

# Visual Agreement Regularized Training for Multi-Modal Machine Translation

Pengcheng Yang,<sup>1,2</sup> Boxing Chen,<sup>3</sup> Pei Zhang,<sup>3</sup> Xu Sun<sup>1,2</sup>

<sup>1</sup>Center for Data Science, Peking University

<sup>2</sup>MOE Key Lab of Computational Linguistics, School of EECS, Peking University

<sup>3</sup>Alibaba DAMO Academy, Hangzhou, China

{yang\_pc, xusun}@pku.edu.cn, {boxing.cbx, xiaoyi.zp}@alibaba-inc.com

## Abstract

Multi-modal machine translation aims at translating the source sentence into a different language in the presence of the paired image. Previous work suggests that additional visual information only provides dispensable help to translation, which is needed in several very special cases such as translating ambiguous words. To make better use of visual information, this work presents visual agreement regularized training. The proposed approach jointly trains the source-to-target and target-to-source translation models and encourages them to share the same focus on the visual information when generating semantically equivalent visual words (e.g. “ball” in English and “ballon” in French). Besides, a simple yet effective multi-head co-attention model is also introduced to capture interactions between visual and textual features. The results show that our approaches can outperform competitive baselines by a large margin on the Multi30k dataset. Further analysis demonstrates that the proposed regularized training can effectively improve the agreement of attention on the image, leading to better use of visual information.

## Introduction

As real-scenarios integrating multiple modal information have become commonplace, an increasing of endeavors have been paid to multi-modal machine translation (MMT). Different from tradition machine translation based solely on textual information, the MMT task aims at translating sentences paired with images into a different target language (Elliott, Frank, and Specia 2016). It not only has plenty of practical applications, e.g., helping improve the translation of ambiguous multi-sense words, but also serves as an ideal testbed for cross-modal text generation (Calixto, Liu, and Campbell 2017).

Despite its importance described above, previous work on MMT typically suffers from two drawbacks. First of all, how to effectively integrate visual information still remains an intractable challenge. Previous work has shown that additional visual information only provides dispensable help, which may be needed in the presence of some special cases (e.g. translating incorrect or ambiguous source words or gener-

FR: footballeur court après le **ballon** EN: soccer player chases the **ball**




Figure 1: A toy example of visual agreement regularized training.  $\vec{\theta}$  denotes forward model translating English (EN) into French (FR) and  $\overleftarrow{\theta}$  is the opposite. When  $\vec{\theta}$  generates “ballon” and  $\overleftarrow{\theta}$  generates “ball”, they should all focus on the region near the football in the image.

ating gender-neutral target words). As a result, most existing approaches tend to ignore such visual information. To remedy this, we present visual agreement regularized training for the MMT task in this work. The proposed training schema jointly trains both source-to-target (forward) and target-to-source (backward) translation models, and encourages them to share the same focus on the image when generating semantic equivalent visual words (e.g. “ball” in English and “ballon” in French). Take Figure 1 as an example, whose translation direction is EN→FR. When generating the target word “ballon”, the forward model is likely to focus on the red box region in the image, on which is also expected to be focused by the backward model when it generates the source word “ball”. By encouraging the agreement of bidirectional (forward and backward) models’ attention (red and blue boxes in Figure 1) to image, the visual information can be more fully utilized.

However, a tricky problem is that semantically equivalent word pairs tend to be unavailable in practice. To address this issue, we present two effective regularization approaches: hard regularization constructing pseudo-aligned

word pairs via external word aligner (Dyer, Chahuneau, and Smith 2013), and soft regularization employing posterior estimation to provide an approximation. Considering that the generation of non-visual words (e.g., “the” and “le” in Figure 1) requires little or no visual information from the image, we further propose an adaptive weighting strategy to adaptively adjust the weight of regularization item based on the dependence of the word to be generated on the image.

The other challenge is how to capture interactions between visual and textual features. Most previous work independently obtains representations of the input image and source sentence, which ignores the complex dependency between two information sources. To tackle this challenge, we present a simple yet effective multi-head co-attention, which can build bidirectional interactions between visual and textual features in multiple subspaces so that two information sources can mutually boost for better representations.

The main contributions of this work are listed as follows:

- In terms of model training, we propose two visual agreement regularized training schemas as well as an adaptive weighting strategy, which encourage bidirectional models to make better use of visual information.
- In terms of model architecture, we introduce a multi-head co-attention model. It aims to capture the interaction between visual and textual features in multiple subspaces so that two information sources can mutually boost for better representations.
- Experimental results show that our approach can outperform competitive baselines by a large margin. Further analysis demonstrates that bidirectional models can more fully utilize visual information by improving the agreement of visual attention.

## Proposed Model

Given the visual features  $\mathbf{v} = (v_1, \dots, v_m)$  and source sentence  $\mathbf{x} = (x_1, \dots, x_n)$ , the MMT task aims to generate the target sentence  $\mathbf{y} = (y_1, \dots, y_l)$ . We first illustrate our approach with Seq2Seq model as base architecture and further extend it to Transformer. Figure 2 presents the sketch of our proposed model, which is elaborated on as follows.

### Visual Encoder and Textual Encoder

The textual encoder is implemented as a recurrent neural network (RNN), which encodes the source sentence  $\mathbf{x}$  into a sequence of hidden states. Formally, the hidden representation of each word  $x_i$  is computed as  $h_i^x = \text{RNN}(h_{i-1}^x, e(x_i))$ , where  $e(x_i)$  denotes the embedding of  $x_i$ . The final textual representation matrix is denoted as  $\mathbf{X} = \{h_1^x, \dots, h_n^x\} \in \mathbb{R}^{n \times d}$ , where  $n$  is the total number of textual representations and  $d$  is the dimension of  $h_i^x$ .

The visual features  $\mathbf{v}$  is pooled ROI features extracted by Faster R-CNN (Ren et al. 2015). The final visual representation matrix is denoted as  $\mathbf{V} = \{h_1^v, \dots, h_m^v\} \in \mathbb{R}^{m \times d}$ , where  $m$  is the total number of visual features.<sup>1</sup>

<sup>1</sup>Here we assume that  $h_i^v$  and  $h_i^x$  share the same dimension. Otherwise, a linear transformation can be introduced to ensure that the dimensions of both are the same.




Figure 2: The illustration of our proposed multi-head co-attention model.

### Multi-Head Co-Attention

To effectively capture the complex dependency between visual and textual features, here we present a simple yet effective multi-head co-attention model, which aims to build bidirectional interactions between two information sources in multiple subspaces. In the  $k$ -th subspace, following Lu et al. (2016), we first connect visual representations  $\mathbf{V}$  and textual representations  $\mathbf{X}$  by computing the similarity matrix  $\mathbf{S}_k \in \mathbb{R}^{m \times n}$  between them as follows:

$$\mathbf{S}_k = \mathbf{V}\mathbf{M}_k\mathbf{X}^T \quad (1)$$

where  $\mathbf{M}_k \in \mathbb{R}^{d \times d}$  is a trainable parameter matrix. The  $(i, j)$  element of  $\mathbf{S}_k$  in Eq. (1) denotes the similarity between the  $i$ -th visual feature and the  $j$ -th textual feature. It is normalized row-wise to produce the image-to-text attention weights  $\mathbf{A}_k^x$ , and column-wise to produce the text-to-image attention weights  $\mathbf{A}_k^v$ :

$$\mathbf{A}_k^x = \text{softmax}(\mathbf{S}_k), \quad \mathbf{A}_k^v = \text{softmax}(\mathbf{S}_k^T) \quad (2)$$

where  $\text{softmax}(\cdot)$  means row-wise normalization. Further, we can obtain the image-aware textual representations  $\hat{\mathbf{X}}_k \in \mathbb{R}^{m \times d}$  and text-aware visual representations  $\hat{\mathbf{V}}_k \in \mathbb{R}^{n \times d}$  by multiplying attention weights with corresponding representations:

$$\hat{\mathbf{X}}_k = \mathbf{A}_k^x\mathbf{X}, \quad \hat{\mathbf{V}}_k = \mathbf{A}_k^v\mathbf{V} \quad (3)$$

In the calculation of  $\hat{\mathbf{X}}_k$  and  $\hat{\mathbf{V}}_k$ ,  $\mathbf{X}$  and  $\mathbf{V}$  mutually guide each other’s attention. Therefore, these two sources of information can mutually boost for better representations.

To capture the interaction between two information sources in multiple subspaces, we set  $K$  different parameter matrices  $\{\mathbf{M}_1, \dots, \mathbf{M}_K\}$ , where  $K$  is a hyper-parameter denoting the number of subspaces. Each  $\mathbf{M}_k$  can yield a set of co-dependent representations  $\hat{\mathbf{X}}_k$  and  $\hat{\mathbf{V}}_k$  in this subspace according to Eq. (1) - Eq. (3). Finally, a mean-pooling layer is used to integrate all co-dependent representations from different subspaces:

$$\hat{\mathbf{X}} = \frac{1}{K} \sum_{k=1}^K \hat{\mathbf{X}}_k, \quad \hat{\mathbf{V}} = \frac{1}{K} \sum_{k=1}^K \hat{\mathbf{V}}_k \quad (4)$$

## Decoder

The decoder implemented as another RNN model is responsible for generating target sentence  $\mathbf{y}$ . The hidden state  $s_t$  of the decoder at time-step  $t$  is computed as:

$$s_t = \text{RNN}(s_{t-1}, [e(y_{t-1}); \bar{c}_t^v; \bar{c}_t^x]) \quad (5)$$

where  $y_{t-1}$  is the word generated at time-step  $t-1$ .  $\bar{c}_t^v$  and  $\bar{c}_t^x$  are time-dependent visual and textual context vectors, respectively. Both  $\bar{c}_t^v$  and  $\bar{c}_t^x$  are obtained via attention mechanism and adaptive fusion, elaborated on as follows.

First, we compute an alignment score  $A_{y_t \rightarrow v_i}$  between each visual representation  $h_i^v$  and the target word  $y_t$  to be generated at the current time-step  $t$  as:

$$A_{y_t \rightarrow v_i} = \frac{\exp(a(s_{t-1}, h_i^v))}{\sum_{j=1}^m \exp(a(s_{t-1}, h_j^v))} \quad (6)$$

where  $a(s_{t-1}, h_i^v)$  is an attention model measuring the dependency between  $s_{t-1}$  and  $h_i^v$ . Readers can refer to Bahdanau, Cho, and Bengio (2014) for the details. Then, the preliminary visual context vector is obtained as:

$$c_t^v = \sum_{i=1}^m A_{y_t \rightarrow v_i} h_i^v \quad (7)$$

The preliminary textual context vector can be computed in a similar way:

$$A_{y_t \rightarrow x_i} = \frac{\exp(a(s_{t-1}, h_i^x))}{\sum_{j=1}^n \exp(a(s_{t-1}, h_j^x))} \quad (8)$$

$$c_t^x = \sum_{i=1}^n A_{y_t \rightarrow x_i} h_i^x \quad (9)$$

$\hat{c}_t^v$  and  $\hat{c}_t^x$  can also be obtained by attending to  $\hat{\mathbf{V}}$  and  $\hat{\mathbf{X}}$  with  $s_{t-1}$  as query, respectively. We compute the final visual vector  $\bar{c}_t^v$  by adaptively fusing  $c_t^v$  and  $\hat{c}_t^v$  as follows:

$$g_t^v = \text{sigmoid}(\mathbf{U}_1 c_t^v + \mathbf{U}_2 \hat{c}_t^v) \quad (10)$$

$$\bar{c}_t^v = g_t^v \odot c_t^v + (1 - g_t^v) \odot \hat{c}_t^v \quad (11)$$

where  $\mathbf{U}_1$  and  $\mathbf{U}_2$  are trainable parameters.  $\odot$  denotes element multiplication. The final textual context vector  $\bar{c}_t^x$  can be obtained by adaptively fusing  $c_t^x$  and  $\hat{c}_t^x$  in a similar way.

## Visual Agreement Regularized Training

In this section, we introduce our proposed visual agreement regularized training, which jointly trains both forward  $P(\mathbf{y}|\mathbf{v}, \mathbf{x}; \vec{\theta})$  and backward  $P(\mathbf{x}|\mathbf{v}, \mathbf{y}; \overleftarrow{\theta})$  translation models on the training corpus  $\mathcal{D} = \{(\mathbf{v}, \mathbf{x}, \mathbf{y})\}$ . Here  $\vec{\theta}$  and  $\overleftarrow{\theta}$  denote corresponding model parameters.

### Overview

Given the instance  $(\mathbf{v}, \mathbf{x}, \mathbf{y})$ , the core idea of visual agreement regularized training is to encourage both forward and backward models to share the same focus on the image when generating semantically equivalent visual words. We use  $A_{y \rightarrow \mathbf{v}}(\vec{\theta}) = (A_{y \rightarrow v_1}, \dots, A_{y \rightarrow v_m})$  to represent the attention vector of the forward model parameterized by  $\vec{\theta}$

on the image when it generates target word  $y$ . Similarly,  $A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta})$  denotes the attention vector of the backward model parameterized by  $\overleftarrow{\theta}$  on the image when it generates source word  $x^*$ , where  $x^* \in \mathbf{x}$  is the true aligned word of  $y$ . Then, the regularized training objective of forward model is defined as:

$$L(\vec{\theta}) = \log P(\mathbf{y}|\mathbf{v}, \mathbf{x}; \vec{\theta}) - \lambda_1 \sum_{y \in \mathbf{y}} \Delta(A_{y \rightarrow \mathbf{v}}(\vec{\theta}), A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta})) \quad (12)$$

where the regularization item  $\Delta(A_{y \rightarrow \mathbf{v}}(\vec{\theta}), A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta}))$  characterizes the difference between two attention vectors  $A_{y \rightarrow \mathbf{v}}(\vec{\theta})$  and  $A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta})$ . We define  $\Delta(\cdot, \cdot)$  as the MSE loss. Similarly, the regularized training objective of the backward model is defined as:

$$L(\overleftarrow{\theta}) = \log P(\mathbf{x}|\mathbf{v}, \mathbf{y}; \overleftarrow{\theta}) - \lambda_2 \sum_{x \in \mathbf{x}} \Delta(A_{x \rightarrow \mathbf{v}}(\overleftarrow{\theta}), A_{y^* \rightarrow \mathbf{v}}(\vec{\theta})) \quad (13)$$

## Visual Agreement Regularization

However, a tricky problem is that true aligned word pair  $(y, x^*)$  or  $(x, y^*)$  tends to be unavailable in practice. To address this issue, here we propose two solutions: hard and soft visual agreement regularization. For simplicity, we illustrate these two approaches based on the forward model.

**Hard visual agreement regularization.** For each training instance  $(\mathbf{v}, \mathbf{x}, \mathbf{y})$ , the hard visual agreement regularization aims to align each target word  $y \in \mathbf{y}$  with a source word  $\hat{x} \in \mathbf{x}$  that has the highest alignment probability. Formally,

$$x^* \approx \hat{x} = \underset{x \in \mathbf{x}}{\text{argmax}} \text{aligner}(x|y) \quad (14)$$

where  $\text{aligner}(\cdot|\cdot)$  denotes the alignment probability that can be obtained by unsupervised word aligner.

**Soft visual agreement regularization.** For some target word  $y$ , there may be multiple source words to be aligned with it. For instance, in Figure 3, the true translation of target word “*footballeur*” is “*soccer player*” consisting of two words. However, the hard regularization can only align one source word with  $y$ . To tackle this problem, here we propose soft visual agreement regularization, which aims at applying posterior estimation to provide an approximation of  $A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta})$ . Formally, we rewrite  $A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta})$  as:

$$A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta}) = \sum_{x \in \mathbf{x}} \mathbb{I}(x = x^*) A_{x \rightarrow \mathbf{v}}(\overleftarrow{\theta}) = \mathbb{E}_{x \sim p^*(x|y)} A_{x \rightarrow \mathbf{v}}(\overleftarrow{\theta}) \quad (15)$$

where  $p^*(x|y)$  denotes the true aligned distribution concentrated<sup>2</sup> to the point  $x^*$ .

For each target word  $y \in \mathbf{y}$ , its probability of being aligned with source word  $x \in \mathbf{x}$  can be characterized by the

<sup>2</sup>It means that  $p^*(x|y) = 1$  holds only when  $x$  is  $x^*$ .

attention weight of forward model to source word  $x$  when generating  $y$ . Therefore, we can use the attention distribution  $A_{y \rightarrow x}$  of target word  $y$  to source sentence  $\mathbf{x}$  to approximate  $p^*(x|y)$  in Eq. (15), yielding the following estimation:

$$\begin{aligned} A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta}) &\approx \mathbb{E}_{x \sim p(x|y, \overrightarrow{\theta})} A_{x \rightarrow \mathbf{v}}(\overleftarrow{\theta}) \\ &= \sum_{x \in \mathbf{x}} p(x|y, \overrightarrow{\theta}) A_{x \rightarrow \mathbf{v}}(\overleftarrow{\theta}) \\ &= \sum_{x \in \mathbf{x}} A_{y \rightarrow x}(\overrightarrow{\theta}) A_{x \rightarrow \mathbf{v}}(\overleftarrow{\theta}) \end{aligned} \quad (16)$$

where  $p(x|y, \overrightarrow{\theta}) = A_{y \rightarrow x}(\overrightarrow{\theta})$  is a posteriori estimation of the true alignment distribution  $p^*(x|y)$ .

### Adaptive Weighting

During the translation process, there usually exist some non-visual words (e.g. “*le*” and “*the*” in Figure 3) whose generation requires little or no visual information from the image, leading to no attention on the input image. Therefore, it is unreasonable to forcibly impose visual agreement regularization on such non-visual words, otherwise it will yield misleading gradient signal. To achieve this, we propose an effective adaptive weighting strategy, which equips the model with the capability of automatically determining whether to introduce visual agreement regularization based on the dependence of the word to be generated on the image.

In more detail, we employ gating mechanism to measure the expected importance of visual context vector  $\overline{c}_t^v$  in relation to the next target word  $y_t$  at time-step  $t$  as follows:

$$\beta_{y_t} = \text{sigmoid}\left(w^T \tanh(\mathbf{W}_1 s_{t-1} + \mathbf{W}_2 \overline{c}_t^v)\right) \quad (17)$$

where  $w, \mathbf{W}_1, \mathbf{W}_2$  are trainable parameters. We replace original visual vector  $\overline{c}_t^v$  in Eq. (5) with the rescaled visual vector  $\beta_{y_t} \overline{c}_t^v$  to update the hidden state of decoder:

$$s_t = \text{RNN}(s_{t-1}, [e(y_{t-1}); \beta_{y_t} \overline{c}_t^v; \overline{c}_t^x]) \quad (18)$$

The gating scalar  $\beta_{y_t}$  characterizes the dependence of the generation of  $y_t$  on visual information. For visual words, the corresponding  $\beta_{y_t}$  presents a larger value. Therefore, the new training objective can be modified to:

$$\begin{aligned} L(\overrightarrow{\theta}) &= \log P(\mathbf{y}|\mathbf{v}, \mathbf{x}; \overrightarrow{\theta}) \\ &\quad - \lambda_1 \sum_{y \in \mathbf{y}} \beta_y \Delta(A_{y \rightarrow \mathbf{v}}(\overrightarrow{\theta}), A_{x^* \rightarrow \mathbf{v}}(\overleftarrow{\theta})) \end{aligned} \quad (19)$$

Different from the constant coefficient of regularization term in Eq. (12), we use  $\beta_y$  to adaptively adjust the weight of regularization in Eq. (19). For visual words,  $\beta_y$  presents a large value so that visual agreement regularization is introduced to make better use of visual information. Conversely, for non-visual words, the value of  $\beta_y$  is small, avoiding the presence of misleading gradient signals caused by unreasonable regularization. Similarly, the training objective of the backward model can be modified to:

$$\begin{aligned} L(\overleftarrow{\theta}) &= \log P(\mathbf{x}|\mathbf{v}, \mathbf{y}; \overleftarrow{\theta}) \\ &\quad - \lambda_2 \sum_{x \in \mathbf{x}} \beta_x \Delta(A_{x \rightarrow \mathbf{v}}(\overleftarrow{\theta}), A_{y^* \rightarrow \mathbf{v}}(\overrightarrow{\theta})) \end{aligned} \quad (20)$$

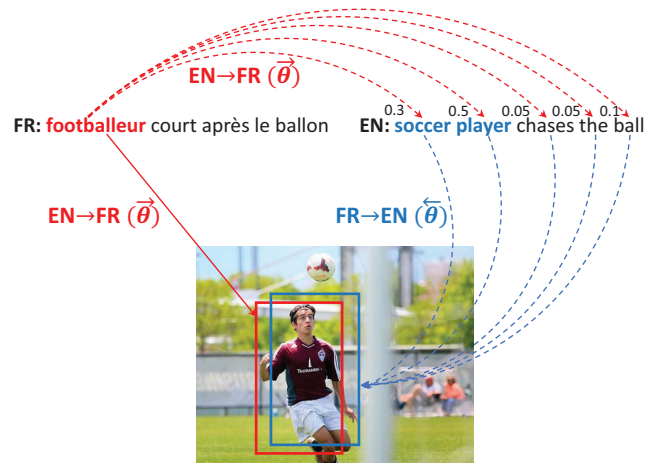


Figure 3: A toy example of soft regularization. The probability value 0.3 that “*footballeur*” is aligned with “*soccer*” can be defined as the attention weight of “*soccer*” when the forward model parameterized by  $\overrightarrow{\theta}$  generates “*footballeur*”.

We also implement our multi-head co-attention model and visual agreement regularized training based on the Transformer (Vaswani et al. 2017). The visual attention vector is calculated as the averaged attention from different heads of the decoder block which performs the best in word alignment. Due to the page limitation and this extension is not the focus of this paper, we will not explain it in more detail. We strongly recommend readers to refer to Vaswani et al. (2017) for more details.

## Experiments

In this section, we introduce the dataset, evaluation metrics, all compared baselines as well as the detailed experiment settings.

### Dataset

Following previous work (Calixto, Liu, and Campbell 2017), we evaluate both our approach and all baselines on the Multi30K dataset (Elliott, Frank, and Specia 2016), which contains 29,000 instances for training and 1,014 for development. We use test-2017 for evaluation, which consists of 1,000 testing instances. Each image is paired with its English (EN) descriptions as well as human translations of German (DE) and French (FR). We use Moses SMT Toolkit to normalize and tokenize all sentences.

### Evaluation Metrics

Following previous work (Elliott and Kádár 2017), we adopt the following automatic evaluation metrics:

- **BLEU**<sup>3</sup> (Papineni et al. 2002) compares the generated text with reference translations by computing overlapping lexical  $n$ -gram units.

<sup>3</sup><https://github.com/moses-smt/mosesdecoder/blob/master/scripts/generic/multi-bleu.perl>

Methods	EN→DE		DE→EN		EN→FR		FR→EN	
	BLEU	METEOR	BLEU	METEOR	BLEU	METEOR	BLEU	METEOR
OnlyText	27.1	49.3	34.2	33.6	51.3	69.1	46.8	40.7
DAMT	27.5	49.4	34.8	34.1	51.2	68.9	47.1	41.1
Imagination	28.2	49.8	35.1	34.3	51.9	69.7	47.4	41.2
Proposal (Hard)	<b>29.3</b>	<b>51.2</b>	35.5	35.1	<b>52.6</b>	<b>69.9</b>	<b>48.9</b>	<b>41.6</b>
Proposal (Soft)	29.2	50.9	<b>35.7</b>	<b>35.2</b>	52.4	69.5	48.7	41.5

Table 1: The experimental results of Seq2Seq-based systems. The best performance is highlighted in bold.

Methods	EN→DE		DE→EN		EN→FR		FR→EN	
	BLEU	METEOR	BLEU	METEOR	BLEU	METEOR	BLEU	METEOR
OnlyText	27.9	49.1	32.2	32.4	52.5	69.9	45.9	40.2
DAMT	28.3	49.7	32.1	32.2	52.2	70.0	46.1	40.1
Imagination	28.4	49.3	32.7	32.9	53.1	70.2	46.8	40.5
Proposal (Hard)	29.3	50.2	<b>33.4</b>	<b>33.5</b>	<b>53.3</b>	<b>70.4</b>	<b>47.7</b>	<b>40.8</b>
Proposal (Soft)	<b>29.5</b>	<b>50.3</b>	32.9	33.1	52.8	70.1	47.3	40.6

Table 2: The experimental results of Transformer-based systems. The best performance is highlighted in bold.

- **METEOR**<sup>4</sup> (Denkowski and Lavie 2014) scores the generated text by aligning them to reference translations based on exact, stem, synonym, and paraphrase matches between words and phrases.

## Settings

We utilize *fast align*<sup>5</sup> (Dyer, Chahuneau, and Smith 2013) to perform word alignment in the hard visual agreement regularization. Faster R-CNN are used to extract pooled ROI features as visual input for each region. For each image, we consistently keep 36 highest probability objects.  $\lambda_1$  in Eq. (19) is set to 0.2 and 0.5 for EN→DE and EN→FR translations, respectively.  $\lambda_2$  in Eq. (20) is set to 0.2 and 0.1 for DE→EN and FR→EN translations, respectively.

For both source and target language, we limit the vocabulary size to 10,000. The size of word embedding is set to 512 and embeddings are learned from scratch. An extra linear layer are utilized to project all visual features into 512-dim. For the Seq2Seq version of our approach, the textual encoder and decoder are all a 2-layer LSTM with hidden size 512. We set the textual encoder to be bidirectional. For the transformer version of our approach, we set the hidden size of multi-head attention layer to 512 and the hidden size of the feed-forward layer to 2,048. The number of heads in multi-head attention is set to 8, while a transformer layer consists of 6 blocks.

We adopt the Adam optimization method with the initial learning rate 0.0003 for training and the learning rate is halved after each epoch. We also make use of dropout to avoid over-fitting.

<sup>4</sup><http://www.cs.cmu.edu/~alavie/METEOR/>

<sup>5</sup>[https://github.com/clab/fast\\_align](https://github.com/clab/fast_align)

## Baselines

We compare our proposed approaches with the following representative and competitive systems:

- **OnlyText** (Bahdanau, Cho, and Bengio 2014) is the traditional sequence-to-sequence framework with attention mechanism, which only encodes the source sentence as the input.
- **DAMT** (Calixto, Liu, and Campbell 2017) employs two separate attention mechanisms to integrate visual and textual features, respectively. Besides, the gate mechanism is further used to rescale visual information.
- **Imagination** (Elliott and Kádár 2017) learns both machine translation and visually grounding task simultaneously so that the vision-language joint semantic embeddings can be constructed.

## Results and Analysis

In this section, we report the experimental results. Besides, further analysis is also provided.

### Experimental Results

The evaluation results of different systems based on the Seq2Seq model are presented in Table 1, showing that both our hard and soft visual agreement regularization achieve better performance than competitive baselines in all translation directions. For instance, our soft regularization model outperforms the text-only Seq2Seq model and the best-performing baseline by 2.2 and 1.1 BLEU score respectively on the EN→DE language pair. This illustrates that visual features are capable of promoting the performance of machine translation, and our approaches can make better use of visual information.

Methods	EN→DE		EN→FR	
	BLEU	METEOR	BLEU	METEOR
OnlyText	27.1	49.3	51.3	69.1
+ Image	27.5	49.4	51.2	68.9
+ MHCA	28.6	50.1	51.7	69.3
Proposal (Hard)	<b>29.3</b>	<b>51.2</b>	<b>52.6</b>	<b>69.9</b>
Proposal (Soft)	29.2	50.9	52.4	69.5

Table 3: The results of incremental analysis. The best performance is highlighted in bold and “MHCA” means our proposed multi-head co-attention.

Table 2 presents the results of different systems based on Transformer, showing that both regularization approaches can substantially outperform the Transformer-based baselines. This demonstrates that our approaches are universal, which can bring about consistent improvements in performance on different base neural architectures. However, it is worth noting that although soft regularization can achieve better performance than all baselines in most translation directions, its performance is not comparable to hard regularization. The reason for this observation, as we suspect, might fall in the multi-head attention mechanism in Transformer, where word alignment is not well performed (Li et al. 2019). This leads to larger error when leveraging attention distribution to approximate the true alignment distribution, and the overall performance would be correspondingly inferior.

### Incremental Analysis

Table 3 presents results of incremental analysis. We treat the text-only Seq2Seq model as the base model and cumulatively add each component until the full model is rebuilt. The results show that our multi-head co-attention model can significantly improve the model performance, illustrating that building bidirectional interactions between textual and visual features can produce better co-dependent representations. Besides, Table 3 shows that either hard or soft regularization further contributes to generating high-quality translations. By constraining bidirectional models to focus on the same region of the image when generating semantically equivalent visual words, the model can make better use of visual information, resulting in more accurate machine translation hypotheses.

### Effectiveness of Adaptive Weighting

Considering that the generation of non-visual words requires little visual information from the image, we propose adaptive weighting, which aims to assign reasonable weights to visual regularization term automatically. Table 4 shows the performance of soft-agreement regularization with two different weighting strategies. In Table 4, “Frozen-WT” and “Adaptive-WT” denote frozen weighting<sup>6</sup> and adaptive weighting, respectively.

<sup>6</sup>Frozen weighting means that different words are assigned the same constant weight, as in Eq. (12).

Methods	EN→DE		EN→FR	
	BLEU	METEOR	BLEU	METEOR
OnlyText	27.1	49.3	51.3	69.1
Frozen-WT	28.8	50.7	52.1	69.4
Adaptive-WT	<b>29.2</b>	<b>50.9</b>	<b>52.4</b>	<b>69.5</b>

Table 4: The results of two weighting (WT) strategies based on the soft regularization.

As shown in Table 4, adaptive weighting can achieve better performance, increasing BLEU score from 28.8 to 29.2 on the EN→DE language pair. The reason is that based on the degree of dependence of the word to be generated on the image, the proposed adaptive weighting can adaptively adjust the weight of the regularization term. The weights corresponding to visual words are larger, which introduces visual agreement regularization to leverage visual information more effectively. For non-visual words, however, a smaller weight will be assigned to the regularization term in the training objective, which prevents the presence of misleading gradient signals.

Figure 4 visualizes the dependence (defined as  $\beta_{y_t}$  in Eq. (18)) of the different words to be generated on the image. The results show that some visual words (such as “man” and “woman”) will be assigned greater weights, while non-visual words have a very small weight. It shows that the proposed adaptive weighting strategy can automatically adjust the corresponding weight according to the degree of dependence of the generated words on the image, resulting in better model performance.

### Effectiveness of Improving Visual Agreement

In order to verify that our approach can effectively improve the agreement of visual attention of bidirectional models, here we introduce a new evaluation metric: visual agreement distance (VAD). The proposed VAD aims at characterizing the difference of visual attention when generating semantically equivalent words. Formally, for the given instance  $(\mathbf{v}, \mathbf{x}, \mathbf{y})$ , we define VAD as:

$$\text{VAD} = \frac{1}{|\mathcal{N}(\mathbf{x}, \mathbf{y})|} \sum_{(x,y) \in \mathcal{N}(\mathbf{x}, \mathbf{y})} \ell_1(A_{x \rightarrow \mathbf{v}}, A_{y \rightarrow \mathbf{v}}) \quad (21)$$

where  $\ell_1(\cdot, \cdot)$  denotes  $\ell_1$  distance.  $\mathcal{N}(\mathbf{x}, \mathbf{y})$  denoted mutually aligned word pairs in  $(\mathbf{x}, \mathbf{y})$ , which is defined as:

$$\mathcal{N}(\mathbf{x}, \mathbf{y}) = \{(x, y) | x \in \mathbf{x}, y \in \mathbf{y}, x \leftrightarrow y\} \quad (22)$$

where  $x \leftrightarrow y$  represents that  $x$  and  $y$  are uniquely aligned with each other. The large VAD means that the model suffers from the poor visual agreement. The evaluation results of VAD for different systems on EN→DE and EN→FR translations are shown in Table 5.

In Table 5, *non-visual* and *visual* denote VAD calculated on non-visual word pairs and the remaining visual word pairs<sup>7</sup>, respectively. The results show that our approach can

<sup>7</sup>We manually labeled such data for testing.




Figure 4: The visualization of the dependence of the right different words to be generated on the left image.

effectively improve the agreement of visual attention to the image of bidirectional models in the generation of semantically equivalent word pairs. In addition, such improvement is more obvious for visual word pairs compared with non-visual word pairs (e.g. 0.44 decline for *visual* v.s. 0.22 decline for *non-visual* on EN→DE). As analyzed in the previous section, the generation of visual words exhibits more dependence on the image. Therefore, the adaptive weighting will assign a greater weight to the corresponding regularization, leading to a more significant improvement.

## Related Work

In summary, this work is mainly related to the following several research lines.

**Multi-modal machine translation.** This task aims at translating the source sentence paired with additional modal (usually visual) information into another different language. Early work focuses on integrating additional visual features as part of inputs or the initialization of model modules. For instance, Huang et al. (2016) present parallel LSTM threads with multiple regional visual features and Calixto and Liu (2017) propose to treat visual features as words in the source sentence or initialization of the encoder. Going a step further, several prior works employ separate attention for different modalities to integrate modal-dependent information. For example, Calixto, Liu, and Campbell (2017) present a doubly-attentive decoder integrating two separate attention over the source information and a more effective hierarchical attention was proposed by Delbrouck and Dupont (2017). Iye, Madhyastha, and Specia (2019) propose an effective translation-and-refine framework, where visual features are only used by a second stage decoder. Inspired by multi-task learning, Elliott and Kádár (2017) perform machine translation while constraining the averaged representations of the shared encoder to be the visual embedding of the paired image. Zhou, Cheng, and Yu (2018) strives to construct a vision-language joint semantic embedding via a novel visual directional attention. There also exist some other endeavors on the MMT task. For example, Calixto, Rios, and Aziz (2019) sets a latent variable as a stochastic embedding which is used in the target-language decoder and to predict visual features. Chen, Jin, and Fu (2019) present a progressive learning approach for image pivoted zero-resource machine translation and Su et al. (2019) investigate the possibility of utilizing images for disambiguation to improve the performance of unsupervised machine translation.

VAD (↓)	EN→DE		EN→FR	
	<i>visual</i>	<i>non-visual</i>	<i>visual</i>	<i>non-visual</i>
<i>w/o agreement</i>	1.18	1.71	1.13	1.52
<i>w/o agreement</i>	0.74	1.49	0.71	1.45

Table 5: The results of visual attention distance (VAD) on the EN→DE and EN→FR translations. The symbol “↓” indicates that lower is better.

**Agreement-based Machine Translation.** Among this line of work, Cheng et al. (2016) present agreement-based joint training for bidirectional neural machine translation systems and Liu et al. (2016) propose to make use of approximate joint search for coupled translation models to produce more balanced outputs. Furthermore, Zhang et al. (2019) present a model regularization approach by minimizing Kullback-Leibler divergence between the probability distributions defined by bidirectional models. Other representative translation tasks also include bilingual lexicon induction (Yang et al. 2019b), unsupervised machine translation (Lample et al. 2018), and so on.

**Cross-modal generation.** This work can also be attributed to the category of cross-modal generation, which aims at generating the desired text in the presence of the source information from multiple modalities. Several typical tasks include image captioning (Xu et al. 2015; Liu et al. 2019), visual storytelling (Wang et al. 2018; Yang et al. 2019a), cross-modal machine commenting (Yang et al. 2019c) and so on.

## Conclusion

This work presents two novel visual agreement regularization approaches as well as an effective adaptive weighting strategy for multi-modal machine translation. By encouraging both forward and backward translation models to share the same focus on the image when generating semantically equivalent visual words, the proposed regularized training is capable of making better use of visual information. In addition, a simple yet effective multi-head co-attention model is also introduced to capture the interaction between visual and textual features. Extensive experimental results show that our approach can outperform competitive baselines by a large margin. Further analysis demonstrates that the proposed regularization approaches can effectively improve the agreement of attention on the image, leading to better use of visual information.

## Acknowledgement

We thank all reviewers for their thoughtful and constructive suggestions. We also thank Kai Fan and Xuancheng Ren for their instructive suggestions and invaluable help. This work was supported in part by National Natural Science Foundation of China (No. 61673028). Xu Sun is the corresponding author of this paper.

## References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Calixto, I., and Liu, Q. 2017. Incorporating global visual features into attention-based neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 992–1003.
- Calixto, I.; Liu, Q.; and Campbell, N. 2017. Doubly-attentive decoder for multi-modal neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 1913–1924.
- Calixto, I.; Rios, M.; and Aziz, W. 2019. Latent variable model for multi-modal translation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, 6392–6405.
- Chen, S.; Jin, Q.; and Fu, J. 2019. From words to sentences: A progressive learning approach for zero-resource machine translation with visual pivots. *arXiv preprint arXiv:1906.00872*.
- Cheng, Y.; Shen, S.; He, Z.; He, W.; Wu, H.; Sun, M.; and Liu, Y. 2016. Agreement-based joint training for bidirectional attention-based neural machine translation. In *IJCAI 2016*, 2761–2767.
- Delbrouck, J., and Dupont, S. 2017. An empirical study on the effectiveness of images in multimodal neural machine translation. *CoRR abs/1707.00995*.
- Denkowski, M. J., and Lavie, A. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, 376–380.
- Dyer, C.; Chahuneau, V.; and Smith, N. A. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *NAACL HLT 2013*, 644–648.
- Elliott, D., and Kádár, Á. 2017. Imagination improves multimodal translation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing*, 130–141.
- Elliott, D.; Frank, S.; and Specia, L. 2016. Multi30k: Multilingual english-german image descriptions. In *Proceedings of the 5th Workshop on Vision and Language*.
- Huang, P.; Liu, F.; Shiang, S.; Oh, J.; and Dyer, C. 2016. Attention-based multimodal neural machine translation. In *Proceedings of the First Conference on Machine Translation*, 639–645.
- Ive, J.; Madhyastha, P.; and Specia, L. 2019. Distilling translations with visual awareness. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, 6525–6538.
- Lample, G.; Ott, M.; Conneau, A.; and Ranzato, M. 2018. Phrase-based & neural unsupervised machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 5039–5049.
- Li, X.; Li, G.; Liu, L.; Meng, M.; and Shi, S. 2019. On the word alignment from neural machine translation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, 1293–1303.
- Liu, L.; Utiyama, M.; Finch, A. M.; and Sumita, E. 2016. Agreement on target-bidirectional neural machine translation. In *NAACL HLT 2016*, 411–416.
- Liu, F.; Ren, X.; Liu, Y.; Lei, K.; and Sun, X. 2019. Exploring and distilling cross-modal information for image captioning. In *IJCAI 2019*, 5095–5101.
- Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2016. Hierarchical question-image co-attention for visual question answering. In *NIPS 2016*, 289–297.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318.
- Ren, S.; He, K.; Girshick, R. B.; and Sun, J. 2015. Faster R-CNN: towards real-time object detection with region proposal networks. In *NIPS 2015*, 91–99.
- Su, Y.; Fan, K.; Bach, N.; Kuo, C.-C. J.; and Huang, F. 2019. Unsupervised multi-modal neural machine translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10482–10491.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. *CoRR abs/1706.03762*.
- Wang, X.; Chen, W.; Wang, Y.; and Wang, W. Y. 2018. No metrics are perfect: Adversarial reward learning for visual storytelling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 899–909.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A. C.; Salakhutdinov, R.; Zemel, R. S.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *ICML 2015*, 2048–2057.
- Yang, P.; Luo, F.; Chen, P.; Li, L.; Yin, Z.; He, X.; and Sun, X. 2019a. Knowledgeable storyteller: A commonsense-driven generative model for visual storytelling. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, 5356–5362.
- Yang, P.; Luo, F.; Chen, P.; Liu, T.; and Sun, X. 2019b. MAAM: A morphology-aware alignment model for unsupervised bilingual lexicon induction. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Volume 1: Long Papers*, 3190–3196.
- Yang, P.; Zhang, Z.; Luo, F.; Li, L.; Huang, C.; and Sun, X. 2019c. Cross-modal commentator: Automatic machine commenting based on cross-modal information. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Volume 1: Long Papers*, 2680–2686.
- Zhang, Z.; Wu, S.; Liu, S.; Li, M.; Zhou, M.; and Xu, T. 2019. Regularizing neural machine translation by target-bidirectional agreement. In *AAAI 2019*, 443–450.
- Zhou, M.; Cheng, R.; and Yu, Z. 2018. A visual attention grounding neural model for multimodal machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3643–3653.