the previously emitted words and the source sentence via an attention mechanism. A multilayer perceptron is used to initialise the decoder’s hidden state s_0 at time step $t = 0$, where the input to this network is the concatenation of the last forward and backward vectors $[\overrightarrow{h}_N; \overleftarrow{h}_1]$.

At each time step t of the decoder, a *time-dependent* source context vector c_t is computed based on the annotation vectors C and the decoder previous hidden state s_{t-1} . This is part of the for-

mulation of the *conditional GRU* and is described further in §2.2. In other words, the encoder is a bi-directional RNN with GRU and the decoder is an RNN with a conditional GRU.

Given a hidden state s_t , the probabilities for the next target word are computed using one projection layer followed by a softmax layer as illustrated in eq. (1), where the matrices L_o , L_s , L_w and L_c are transformation matrices and c_t is a time-dependent source context vector generated by the conditional GRU.

2.2 Conditional GRU

The conditional GRU¹, illustrated in Figure 1, has three main components computed at each time step t of the decoder:

- REC₁ computes a hidden state proposal s'_t based on the previous hidden state s_{t-1} and the previously emitted word \hat{y}_{t-1} ;
- ATT_{src}² is an attention mechanism over the hidden states of the source-language RNN and computes c_t using all source annotation vectors C and the hidden state proposal s'_t ;
- REC₂ computes the final hidden state s_t using the hidden state proposal s'_t and the time-dependent source context vector c_t .

First, a single-layer feed-forward network is used to compute an *expected alignment* $e_{t,i}^{\text{src}}$ between each source annotation vector \mathbf{h}_i and the target word \hat{y}_t to be emitted at the current time step t , as shown in Equations (2) and (3):

$$e_{t,i}^{\text{src}} = (\mathbf{v}_a^{\text{src}})^T \tanh(\mathbf{U}_a^{\text{src}} s'_t + \mathbf{W}_a^{\text{src}} \mathbf{h}_i), \quad (2)$$

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

where $\alpha_{t,i}^{\text{src}}$ is the normalised alignment matrix between each source annotation vector \mathbf{h}_i and the word \hat{y}_t to be emitted at time step t , and $\mathbf{v}_a^{\text{src}}$, $\mathbf{U}_a^{\text{src}}$ and $\mathbf{W}_a^{\text{src}}$ are model parameters.

Finally, a time-dependent source context vector c_t is computed as a weighted sum over the source annotation vectors, where each vector is weighted by the attention weight $\alpha_{t,i}^{\text{src}}$, as in eq. (4):

$$\mathbf{c}_t = \sum_{i=1}^N \alpha_{t,i}^{\text{src}} \mathbf{h}_i. \quad (4)$$

¹<https://github.com/nyu-dl/dl4mt-tutorial/blob/master/docs/cgru.pdf>.

²ATT_{src} is named ATT in the original technical report.

$$p(y_t = k | \mathbf{y}_{<t}, \mathbf{c}_t) \propto \exp(\mathbf{L}_o \tanh(\mathbf{L}_s \mathbf{s}_t + \mathbf{L}_w \mathbf{E}_y[\hat{y}_{t-1}] + \mathbf{L}_c \mathbf{c}_t)). \quad (1)$$




Figure 1: An illustration of the conditional GRU: the steps taken to compute the current hidden state s_t from the previous state s_{t-1} , the previously emitted word \hat{y}_{t-1} , and the source annotation vectors C , including the candidate hidden state s'_t and the source-language attention vector c_t .

3 Multi-modal NMT

Our MNMT model can be seen as an expansion of the attention-based NMT framework described in §2.1 with the addition of a *visual component* to incorporate spatial visual features.

We use publicly available pre-trained CNNs for image feature extraction. Specifically, we extract spatial image features for all images in our dataset using the 50-layer Residual network (ResNet-50) of He et al. (2015). These spatial features are the activations of the `res4f` layer, which can be seen as encoding an image in a 14×14 grid, where each of the entries in the grid is represented by a 1024D feature vector that only encodes information about that specific region of the image. We vectorise this 3-tensor into a 196×1024 matrix $A = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_L)$, $\mathbf{a}_l \in \mathbb{R}^{1024}$ where each of the $L = 196$ rows consists of a 1024D feature vector and each column, i.e. feature vector, represents one grid in the image.

3.1 NMT_{SRC+IMG}: decoder with two independent attention mechanisms

Model NMT_{SRC+IMG} integrates two separate attention mechanisms over the source-language words and visual features in a single decoder RNN. Our doubly-attentive decoder RNN is conditioned on

the previous hidden state of the decoder and the previously emitted word, as well as the source sentence and the image via two independent attention mechanisms, as illustrated in Figure 2.

We implement this idea expanding the conditional GRU described in §2.2 onto a *doubly-conditional* GRU. To that end, in addition to the source-language attention, we introduce a new attention mechanism ATT_{img} to the original conditional GRU proposal. This visual attention computes a *time-dependent* image context vector $\hat{\mathbf{i}}_t$ given a hidden state proposal s'_t and the image annotation vectors $A = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_L)$ using the “soft” attention (Xu et al., 2015).

This attention mechanism is very similar to the source-language attention with the addition of a *gating scalar*, explained further below. First, a single-layer feed-forward network is used to compute an *expected alignment* $e_{t,l}^{\text{img}}$ between each image annotation vector \mathbf{a}_l and the target word to be emitted at the current time step t , as in eqs. (5) and (6):

$$e_{t,l}^{\text{img}} = (\mathbf{v}_a^{\text{img}})^T \tanh(\mathbf{U}_a^{\text{img}} s'_t + \mathbf{W}_a^{\text{img}} \mathbf{a}_l), \quad (5)$$

$$\alpha_{t,l}^{\text{img}} = \frac{\exp(e_{t,l}^{\text{img}})}{\sum_{j=1}^L \exp(e_{t,j}^{\text{img}})}, \quad (6)$$

where $\alpha_{t,l}^{\text{img}}$ is the normalised alignment matrix between all the image patches \mathbf{a}_l and the target word to be emitted at time step t , and $\mathbf{v}_a^{\text{img}}$, $\mathbf{U}_a^{\text{img}}$ and $\mathbf{W}_a^{\text{img}}$ are model parameters. Note that Equations (2) and (3), that compute the expected source alignment $e_{t,i}^{\text{src}}$ and the weight matrices $\alpha_{t,i}^{\text{src}}$, and eqs. (5) and (6) that compute the expected image alignment $e_{t,l}^{\text{img}}$ and the weight matrices $\alpha_{t,l}^{\text{img}}$, both compute similar statistics over the source and image annotations, respectively.

In eq. (7) we compute $\beta_t \in [0, 1]$, a gating scalar used to weight the expected importance of the image context vector in relation to the next target word at time step t :

$$\beta_t = \sigma(\mathbf{W}_\beta \mathbf{s}_{t-1} + \mathbf{b}_\beta), \quad (7)$$

where \mathbf{W}_β , \mathbf{b}_β are model parameters. It is in turn used to compute the time-dependent image context vector $\hat{\mathbf{i}}_t$ for the current decoder time step t , as in eq. (8):

$$\hat{\mathbf{i}}_t = \beta_t \sum_{l=1}^L \alpha_{t,l}^{\text{img}} \mathbf{a}_l. \quad (8)$$

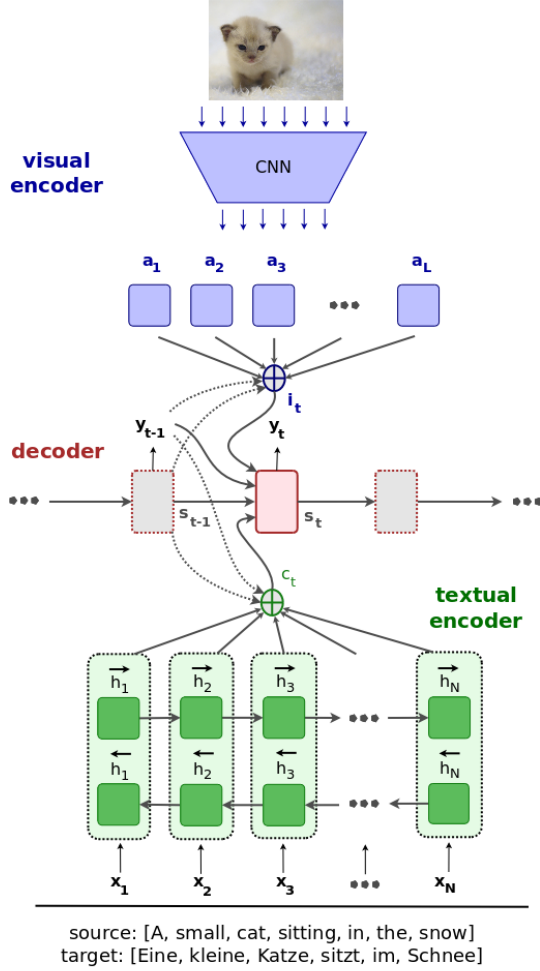


Figure 2: A doubly-attentive decoder learns to attend to image patches and source-language words independently when generating translations.

The only difference between Equations (4) (source context vector) and (8) (image context vector) is that the latter uses a gating scalar, whereas the former does not. We use β following Xu et al. (2015) who empirically found it to improve the variability of the image descriptions generated with their model.

Finally, we use the time-dependent image context vector i_t as an additional input to a modified version of REC_2 (§2.2), which now computes the final hidden state s_t using the hidden state proposal s'_t , and the time-dependent source and image context vectors c_t and i_t , as in eq. (9):

$$\begin{aligned}
 z_t &= \sigma(\mathbf{W}_z^{\text{src}} c_t + \mathbf{W}_z^{\text{img}} i_t + U_z s'_t), \\
 r_t &= \sigma(\mathbf{W}_r^{\text{src}} c_t + \mathbf{W}_r^{\text{img}} i_t + U_r s'_t), \\
 \underline{s}_t &= \tanh(\mathbf{W}^{\text{src}} c_t + \mathbf{W}^{\text{img}} i_t + r_t \odot (U s'_t)), \\
 s_t &= (1 - z_t) \odot \underline{s}_t + z_t \odot s'_t.
 \end{aligned} \tag{9}$$

In Equation (10), the probabilities for the next target word are computed using the new multi-modal hidden state s_t , the previously emitted word \hat{y}_{t-1} , and the two context vectors c_t and i_t , where L_o , L_s , L_w , L_{cs} and L_{ci} are projection matrices and trained with the model.

4 Data

The Flickr30k data set contains 30k images and 5 descriptions in English for each image (Young et al., 2014). In this work, we use the Multi30k dataset (Elliott et al., 2016), which consists of two multilingual expansions of the original Flickr30k: one with translated data and another one with comparable data, henceforth referred to as M30k_T and M30k_C , respectively.

For each of the 30k images in the Flickr30k, the M30k_T has one of the English descriptions manually translated into German by a professional translator. Training, validation and test sets contain 29k, 1,014 and 1k images respectively, each accompanied by a sentence pair (the original English sentence and its translation into German). For each of the 30k images in the Flickr30k, the M30k_C has five descriptions in German collected independently from the English descriptions. Training, validation and test sets contain 29k, 1,014 and 1k images respectively, each accompanied by five sentences in English and five sentences in German.

We use the entire M30k_T training set for training our MNMT models, its validation set for model selection with BLEU (Papineni et al., 2002), and its test set for evaluation. In addition, since the amount of training data available is small, we build a back-translation model using the text-only NMT model described in §2.1 trained on the Multi30k_T data set (German→English and English→German), without images. We use this model to back-translate the 145k German (English) descriptions in the Multi30k_C into English (German) and include the triples (synthetic English description, German description, image) when translating into German, and the triples (synthetic German description, English description, image) when translating into English, as additional training data (Sennrich et al., 2016a).

We also use the WMT 2015 text-only parallel corpora available for the English–German language pair, consisting of about 4.3M sentence pairs (Bojar et al., 2015). These include the Eu-

$$p(y_t = k \mid \mathbf{y}_{<t}, C, A) \propto \exp(\mathbf{L}_o \tanh(\mathbf{L}_s \mathbf{s}_t + \mathbf{L}_w \mathbf{E}_y[\hat{y}_{t-1}] + \mathbf{L}_{cs} \mathbf{c}_t + \mathbf{L}_{ci} \mathbf{i}_t)). \quad (10)$$

roparl v7 (Koehn, 2005), News Commentary and Common Crawl corpora, which are concatenated and used for pre-training.

We use the scripts in the Moses SMT Toolkit (Koehn et al., 2007) to normalise 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 83,093 English and 91,141 German subword tokens. If sentences in English or German are longer than 80 tokens, they are discarded. We train models to translate from English into German, as well as for German into English, and report evaluation of cased, tokenized sentences with punctuation.

5 Experimental setup

Our encoder is a bidirectional RNN with GRU, one 1024D single-layer forward and one 1024D single-layer backward RNN. Source and target word embeddings are 620D each and trained jointly with the model. Word embeddings and other non-recurrent matrices are initialised by sampling from a Gaussian $\mathcal{N}(0, 0.01^2)$, recurrent matrices are random orthogonal and bias vectors are all initialised to zero.

Visual features are obtained by feeding images to the pre-trained ResNet-50 and using the activations of the `res4f` layer (He et al., 2015). We apply dropout with a probability of 0.5 in the encoder bidirectional RNN, the image features, the decoder RNN and before emitting a target word. We follow Gal and Ghahramani (2016) and apply dropout to the encoder bidirectional and the decoder RNN using one same mask in all time steps.

All models are trained using stochastic gradient descent with ADADELTA (Zeiler, 2012) with minibatches of size 80 (text-only NMT) or 40 (MNMT), where each training instance consists of one English sentence, one German sentence and one image (MNMT). We apply early stopping for model selection based on BLEU4, so that if a model does not improve on BLEU4 in the validation set for more than 20 epochs, training is halted.

The translation quality of our models is evaluated quantitatively in terms of BLEU4 (Papineni et al., 2002), METEOR (Denkowski and Lavie, 2014), TER (Snover et al., 2006), and

chrF3 (Popović, 2015).³ We report statistical significance with approximate randomisation for the first three metrics with MultEval (Clark et al., 2011).

5.1 Baselines

We train a text-only phrase-based SMT (PBSMT) system and a text-only NMT model for comparison (English→German and German→English). Our PBSMT baseline is built with Moses and uses a 5-gram LM with modified Kneser-Ney smoothing (Kneser and Ney, 1995). It is trained on the English→German (German→English) descriptions of the M30k_T, whereas its LM is trained on the German (English) descriptions only. We use minimum error rate training to tune the model with BLEU (Och, 2003). The text-only NMT baseline is the one described in §2.1 and is trained on the M30k_T’s English–German descriptions, again in both language directions.

When translating into German, we also compare our model against two publicly available results obtained with multi-modal attention-based NMT models. The first model is Huang et al. (2016)’s best model trained on the same data, and the second is their best model using additional object detections, respectively models `m1` (image at head) and `m3` in the authors’ paper.

5.2 Results

In Table 1, we show results for the two text-only baselines NMT and PBSMT, the multi-modal models of Huang et al. (2016), and our MNMT models trained on the M30k_T and pre-trained on the in-domain back-translated M30k_C and the general-domain text-only English-German MT corpora from WMT 2015. All models are trained to translate from English into German.

Training on M30k_T One main finding is that our model consistently outperforms the comparable model of Huang et al. (2016) when translating into German, with improvements of +1.4 BLEU and +2.7 METEOR. In fact, even when their model has access to more data our model still improves by +0.9 METEOR.

Moreover, we can also conclude from Table 1 that PBSMT performs better at recall-oriented

³We specifically compute character 6-gram F3, and additionally character precision and recall for comparison.

English→German					
Model	Training data	BLEU4↑	METEOR↑	TER↓	chrF3↑ (prec. / recall)
NMT	M30k _T	33.7	52.3	46.7	65.2 (67.7 / 65.0)
PBSMT	M30k _T	32.9	<u>54.3</u> [†]	<u>45.1</u> [†]	67.4 (66.5 / 67.5)
Huang et al. (2016)	M30k _T	35.1 (↑ 1.4)	52.2 (↓ 2.1)	—	—
	+ RCNN	36.5 (↑ 2.8)	54.1 (↓ 0.2)	—	—
NMT _{SRC+IMG}	M30k _T	36.5 ^{†‡}	55.0 [†]	43.7 ^{†‡}	67.3 (66.8 / 67.4)
Improvements					
NMT _{SRC+IMG} vs. NMT		↑ 2.8	↑ 2.7	↓ 3.0	↑ 2.1 ↓ 0.9 / ↑ 2.4
NMT _{SRC+IMG} vs. PBSMT		↑ 3.6	↑ 0.7	↓ 1.4	↓ 0.1 ↑ 0.3 / ↓ 0.1
NMT _{SRC+IMG} vs. Huang		↑ 1.4	↑ 2.8	—	—
NMT _{SRC+IMG} vs. Huang (+RCNN)		↑ 0.0	↑ 0.9	—	—
Pre-training data set: back-translated M30k _C (in-domain)					
PBSMT (LM)	M30k _T	34.0	55.0 [†]	44.7	68.0 (66.8 / 68.1)
NMT	M30k _T	<u>35.5</u> [‡]	53.4	<u>43.3</u> [‡]	65.2 (67.7 / 65.0)
NMT _{SRC+IMG}	M30k _T	37.1 ^{†‡}	54.5 [†]	42.8 ^{†‡}	66.6 (67.2 / 66.5)
NMT _{SRC+IMG} vs. best PBSMT		↑ 3.1	↓ 0.5	↓ 1.9	↓ 1.4 ↑ 0.4 / ↓ 1.6
NMT _{SRC+IMG} vs. NMT		↑ 1.6	↑ 1.1	↓ 0.5	↑ 1.4 ↓ 0.5 / ↑ 1.5
Pre-training data set: WMT'15 English-German corpora (general domain)					
PBSMT (concat)	M30k _T	32.6	53.9	46.1	67.3 (66.3 / 67.4)
PBSMT (LM)	M30k _T	32.5	54.1	46.0	67.3 (66.0 / 67.4)
NMT	M30k _T	37.8	<u>56.7</u>	<u>41.0</u>	<u>69.2</u> (69.7 / 69.1)
NMT _{SRC+IMG}	M30k _T	39.0 ^{†‡}	56.8 [‡]	40.6 [‡]	69.6 (69.6 / 69.6)
NMT _{SRC+IMG} vs. best PBSMT		↑ 6.4	↑ 2.7	↓ 5.4	↑ 2.3 ↑ 3.3 / ↑ 2.2
NMT _{SRC+IMG} vs. NMT		↑ 1.2	↑ 0.1	↓ 0.4	↑ 0.4 ↓ 0.1 / ↑ 0.5

Table 1: BLEU4, METEOR, chrF3, character-level precision and recall (higher is better) and TER scores (lower is better) on the translated Multi30k (M30k_T) test set. Best text-only baselines results are underlined and best overall results appear in bold. We show Huang et al. (2016)’s improvements over the best text-only baseline in parentheses. Results are significantly better than the NMT baseline ([†]) and the SMT baseline ([‡]) with $p < 0.01$ (no pre-training) or $p < 0.05$ (when pre-training either on the back-translated M30k_C or WMT’15 corpora).

metrics, i.e. METEOR and chrF3, whereas NMT is better at precision-oriented ones, i.e. BLEU4. This is somehow expected, since the attention mechanism in NMT (Bahdanau et al., 2015) does not explicitly take attention weights from previous time steps into account, and thus lacks the notion of source coverage as in SMT (Koehn et al., 2003; Tu et al., 2016). We note that these ideas are complementary and incorporating coverage into model NMT_{SRC+IMG} could lead to more improvements, especially in recall-oriented metrics. Nonetheless, our doubly-attentive model shows consistent gains in both precision- and recall-oriented metrics in comparison to the text-only NMT baseline, i.e. it is significantly better according to BLEU4, METEOR and TER ($p < 0.01$), and it also improves chrF3 by +2.1. In comparison to the PBSMT baseline, our proposed model still significantly improves according to both BLEU4 and TER ($p <$

0.01), also increasing METEOR by +0.7 but with an associated p -value of $p = 0.071$, therefore not significant for $p < 0.05$. Although chrF3 is the only metric in which the PBSMT model scores best, the difference between our model and the latter is only 0.1, meaning that they are practically equivalent. We note that model NMT_{SRC+IMG} consistently increases character recall in comparison to the text-only NMT baseline. Although it can happen at the expense of character precision, gains in recall are always much higher than any eventual loss in precision, leading to consistent improvements in chrF3.

In Table 2, we observe that when translating into English and training on the original M30k_T, model NMT_{SRC+IMG} outperforms both baselines by a large margin, according to all four metrics evaluated. We also note that both model NMT_{SRC+IMG}’s character-level precision and re-

German→English				
Model	BLEU4 \uparrow	METEOR \uparrow	TER \downarrow	chrF3 \uparrow
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>
NMT _{SRC+IMG}	40.6\dagger	37.5\dagger	37.7\dagger	65.2
Improvements				
Ours vs. NMT	\uparrow 2.4	\uparrow 1.7	\downarrow 2.5	\uparrow 2.4
Ours vs. PBSMT	\uparrow 7.8	\uparrow 2.7	\downarrow 6.2	\uparrow 3.4
Pre-training data set: back-translated M30k _C (in-domain)				
PBSMT	36.8	36.4	40.8	64.5
NMT	<u>42.6</u>	<u>38.9</u>	<u>36.1</u>	<u>67.6</u>
NMT _{SRC+IMG}	43.2\ddagger	39.0\ddagger	35.5\ddagger	67.7
Improvements				
Ours vs. PBSMT	\uparrow 6.4	\uparrow 2.6	\downarrow 5.3	\uparrow 3.2
Ours vs. NMT	\uparrow 0.6	\uparrow 0.1	\downarrow 0.6	\uparrow 0.1

Table 2: BLEU4, METEOR, chrF3 (higher is better), and TER scores (lower is better) on the translated Multi30k (M30k_T) test set. Best text-only baselines results are underlined and best overall results appear in bold. Results are significantly better than the NMT baseline (\dagger) and the SMT baseline (\ddagger) with $p < 0.01$.

call are higher than those of the two baselines, in contrast to results obtained when translating from English into German. This suggests that model NMT_{SRC+IMG} might better integrate the image features when translating into an “easier” language, i.e. a language with less morphology, although experiments involving more language pairs are necessary to confirm whether this is indeed the case.

Pre-training We now discuss results for models pre-trained using different data sets. We first pre-trained the two text-only baselines PBSMT and NMT, and our MNMT model on the back-translated M30k_C, a medium-sized in-domain image description data set (145k training instances), in both directions. We also pre-trained the same models on the English–German parallel sentences of much larger MT data sets, i.e. the concatenation of the Europarl (Koehn, 2005), Common Crawl and News Commentary corpora, used in WMT 2015 (\sim 4.3M parallel sentences). Model PBSMT (concat.) used the concatenation of the pre-training and training data for training, and model PBSMT (LM) used the general-domain German sentences as additional data to train the LM. From Tables 1 and 2, it is clear that model NMT_{SRC+IMG} can learn from both in-domain, multi-modal pre-training data sets as well as text-only, general domain ones.

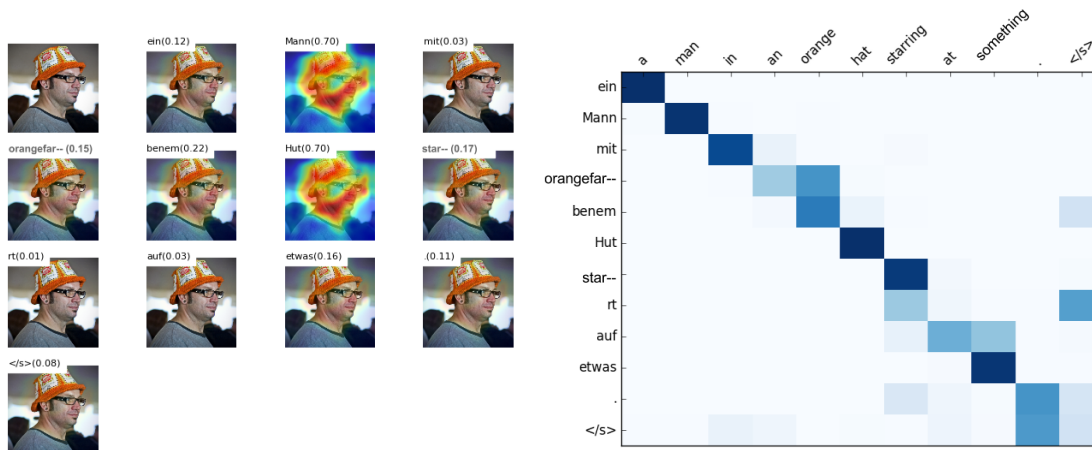
Pre-training on M30k_C When pre-training on the back-translated M30k_C and translating into German, the recall-oriented chrF3 shows a difference of 1.4 points between PBSMT and our model, mostly due to character recall; nonetheless, our model still improved by the same margin on the text-only NMT baseline. Our model still outperforms the PBSMT baseline according to BLEU4 and TER, and the text-only NMT baseline according to all metrics ($p < .05$).

When translating into English, model NMT_{SRC+IMG} still consistently scores higher according to all metrics evaluated, although the differences between its translations and those obtained with the NMT baseline are no longer statistically significant ($p < 0.01$).

Pre-training on WMT 2015 corpora We also pre-trained our English–German models on the WMT 2015 corpora, which took 10 days, i.e. \sim 6–7 epochs. Results show that model NMT_{SRC+IMG} improves significantly over the NMT baseline according to BLEU4, and is consistently better than the PBSMT baseline according to all four metrics.⁴ This is a strong indication that model NMT_{SRC+IMG} can exploit the additional pre-training data efficiently, both general- and in-domain. While the PBSMT model is still competitive when using additional in-domain data—according to METEOR and chrF3—the same cannot be said when using general-domain pre-training corpora. From our experiments, NMT models in general, and especially model NMT_{SRC+IMG}, thrive when training and test domains are mixed, which is a very common real-world scenario.

Textual and visual attention In Figure 3, we visualise the visual and textual attention weights for an entry of the M30k_T test set. In the visual attention, the β gate (written in parentheses after each word) caused the image features to be used mostly to generate the words *Mann* (man) and *Hut* (hat), two highly visual terms in the sentence. We observe that in general visually grounded terms, e.g. Mann and Hut, usually have a high associated β value, whereas other less visual terms like *mit* (with) or *auf* (at) do not. That causes the model to use the image features when it is describing a visual concept in the sentence, which is an interest-

⁴In order for PBSMT models to remain competitive, we believe more advanced data selection techniques are needed, which are out of the scope of this work.



(a) Image–target word alignments.

(b) Source–target word alignments.

Figure 3: Visualisation of image– and source–target word alignments for the M30k_T test set.

ing feature of our model. Interestingly, our model is very selective when choosing to use image features: it only assigned $\beta > 0.5$ for 20% of the outputted target words, and $\beta > 0.8$ to only 8%. A manual inspection of translations shows that these words are mostly concrete nouns with a strong visual appeal.

Lastly, using two independent attention mechanisms is a good compromise between model compactness and flexibility. While the attention-based NMT model baseline has $\sim 200\text{M}$ parameters, model NMT_{SRC+IMG} has $\sim 213\text{M}$, thus using just $\sim 6.6\%$ more parameters than the latter.

6 Related work

Multi-modal MT was just recently addressed by the MT community by means of a shared task (Specia et al., 2016). However, there has been a considerable amount of work on natural language generation from non-textual inputs. Mao et al. (2014) introduced a multi-modal RNN that integrates text and visual features and applied it to the tasks of image description generation and image–sentence ranking. In their work, the authors incorporate global image features in a separate multi-modal layer that merges the RNN textual representations and the global image features. Vinyals et al. (2015) proposed an influential neural IDG model based on the sequence-to-sequence framework, which is trained end-to-end. Elliott et al. (2015) put forward a model to generate multilingual descriptions of images by learning and transferring features between two in-

dependent, non-attentive neural image description models.⁵ Venugopalan et al. (2015) introduced a model trained end-to-end to generate textual descriptions of open-domain videos from the video frames based on the sequence-to-sequence framework. Finally, Xu et al. (2015) introduced the first attention-based IDG model where an attentive decoder learns to attend to different parts of an image as it generates its description in natural language.

In the context of NMT, Zoph and Knight (2016) introduced a multi-source attention-based NMT model trained to translate a pair of sentences in two different source languages into a target language, and reported considerable improvements over a single-source baseline. Dong et al. (2015) proposed a multi-task learning approach where a model is trained to translate from one source language into multiple target languages. Firat et al. (2016) put forward a multi-way model trained to translate between many different source and target languages. Instead of one attention mechanism per language pair as in Dong et al. (2015), which would lead to a quadratic number of attention mechanisms in relation to language pairs, they use a shared attention mechanism where each target language has one attention shared by all source languages. Luong et al. (2016) proposed a multi-task approach where they train a model using two tasks and a shared decoder: the main task is to translate from German into English and the sec-

⁵Although their model has not been devised with translation as its primary goal, theirs is one of the baselines of the first shared task in multi-modal MT in WMT 2016 (Specia et al., 2016).

ondary task is to generate English image descriptions. They show improvements in the main translation task when also training for the secondary image description task. Although not an NMT model, Hitschler et al. (2016) recently used image features to re-rank translations of image descriptions generated by an SMT model and reported significant improvements.

Although no purely neural multi-modal model to date significantly improves on both text-only NMT and SMT models (Specia et al., 2016), different research groups have proposed to include global and spatial visual features in re-ranking n -best lists generated by an SMT system or directly in an NMT framework with some success (Caglayan et al., 2016a; Calixto et al., 2016; Huang et al., 2016; Libovický et al., 2016; Shah et al., 2016). To the best of our knowledge, the best published results of a purely MNMT model are those of Huang et al. (2016), who proposed to use global visual features extracted with the VGG19 network (Simonyan and Zisserman, 2015) for an entire image, and also for regions of the image obtained using the RCNN of Girshick et al. (2014). Their best model improves over a strong text-only NMT baseline and is comparable to results obtained with an SMT model trained on the same data. For that reason, their models are used as baselines in our experiments whenever possible.

Our work differs from previous work in that, first, we propose attention-based MNMT models. This is an important difference since the use of attention in NMT has become standard and is the current state-of-the-art (Jean et al., 2015; Luong et al., 2015; Firat et al., 2016; Sennrich et al., 2016b). Second, we propose a *doubly-attentive model* where we effectively fuse two mono-modal attention mechanisms into one multi-modal decoder, training the entire model jointly and end-to-end. Additionally, we are interested in how to merge textual and visual representations into multi-modal representations when generating words in the target language, which differs substantially from text-only translation tasks even when these translate from many source languages and/or into many target languages (Dong et al., 2015; Firat et al., 2016; Zoph and Knight, 2016). To the best of our knowledge, we are among the first⁶ to integrate multi-modal inputs in NMT via

⁶As pointed out by an anonymous reviewer, Caglayan et al. (2016b) have also experimented with attention-based

independent attention mechanisms.

Applications Initial experiments with model $NMT_{SRC+IMG}$ have been reported in Calixto et al. (2016). Additionally, $NMT_{SRC+IMG}$ has been applied to the machine translation of user-generated product listings from an e-commerce website, while also making use of the product images to improve translations (Calixto et al., 2017b,a).

7 Conclusions and Future Work

We have introduced a novel attention-based, multi-modal NMT model to incorporate spatial visual information into NMT. We have reported state-of-the-art results on the M30k_T test set, improving on previous multi-modal attention-based models. We have also showed that our model can be efficiently pre-trained on both medium-sized back-translated in-domain multi-modal data as well as also large general-domain text-only MT corpora, finding that it is able to exploit the additional data regardless of the domain. Our model also compares favourably to both NMT and PB-SMT baselines evaluated on the same training data. In the future, we will incorporate coverage into our model and study how to apply it to other Natural Language Processing tasks.

Acknowledgements

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). The authors would like to thank Chris Hokamp, Peyman Passban, and Dasha Bogdanova for insightful discussions at early stages of this work, Andy Way for proof-reading and providing many good suggestions of improvements, as well as our anonymous reviewers for their valuable comments and feedback.

Reproducibility

Code and pre-trained models for this paper are available at https://github.com/iacercalixto/nmt_doubly_attentive.

multi-modal NMT.

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.
- 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>.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. 2015. Findings of the 2015 workshop on statistical machine translation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*. Lisbon, Portugal, pages 1–46. <http://aclweb.org/anthology/W15-3001>.
- Ozan Caglayan, Walid Aransa, Yaxing Wang, Marc Masana, Mercedes García-Martínez, Fethi Bougares, Loïc Barrault, and Joost van de Weijer. 2016a. 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>.
- Ozan Caglayan, Loïc Barrault, and Fethi Bougares. 2016b. Multimodal Attention for Neural Machine Translation. *CoRR* abs/1609.03976. <http://arxiv.org/abs/1609.03976>.
- 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, Daniel Stein, Evgeny Matusov, Sheila Castilho, and Andy Way. 2017a. Human Evaluation of Multi-modal Neural Machine Translation: A Case-Study on E-Commerce Listing Titles. In *Proceedings of the Sixth Workshop on Vision and Language*. Valencia, Spain, pages 31–37. <http://www.aclweb.org/anthology/W17-2004>.
- Iacer Calixto, Daniel Stein, Evgeny Matusov, Pintu Lohar, Sheila Castilho, and Andy Way. 2017b. Using Images to Improve Machine-Translating E-Commerce Product Listings. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Valencia, Spain, pages 637–643. <http://www.aclweb.org/anthology/E17-2101>.
- 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*.
- 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, 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>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.
- Julian Hitschler, Shigehiko Schamoni, and Stefan Riezler. 2016. Multimodal Pivots for Image Caption Translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany, pages 2399–2409. <http://www.aclweb.org/anthology/P16-1227>.
- 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>.
- Sébastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On Using Very Large Target Vocabulary for Neural Machine 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 1–10. <http://www.aclweb.org/anthology/P15-1001>.
- 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, US., pages 1700–1709.
- 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. 2005. Europarl: A Parallel Corpus for Statistical Machine Translation. In *Conference Proceedings: the tenth Machine Translation Summit*. AAMT, AAMT, Phuket, Thailand, pages 79–86. <http://mt-archive.info/MTS-2005-Koehn.pdf>.
- 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*. Prague, Czech Republic, ACL '07, pages 177–180. <http://dl.acm.org/citation.cfm?id=1557769.1557821>.
- Philipp Koehn, Franz Josef Och, and Daniel Marcu. 2003. Statistical Phrase-based Translation. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - Volume 1*. Edmonton, Canada, NAACL '03, pages 48–54. <https://doi.org/10.3115/1073445.1073462>.
- 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>.
- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. Multi-Task Sequence to Sequence Learning. In *Proceedings of the International Conference on Learning Representations (ICLR), 2016*. San Juan, Puerto Rico.
- 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>.
- 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>.
- 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. 2015. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations (ICLR)*. San Diego, CA.
- 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, collocated with ACL 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>.
- Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond J. Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to sequence - video to text. In *2015 IEEE International Conference on Computer Vision, ICCV 2015*. Santiago, Chile, pages 4534–4542. <https://doi.org/10.1109/ICCV.2015.515>.
- 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>.
- Barret Zoph and Kevin Knight. 2016. Multi-Source Neural Translation. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California, pages 30–34. <http://www.aclweb.org/anthology/N16-1004>.