

Supervised Visual Attention for Simultaneous Multimodal Machine Translation

Veneta Haralampieva

VENETA.L.HARALAMPIEVA@GMAIL.COM

Ozan Caglayan (Corresponding author)

O.CAGLAYAN@IC.AC.UK

Lucia Specia

L.SPECIA@IC.AC.UK

Department of Computing, Imperial College London, UK

Abstract

There has been a surge in research in multimodal machine translation (MMT), where additional modalities such as images are used to improve translation quality of textual systems. A particular use for such multimodal systems is the task of simultaneous machine translation, where visual context has been shown to complement the partial information provided by the source sentence, especially in the early phases of translation. In this paper, we propose the first Transformer-based simultaneous MMT architecture, which has not been previously explored in simultaneous translation. Additionally, we extend this model with an auxiliary supervision signal that guides the visual attention mechanism using labelled phrase-region alignments. We perform comprehensive experiments on three language directions and conduct thorough quantitative and qualitative analyses using both automatic metrics and manual inspection. Our results show that (i) supervised visual attention consistently improves the translation quality of the simultaneous MMT models, and (ii) fine-tuning the MMT with supervision loss enabled leads to better performance than training the MMT from scratch. Compared to the state-of-the-art, our proposed model achieves improvements of up to 2.3 BLEU and 3.5 METEOR points.

1. Introduction

Simultaneous machine translation (MT) aims at providing a computational framework that reproduces how human interpreters perform *simultaneous interpretation*. In this task, the duty of the interpreter is to translate speech in near real-time, by constantly maintaining a balance between the time needed to accumulate sufficient context and the translation *latency* the listeners experience in return. This *streaming* property is what differentiates simultaneous MT from the conventional MT approaches, which process complete source *sentences*. Traditional work in simultaneous translation have dealt with this *streaming* property by relying on syntactic or heuristic constraints (Bub et al., 1997; Ryu et al., 2006; Bangalore et al., 2012) to determine the amount of *waiting* prior to *committing* a partial translation. Similar approaches have also been explored using state-of-the-art neural MT (NMT) architectures (Sutskever et al., 2014; Bahdanau et al., 2015; Vaswani et al., 2017), such as rule-based deterministic policies implemented at (greedy) decoding time (Cho & Esipova, 2016) or the WAIT-K policy which sequentially switches between reading a new word and committing a translation (Ma et al., 2019). Adaptive policies, which attempt to learn when to *read* or *commit* depending on the context, have also been explored mostly through reinforcement learning-based techniques (Gu et al., 2017; Alinejad et al., 2018; Ive et al., 2021).

In this work, we focus on the *translation quality* aspect of the simultaneous translation framework and explore whether input contexts other than the linguistic signal can improve the performance of simultaneous MT models. Although such additional information may naturally occur in human simultaneous interpretation through the *a priori* knowledge of factors such as the topic, speaker or even the venue of the speech, in a computational model any additional context should be explicitly and carefully integrated to explore different inductive biases during model learning. Therefore, to mimic the availability of multiple input modalities for simultaneous MT, we follow the multimodal machine translation (MMT) framework where the objective is to translate image descriptions into different languages, by integrating the images themselves as additional context (Specia et al., 2016; Sulubacak et al., 2020). Intuitively, the expectation is that as long as there is a correlation between the language and the visual semantics, this way of grounding language can help anticipate future context for better translation quality, and even reduce the *latency* for adaptive policies.

Although relatively few, several works have explored a similar framework to analyse the benefits of visual grounding to simultaneous MMT: Caglayan et al. (2020a) and Imankulova et al. (2020) approached the problem by integrating visual features into the encoder and/or the decoder of recurrent MMT architectures and coupling them with deterministic WAIT-K policy (Ma et al., 2019) and rule-based decoding algorithms (Cho & Esipova, 2016). Ive et al. (2021) introduced reinforcement learning (RL)-based adaptive policies for recurrent MMT architectures.

In this paper, we first propose a Transformer-based (Vaswani et al., 2017) simultaneous MMT model where regional features extracted from a state-of-the-art object detector (Anderson et al., 2018), are fused with the source language representations using an attention based cross-modal interaction (CMI) layer. To implement simultaneous multimodal translation, we adopt the *prefix training* approach (Niehues et al., 2018; Arivazhagan et al., 2020) and evaluate its performance along with the deterministic WAIT-K policy (Ma et al., 2019). Next, we propose a novel approach to simultaneous MMT to improve the grounding – and therefore the anticipation – ability of our model. The proposed method involves *supervising* the alignment between the source language representations and the image regions, through the use of labelled phrase-region correspondences. We devise two multi-task learning settings to enable the visual supervision: (i) fine-tuning a pre-trained MMT for a fixed number of epochs, or (ii) training the MMT from scratch. We perform extensive experiments on the three different language pairs (English→{Czech,French,German}) of the Multi30k dataset (Elliott et al., 2016) and conduct thorough quantitative and qualitative analyses to understand the potential impacts of attention supervision. Our results show that (i) *prefix training* achieves substantially better scores than the WAIT-K approach, (ii) supervised visual attention consistently improves the translation quality of the MMT models, and (iii) fine-tuning the MMT offers better performance than training the MMT from scratch. Finally, our supervised models achieve up to 2.3 BLEU (Papineni et al., 2002) and 3.5 METEOR (Denkowski & Lavie, 2014) points improvements over the current state-of-the-art on Multi30k dataset.

The remaining of the paper is organised as follows: we provide a detailed account of literature in § 2 and describe our methodology and resources in § 3. In § 4, we present the quantitative and qualitative results across different language pairs and test sets of

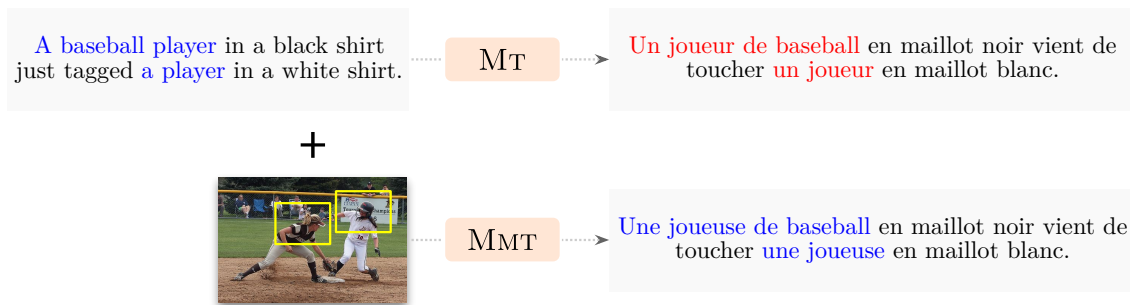


Figure 1: Multimodal machine translation (MMT) can help disambiguate the source sentence when translating into gender-marked languages. The example is taken from the Multi30k dataset (Elliott et al., 2016).

Multi30k. Finally, we conclude our work in § 5 with a discussion on possible directions for future research.

2. Related Work

The primary goal of Multimodal Machine Translation (MMT) is to improve the quality of machine translation (MT) by incorporating information from additional sources such as images or videos (Sulubacak et al., 2020). Of these two approaches, image-guided MMT (Figure 1) is substantially more researched than the video-guided MMT, simply due to the availability of more training resources. Since this paper heavily relies on image-guided MMT, we begin this section with a detailed literature overview on MMT first, and then continue with simultaneous MT.

2.1 Multimodal Machine Translation

Attentive models. Inspired by the success of the textual attention (Bahdanau et al., 2015) in NMT models, a considerable amount of research has focused on coupling visual attention and textual attention altogether to perform image-guided MMT. These works generally encode the images using state-of-the-art CNN models pre-trained on large computer vision datasets such as ImageNet (Deng et al., 2009). This way, an image is represented as a set of convolutional feature maps of size $K \times K \times C$ i.e. each output channel C_i encodes activations across a uniformly partitioned $K \times K$ grid. Using these features, Caglayan et al. (2016), Calixto et al. (2017) explore shared and separate attention mechanisms to compute a cross-modal latent representation in the multi-modal space. In addition to the language context computed by the textual attention, the translation decoder is now also conditioned on the multi-modal representation. As a follow-up, Libovický and Helcl (2017) proposed various extensions, of which the hierarchical attention variant gained popularity. This method weighs the relevance of textual and visual contexts using a third attention mechanism, instead of simpler fusion strategies such as addition or concatenation. Libovický et al. (2018) later adapted these extensions to Transformer-based (Vaswani et al., 2017) NMT models

to enable image-guided Transformer MMTs. Although the majority of the approaches in attentive MMT rely on decoder-side visual attention, encoder-side grounding was also explored, but only for recurrent architectures: Delbrouck and Dupont (2017) propose adding a visual attention layer in the encoder, where its states act as the *query* and the visual features are taken as *keys* and *values*.

Simple grounding. Finally, another line of work investigates the use of *pooled* visual representations ($\in \mathbb{R}^C$) for image-guided MMT, instead of the dense convolutional feature maps. These approaches usually initialise the hidden state of the encoder and/or the decoder in the model, using a projection of the visual feature vector (Calixto et al., 2016; Calixto & Liu, 2017; Caglayan et al., 2017). Multi-task learning is also explored (Elliott & Kádár, 2017; Zhou et al., 2018) using these simple vectorial representations, where the model is tasked with the reconstruction of visual features using the encoder’s output.

2.2 Supervision of Attention

Previous work has explored whether supervising the encoder-decoder attention can improve the alignment and translation quality of text-only NMT systems. Liu et al. (2016), Mi et al. (2016) investigate supervising the attention of a recurrent NMT model by adding an alignment loss which is jointly optimised alongside the negative log-likelihood objective. Garg et al. (2019) later extend this to Transformers-based text-only NMT models, showing that supervising a single attention head from the cross-attention layers of the decoder, can outperform existing alignment models without a significant degradation in translation quality.

Recently, Nishihara et al. (2020), Specia et al. (2021) investigate a multimodal co-attention mechanism (Lu et al., 2016) in the encoder, which first uses an affinity matrix to capture the relationship between the source tokens and image features and then computes the visual and textual attention weights. Specia et al. (2021) explore the impact of supervising the visual attention weights using ground-truth alignments obtained from the Flickr30k Entities dataset (Plummer et al., 2015) and show that their multimodal systems are better at disambiguating words, compared to their text-only baseline system. For Nishihara et al. (2020), the improvements are marginal unless cross-language alignments are also supervised in addition to the visual attention.

2.3 Simultaneous Machine Translation

Early works in simultaneous neural MT (SiMT) explore using a pre-trained full-sentence NMT model at test time, by relying on specific decoding approaches designed for simultaneous interpretation. Of these, Cho and Esipova (2016) propose a greedy decoding algorithm with two different waiting criteria based on simple heuristics which determine whether the model should *READ* a source token or *WRITE* a target one. Rather than relying on hand-crafted heuristics, several works investigate using Reinforcement Learning (RL) to learn an adaptive policy which maximises the translation quality and minimises the delay/latency. Satija and Pineau (2016) train an agent using Deep Q-Learning while Gu et al. (2017) rely on the policy gradient algorithm (Williams, 1992). Alinejad et al. (2018) later extend

the latter by adding a *PREDICT* action, which enriches the available context by utilising predictions of future source tokens.

A drawback of the approaches discussed so far is the discrepancy between training and test times of the underlying NMT model: the model being trained on full-sentence source contexts, are later exposed to partial contexts at test time. Dalvi et al. (2018) propose mitigating this by fine-tuning the model using either chunked segments or prefix pairs. Next, Ma et al. (2019) explore a fixed WAIT-K policy which can be used at both training and test times, with the model initially reading k source tokens before proceeding to alternate between writing a single target token and reading a source one. Later, Arivazhagan et al. (2019) extend this to an adaptive policy using an advanced attention mechanism in a Recurrent NMT model and a weighted latency loss, while Ma et al. (2020) further develop this for the multi-head attention used in Transformers. More recently, Arivazhagan et al. (2020) investigate whether training a model on prefix pairs and re-translating previously emitted target words at decoding time improves translation quality or not. Their results show that augmenting the training data with prefix pairs can outperform the WAIT-K trained systems, with re-translation further increasing the quality.

An alternative method for obtaining an adaptive policy is to train a policy model using Supervised Learning and ground-truth action sequences which Zheng et al. (2019) propose generating using a pre-trained NMT Transformer model, with the ideal action being a *WRITE* when the ground truth target word is ranked within the top k next word candidates, or *READ* otherwise. Arthur et al. (2021) rely on a statistical word alignment system to obtain the ground-truth actions and use them to jointly train the translation and policy models.

Simultaneous MMT. Previous work in simultaneous MMT mostly rely on rule-based strategies (Cho & Esipova, 2016; Ma et al., 2019) on the Multi30k dataset for recurrent MMT models. Of these approaches, Imankulova et al. (2020) explored the WAIT-K policy using a recurrent MMT equipped with a hierarchical multimodal attention. Specifically, for each k , they first conduct a textual pre-training (i.e. without the visual features) until convergence, and then fine-tune the checkpoint with visual features enabled. Caglayan et al. (2020a) conducted a study where they compare object classification and object detection features for two different multimodal architectures: decoder-level visual attention and encoder-level visual attention. As for the simultaneous translation part, they explore both WAIT-K and rule-based decoding (Cho & Esipova, 2016) methods. Finally, Ive et al. (2021) attempted to learn a multimodal policy through reinforcement learning, for deciding the *READ/WRITE* actions during simultaneous translation.

Our work resembles to Caglayan et al. (2020a) as we explore the deterministic WAIT-K policy along with the object detection features extracted from salient regions. We also add another policy to our inventory and more importantly, we investigate the impact of supervising the visual attention using human-labelled annotations.

3. Method

This section presents the approaches explored in this work. We first begin with a description of the text-only Transformers NMT (Vaswani et al., 2017) and how we extended it to accommodate simultaneous translation. Next, we describe our baseline multimodal Trans-

former architecture that incorporates visual attention. Finally, we introduce our approach to supervise the visual attention in simultaneous MMT.

In what follows, the source sentence and target sentence tokens are denoted with $\mathbf{x} = [x_1, x_2, \dots, x_N]$ and $\mathbf{y} = [y_1, y_2, \dots, y_M]$, respectively.

3.1 Transformer-based NMT

Transformers NMT (Vaswani et al., 2017) are the state-of-the-art sequence-to-sequence architectures equipped with deep encoder and decoder stacks that rely heavily on feed-forward layers in contrast to recurrent NMTs (Sutskever et al., 2014; Bahdanau et al., 2015). Combined with the use of *self-attention* layers, these changes allow for (i) better gradient dynamics during training and thus deeper architectures, and (ii) different inductive biases than the left-to-right processing nature of recurrent NMTs. The overall diagram of a Transformers-based NMT is given in Figure 2. In the following, we briefly explain the encoder and the decoder blocks of Transformers.

3.1.1 ENCODER

The transformer encoder $f()$ encodes the source sentence \mathbf{x} into a latent representation $\mathbf{h} = f(\mathbf{x})$ using a series of operations. The basic block in the encoder transforms the input using a self-attention layer followed by a feed-forward layer. Layer normalisation (Ba et al., 2016) and residual connections (He et al., 2016) are incorporated to achieve stability and improved training dynamics. This basic block (shown in the left part of Figure 2) is further replicated B times in a vertical fashion so that each layer consumes the output of the previous one. The inputs to the first layer are the embeddings of the source sentence tokens $\mathbf{x} = [x_1, x_2, \dots, x_N]$, additively shifted by positional encodings. The latter is crucial to embed the positional information of words into the embeddings so that the representation is not invariant to the word order. The output of the last encoder layer is passed through a final layer normalisation.

3.1.2 ATTENTION

A key aspect of the Transformer architecture is the extensive use of several types of attention layers. In its simplest form, the attention mechanism (Bahdanau et al., 2015) allows computing the weighted sum of a set of vectors $\mathbf{V} \in \mathbb{R}^{N \times D}$ where the weights are obtained based on dot product scores between keys $\mathbf{K} \in \mathbb{R}^{N \times D}$ and queries $\mathbf{Q} \in \mathbb{R}^{N \times D}$. The scores are then normalised with the softmax operator to produce a valid probability distribution and finally multiplied with \mathbf{V} as follows:

$$A(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{D}} \right) \mathbf{V} \quad (1)$$

To further enrich the learned representations, multiple attention representations are computed in parallel, by supplying different projections of queries, keys and values as inputs to each attention function $A^{(i)}$ (Equation 2). With n attention heads, the final multi-head attention output is computed by projecting the concatenation of all attention head outputs

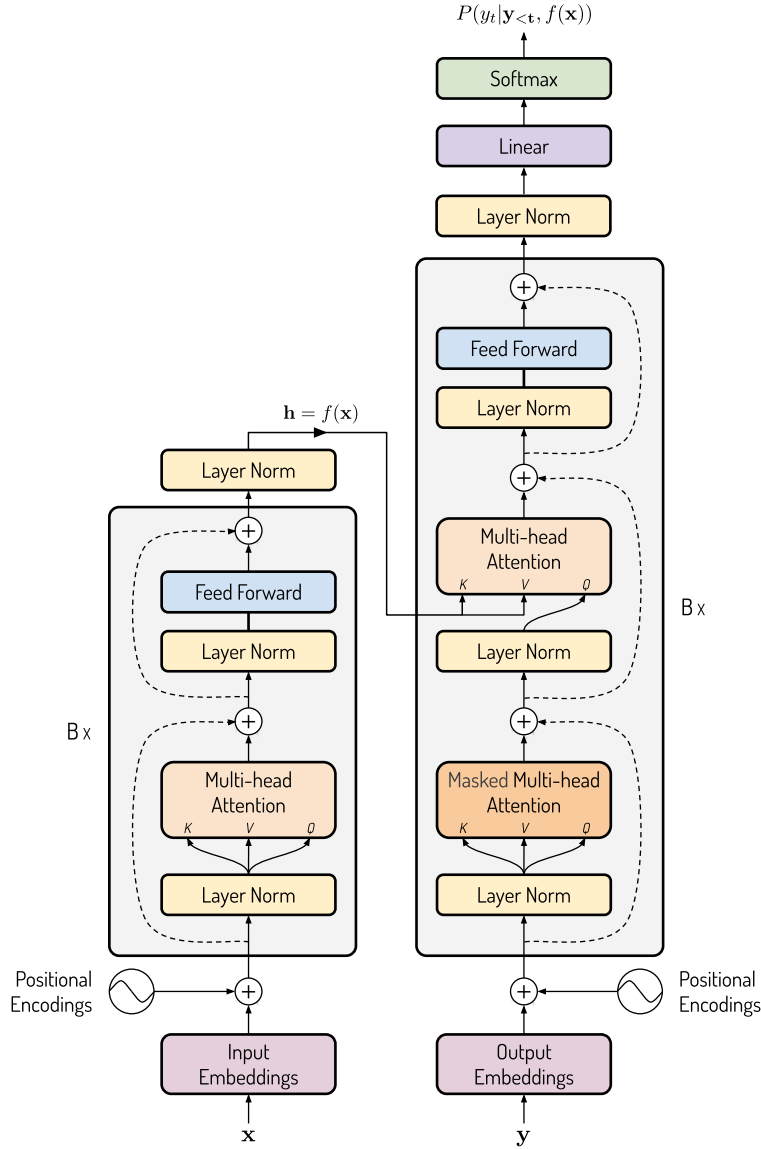


Figure 2: The architecture of a Transformers NMT (Vaswani et al., 2017): each block is repeated B times to create a deep model. This is the more stable “pre-norm” variant (Wang et al., 2019) where layer normalisation is applied prior to each sub-layer. The dashed lines denote the residual connections.

as follows:

$$A^{(i)} = A(\mathbf{W}_q^{(i)} \mathbf{Q}, \mathbf{W}_k^{(i)} \mathbf{K}, \mathbf{W}_v^{(i)} \mathbf{V}) \tag{2}$$

$$\text{MHA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{W}_o \text{Concat}([A^{(1)}, \dots, A^{(n)}]) \tag{3}$$

In the Transformer-based NMT, each encoder block includes a multi-head self-attention layer to capture the relationship between different source positions. Additionally, the de-

coder uses (i) *masked* multi-head self-attention to model the causal relationship between each position and the ones preceding it and (ii) cross-attention to integrate source semantics crucial to perform the translation. In fact, the only difference between self-attention and cross-attention is that the former sets $\mathbf{Q} = \mathbf{K} = \mathbf{V}$ to the output of the previous layer whereas the latter uses the output of the **encoder** to set $\mathbf{K} = \mathbf{V} = f(\mathbf{x})$ (Figure 2).

3.1.3 DECODER

Once the source sentence is encoded into the latent representation \mathbf{h} , the decoder sequentially generates the target sentence in an auto-regressive way. This means that the probability of the next target token ($P(y_t|\mathbf{y}_{<t}, \mathbf{h})$) is conditioned on the history ($\mathbf{y}_{<t}$) of target words predicted so far, in addition to the source sentence semantics encoded in \mathbf{h} . In this formulation, the whole decoder can be thought as a building block that implements the aforementioned probability term $P()$.

In terms of computation, a decoder block is very similar to an encoder one except that (i) the self-attention is now *masked* to enforce that the decoder is causal i.e. it does not mistakenly look at future positions and (ii) a secondary multi-head attention known as **cross-attention**, integrates information from the **encoder** through the latent sentence representations \mathbf{h} (§ 3.1.2). Finally, we train the model in an end-to-end way and minimise the negative log-likelihood of the sentence pairs in the training set \mathcal{D} :

$$\mathcal{L}_{\text{MT}} = - \sum_i^{\mathcal{D}} \log \left(P(\mathbf{y}^{(i)}|\mathbf{x}^{(i)}) \right) \quad \text{where} \quad P(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^{|\mathbf{y}|} P(y_t|\mathbf{y}_{<t}, \mathbf{h}) \quad (4)$$

3.1.4 OUR BASELINE

We use the **Base** Transformer (Vaswani et al., 2017) configuration in all our experiments, where both the encoder and decoder have 6 layers ($B = 6$ in Figure 2), each attention layer has 8 heads, the model dimension is 512 and the feed forward layer size is 2048. Additionally, we share the parameters of the target and output language embedding matrix (Press & Wolf, 2017). We should note that our implementation applies the “pre-norm” (Wang et al., 2019) formulation where the layer normalisation is placed prior to each sub-layer rather than after, to increase stability.

During training, we optimise the models using Adam (Kingma & Ba, 2014) and decay the learning rate with the *noam* scheduler (Vaswani et al., 2017). The initial learning rate, β_1 and β_2 are 0.2, 0.9 and 0.98, respectively. The learning rate is warmed up for 4,000 steps. We use a batch size of 32, apply label smoothing with $\epsilon = 0.1$ (Szegedy et al., 2016) and clip the gradients so that their norm is 1 (Pascanu et al., 2014). We train each system 3 times with different random seeds for a maximum of 100 epochs, with early stopping based on the validation METEOR (Denkowski & Lavie, 2014) score, which is the official metric used in all shared tasks in MMT (Barrault et al., 2018). The best checkpoint with respect to validation METEOR is selected to decode test set translations using the greedy search algorithm.




Figure 3: WAIT-K decoding (Ma et al., 2019): Initially, the decoder waits k words to be read before committing the first translation. Afterwards, the algorithm switches back and forth between read and write actions until the end-of-sentence marker is generated by the decoder.

3.2 Simultaneous NMT

This section describes the different training strategies that are used in this work to realise simultaneous machine translation. Following the notation from Ma et al. (2019), we first define a function $g(t)$ that returns the number of source tokens read so far by the encoder, at a particular decoding timestep t . By definition, $0 \leq g(t) \leq |\mathbf{x}|$ for all values of t . We can formulate this by modifying Equation 4 so that the source representation now depends on $g(t)$ as follows:

$$P(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^{|\mathbf{y}|} P(y_t | \mathbf{y}_{<t}, f(\mathbf{x}_{\leq g(t)})) \quad (5)$$

This generalisation allows us to define the conventional full sentence (consecutive) NMT as well, by using a constant function $g(t) = |\mathbf{x}|$ for all t and all sentence pairs.

3.2.1 WAIT-K DECODING

The WAIT-K policy (Ma et al., 2019) simply amounts to selecting a particular function $g_k(t)$ that depends on a pre-determined k value. This value determines the number of source tokens to be initially read by the encoder, before beginning the translation decoding process. Afterwards, the algorithm switches back and forth between read and write actions until the end-of-sentence marker (EOS) is predicted by the decoder. Mathematically speaking, the definition of $g_k(t)$ is as follows:

$$g_k(t) = \min(k + t - 1, |\mathbf{x}|) \quad (6)$$

The WAIT-K decoding algorithm is a *fixed-delay* policy because the translator always lags behind the speaker by k tokens. According to Ma et al. (2019), this policy is inspired by human interpreters who intuitively wait for some context to accumulate prior to starting the translation. For our multimodal purposes, having a fixed policy rather than an adaptive one is useful as fixed latency allows focused systematic analysis of quality improvements. A depiction of the algorithm is given in Figure 3.

Traditionally, the encoder in the Transformer NMT is bi-directional due to the self-attention mechanism which allows each position to attend to each other one. In the simultaneous translation framework, however, this can be challenging as it implies that every time

Algorithm 1: Prefix training (Niehues et al., 2018; Arivazhagan et al., 2020)

```

inputs : The current mini-batch  $B$ 
output: The modified mini-batch  $\hat{B}$ 
1  $\hat{B} \leftarrow []$ 
2 for  $(\mathbf{x}, \mathbf{y})$  in  $B$  do
3    $c \sim \text{Uniform}(0, 1)$            // Apply truncation with  $p = 0.5$ 
4   if  $c < 0.5$  then
5      $l_x \sim \text{Uniform}(1, |\mathbf{x}|)$    // Randomly sample a source prefix length
6      $l_y \leftarrow \text{Round}\left(\frac{l_x \cdot |\mathbf{y}|}{|\mathbf{x}|}\right)$  // Keep the same proportion for the target
7      $l_y \leftarrow \max(2, l_y)$    // Always include the BOS token
8   else
9      $\hat{\mathbf{x}} \leftarrow \mathbf{x}; \hat{\mathbf{y}} \leftarrow \mathbf{y}$  // No truncation
10  end
11   $\hat{B}.\text{append}((\hat{\mathbf{x}}, \hat{\mathbf{y}}))$ 
12 end

```

a new source token is read, the encoder’s representation \mathbf{h} would need to be recomputed, leading to a quadratic runtime cost with respect to the source sentence length. Instead, we implement an **uni-directional** encoder following Elbayad et al. (2020), by employing a **masked** self-attention in the encoder. This prevents future positions that are not read so far from being attended to, similar to the self-attention layers in the decoder block.

3.2.2 SIMULTANEOUS-AWARE TRAINING

A simple way to perform simultaneous NMT is by first training a consecutive NMT model (i.e. $g(t) = |\mathbf{x}|$) and then translating the test sets following the WAIT- k policy, where for each k we would use a different choice of $g_k(t)$ at inference time. We refer to this approach as WAIT- k DECODING. However, this creates a discrepancy between training and testing time, as the model will be exposed to partial source sentences at test time, although it was always trained on full sentences. For this reason, Ma et al. (2019) also proposed to use the same $g_k(t)$ function at both training and inference time, an approach that we refer to as WAIT- k TRAINING. We follow this approach for our initial set of experiments.

Furthermore, we adopt a second simultaneous-aware training recipe called **prefix training** (Niehues et al., 2018; Arivazhagan et al., 2020) which employ a simple data processing strategy to mitigate the aforementioned exposure bias between training and testing time. Specifically, for each sentence pair (\mathbf{x}, \mathbf{y}) in the mini-batch, we flip a fair coin to decide whether we will consider it for truncation or not. If it is considered, we first randomly sample a prefix length l_x for the source sentence and truncate the corresponding target sentence with the same proportion as the source side (Algorithm 1). Unlike WAIT- k TRAINING which requires training a separate model for each value of k , we train a single *prefix* model and decode the final checkpoint using WAIT- k DECODING across different values of k .

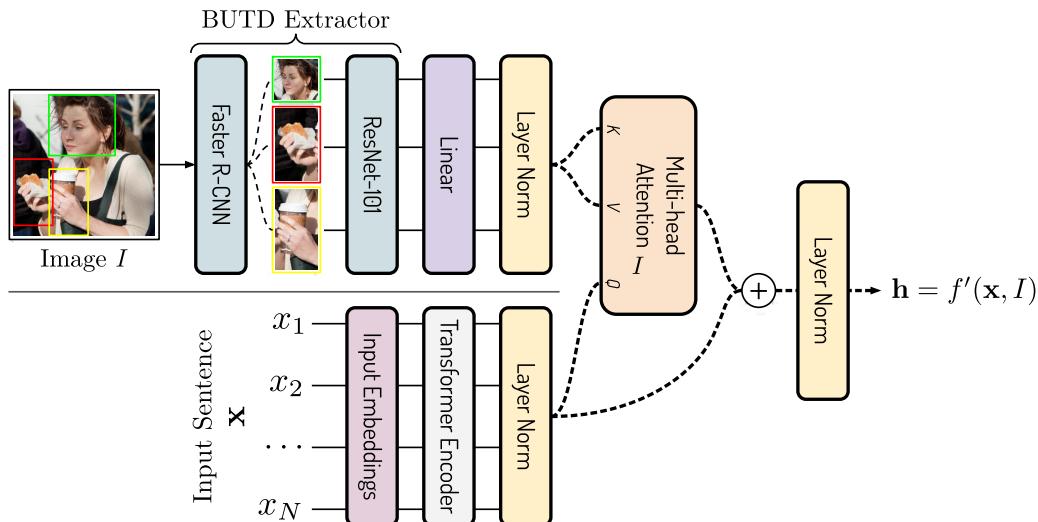


Figure 4: Transformer-based multimodal **encoder** for simultaneous NMT: The upper stream encodes visual features, whereas the bottom one is the usual Transformer encoder for the input sentence. Dashed lines denote a collection of vectors.

3.3 Multimodal NMT

In this section, we describe our take on integrating visual information to our consecutive and simultaneous NMT models. For that, we reformulate the **encoder-attention** variant of Caglayan et al. (2020a) for Transformer-based models. In what follows, we first describe the visual feature extraction pipeline and then present our multimodal NMT architecture in detail.

3.3.1 VISUAL FEATURES

To represent visual semantics, we explore regional object-detection features extracted using the popular *bottom up top down* (BUTD) approach (Anderson et al., 2018). BUTD combines the Faster R-CNN object detector (Ren et al., 2015) with a ResNet-101 backend (He et al., 2016) to perform feature extraction. We use the provided model¹ which is pre-trained on the large-scale Visual Genome dataset (Krishna et al., 2017). Having 1,600 object labels in its inventory, the BUTD detector is quite rich and for that reason has been used in most previous work in cross-modal pre-training (Lu et al., 2019; Tan & Bansal, 2019).

For our purposes, we use the default settings of the detector and extract 36 regional features per each image. In other words, alongside the language representation $\mathbf{x} \in \mathbb{R}^{N \times D}$ of a given source sentence, the associated image is represented with a set of regional features $\mathbf{v} \in \mathbb{R}^{36 \times 2048}$. The BUTD extractor is not further trained/fine-tuned during model training.

1. <https://hub.docker.com/r/airsplay/bottom-up-attention>

3.3.2 MULTIMODALITY

Our multimodal MT approach reformulates the **encoder-attention** variant of Caglayan et al. (2020a) for Transformer-based architectures. The main motivation for implementing an encoder-side cross-modal interaction is the nature of the simultaneous MT problem: since the source context is partial and grows gradually, the additional visual information can complement the missing context and allow the model to anticipate target words in a grounded way. Moreover, incorporating cross-modality at encoder side is also crucial for our second set of experiments regarding the supervision of the visual attention module, which aims at better language grounding (§ 3.3.3).

Figure 4 summarises the overall architecture of our MMT model, where the upper stream implements the visual representation module. This module (i) extracts the set of regional feature vectors $\mathbf{v} \in \mathbb{R}^{36 \times 2048}$, (ii) projects them to the D -dimensional space yielding $\mathbf{v}' \in \mathbb{R}^{36 \times D}$ and finally (iii) employs a multi-head² attention layer for cross-modal interaction (CMI). The key, value, query configuration of the cross-modal attention determines the nature of the interaction i.e. we set $\mathbf{K} = \mathbf{V} = \mathbf{v}'$ whereas the query \mathbf{Q} receives the text encoder’s output. This way, we get a cross-modal representation out of the attention layer which computes the weighted sum of regional feature projections (\mathbf{V}) based on the similarity between language representations (\mathbf{Q}) and regional features (\mathbf{K}). Since this output is a linear combination of visual vectors only, we augment it with the text encoder’s outputs using element-wise addition, similar to a residual connection. A final layer normalisation is applied on top to obtain the multimodal encoding $\mathbf{h} = f'(\mathbf{x}, I)$. The rest of the architecture is the same as Figure 2 in the sense that \mathbf{h} is passed to the cross-attention layer of the decoder.

3.3.3 SUPERVISING THE VISUAL ATTENTION

The cross-modal relationship between language and visual representations in attentive MMT models is generally learned in an unsupervised way, as the MMT model is only trained through a sequence-to-sequence cross-entropy objective. This is also the approach that we followed for our MMT model in § 3.3.2.

In this section, we are interested in whether explicit supervision of the visual attention could introduce an inductive bias so that the anticipation ability – and therefore the translation quality – of our simultaneous MMT models is further improved. We devise a **multi-task learning** scheme where in addition to the cross-entropy objective used for the MT task (Equation 4), we employ an auxiliary loss which tries to bring together the cross-modal attention distribution \mathbf{M}' computed by the CMI module and the ground-truth alignment \mathbf{M} . An illustration of the predicted attention and its ground-truth distribution is provided in Figure 5.

Ground-truth alignments. In order to supervise the visual attention during training, we require a labeled dataset of phrase-region annotations for the Multi30k dataset. Since the Multi30k dataset is derived from the Flickr30k image captioning dataset, we rely on the Flickr30k Entities dataset (Plummer et al., 2015) which provides human-annotated bounding boxes for noun phrases in the English image captions. For instance in Figure 5,

2. Throughout this work, we set the number of heads to 1 for the cross-modal interaction module.

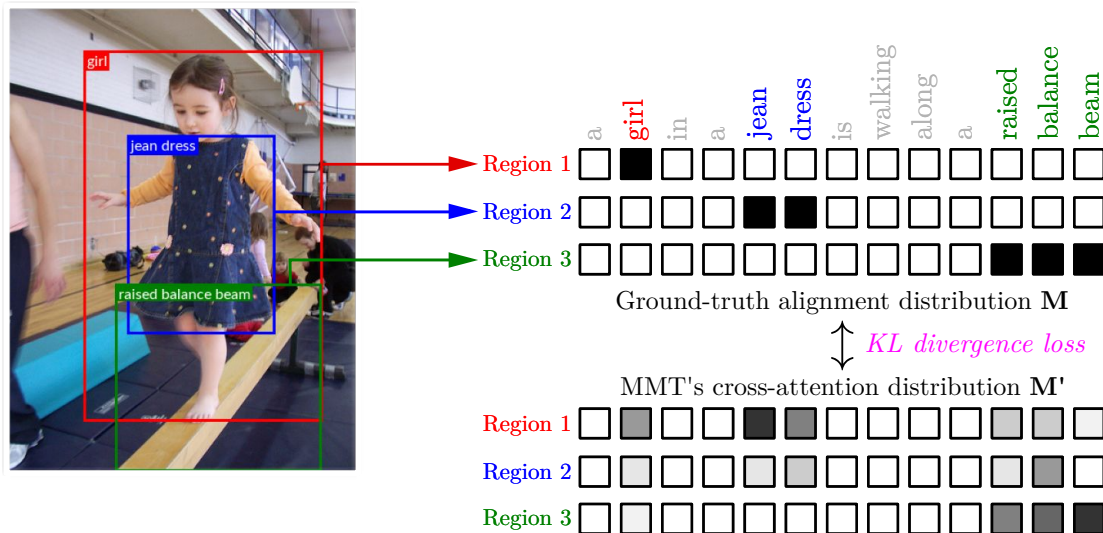


Figure 5: Cross-modal supervision of attention: The attention distribution from the model is pulled towards the ground-truth alignment matrix obtained from the Flickr30k entities dataset, using KL-divergence. Greyed out token positions do not contribute to the loss.

we can see that each of the phrases “a girl”, “a jean dress” and “a raised balance beam” is mapped to a bounding box³ that denotes the object referred.

Since the ground-truth regions are different from the regions predicted by the pre-trained BUTD detector (§ 3.3.1), we **re-extract** regional features using the same pre-trained BUTD detector but from the set of regions provided by the Entities dataset. On average, we end up with 4.3 bounding box annotations per image, which is substantially lower compared to the fixed number of 36 regions that we extracted from the pre-trained BUTD detector. Finally, we should note that the Entities dataset does not provide any annotations for the `test2017` and `testCOCO` test splits of Multi30k dataset.

Training. For a given source sentence \mathbf{x} with N words, first an alignment matrix $\mathbf{M} \in \mathbb{R}^{R \times N}$ is formed where R denotes the total number of ground-truth bounding box annotations for the overall sentence. We set $\mathbf{M}_{[i,j]} = 1$ if the word x_j is associated with the region i . If $k > 1$ region associations exist for the word x_j , the probability mass assigned to each region becomes $1/k$. The columns of the matrix that refer to words without any bounding box associations are not taken into account when computing the final alignment loss.

Alignment loss. The Kullback-Leibler (KL) divergence is a measure of how much a given probability distribution Q differs from a reference distribution P . In other words, minimising the KL-divergence between the predicted cross-modal attention distribution (Q) and the ground-truth alignments from the entities dataset (P) allows the model to generate a visual attention distribution that is closer to the human-labeled annotations. The final multi-task

3. The number of bounding boxes for a given noun phrase is not limited to one.

Split	English		German		French		Czech		# Sents
	Words	Len	Words	Len	Words	Len	Words	Len	
train	380K	13.1	364K	12.6	416K	14.4	298K	10.3	29,000
val	13.4K	13.2	13.1K	12.9	14.6K	14.4	10.4K	10.2	1,014
test2016	13.0K	13.1	12.2K	12.2	14.2K	14.2	10.5K	10.5	1,000
test2017	11.4K	11.4	10.9K	10.9	12.8K	12.8	-	-	1,000
testCOCO	5.2K	11.4	5.2K	11.2	5.8K	12.5	-	-	461

Table 1: Multi30k dataset statistics: “Words” denote the total number of words in a split whereas “Len” is the average number of words per sentence in that split.

objective combines the MT loss and the alignment loss with coefficients $\alpha = \beta = 1$:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{MT}} + \beta D_{\text{KL}}(P = \mathbf{M} \parallel Q = \mathbf{M}') \quad (7)$$

Fine-tuned supervision. In addition to training the supervised MMT from scratch, we also experiment with taking the best checkpoints of the MMT models with unsupervised visual attention ($\alpha = 1, \beta = 0$) and fine-tuning them with the alignment loss enabled ($\alpha = \beta = 1$). For these particular variants, we disable the learning rate scheduling and lower the learning rate to $1e - 5$. We track performance with respect to validation set METEOR and keep the best checkpoint for further decoding of the test sets.

3.4 Dataset

We use the Multi30k dataset (Elliott et al., 2016)⁴, which is a multi-lingual extension to the Flickr30k image captioning dataset (Young et al., 2014). Specifically, Multi30k provides a training set of 29,000 examples where each example corresponds to an image I and its English caption \mathbf{x} from Flickr30k, extended with the German translation \mathbf{y} of the English caption. In other words, both training and testing samples are triplets in the form of $\{I, \mathbf{x}, \mathbf{y}\}$. The dataset is later extended with French (Elliott et al., 2017) and Czech (Barrault et al., 2018) translations, making it the standard dataset for work on MMT, simultaneous MMT (Caglayan et al., 2020a; Imankulova et al., 2020), as well as other multilingual multimodal tasks.

We experiment with all three language pairs, namely, English→German, English→Czech and English→French. The training, validation and the **test2016** test sets are available for all language directions, whereas the test sets **test2017** and **testCOCO** sets are only available for German and French. The latter test set is specifically designed (Elliott et al., 2017) to contain at least one ambiguous word per image caption where images are selected from the COCO dataset (Chen et al., 2015) instead of the in-domain Flickr30k dataset. Some common statistics of the dataset are provided in Table 1.

Preprocessing. We use Moses tools (Koehn et al., 2007) to lowercase, punctuation-normalise and tokenise the sentences with *hyphen splitting* option enabled. We then create

4. <https://github.com/multi30k/dataset>

		2016			2017			COCO		
		CONS	TR	PREF	CONS	TR	PREF	CONS	TR	PREF
$k=1$	EN-Cs	17.1	13.0	22.4	-	-	-	-	-	-
	EN-DE	40.3	41.3	48.8	37.2	35.8	43.2	35.9	33.3	39.6
	EN-FR	59.5	58.7	63.0	55.6	52.6	58.0	52.6	48.9	53.7
$k=2$	EN-Cs	21.2	19.0	24.1	-	-	-	-	-	-
	EN-DE	46.7	48.7	52.4	42.4	42.4	46.3	41.3	39.4	42.4
	EN-FR	66.9	66.3	68.1	62.0	60.7	62.5	58.7	56.7	58.4
$k=3$	EN-Cs	23.9	23.3	26.5	-	-	-	-	-	-
	EN-DE	51.4	51.8	53.9	46.0	45.1	47.8	44.4	41.4	43.6
	EN-FR	71.4	71.3	71.5	65.9	65.4	65.2	61.1	59.8	60.1

Table 2: METEOR comparison of consecutive training (CONS), WAIT-K training (TR) and prefix training (PREF) approaches for **unimodal** simultaneous MT. All systems are decoded with the WAIT-K policy where $k \in \{1, 2, 3\}$. The scores are averages of three runs with different seeds. Best systems are indicated in **bold** typeface.

word vocabularies using the *training* subset of each language. The resulting English, French, German and Czech vocabularies have 9.8K, 11K, 18K and 22.2K words, respectively. We do not use sub-word segmentation approaches to avoid their potential side effects for simultaneous MT and to be able to analyse the grounding capability of the models more easily when it comes to cross-modal attention. We note that word-level MT performs as well as sub-word segmentation on this particular dataset, according to (Caglayan, 2019).

4. Results

4.1 Unimodal Simultaneous MT

Our first set of experiments focus on how different training regimes impact the translation quality for Transformer-based unimodal (i.e. text-only) simultaneous MT. For this, we compute METEOR scores of English-Czech, English-German and English-French MT systems, across three different test sets of Multi30k. Table 2 summarises the results obtained by performing WAIT-K decoding policy with $k \in \{1, 2, 3\}$, across the methods described in § 3.2.

First, we observe that the prefix training method yields substantially higher METEOR scores than the other two approaches in general. WAIT-K training achieves the lowest quality across all language pairs and test sets, an observation in line with the previous findings of Caglayan et al. (2020a). Overall, the quantitative results corroborate our initial hypothesis regarding the exposure bias between training and test time (§ 3.2.2) and show that the prefix training is indeed a good choice in-between the full-sentence (consecutive) training and the aggressive WAIT-K training approaches.

For **English-Czech** particularly, we observe that when $k = 1$, WAIT-K training is considerably worse ($\downarrow 4.1$ METEOR) than the other approaches. We hypothesise that the lower

SRC:	A woman holding a bowl of food in a kitchen.
CS:	Muž v červeném triku drží své jídlo.
DE:	Ein Frau hält eine schüssel mit Essen in einer Küche.
FR:	Un femme tenant un bol de nourriture dans une cuisine.

Table 3: Examples showing the effect of the lack of context for the decoder early in the sentence when using WAIT-1 decoding: greyed out words are those that have not been read or written yet. Red and blue indicate incorrect and correct translations, respectively.

scores for the English-Czech pair in general is likely to be caused by the fact that, unlike English, German and French, Czech has no articles preceding a noun. This means that when generating the first word of the Czech translation, the probability distribution will always generate the same word, which is likely to be the most common target sentence prefix in the training set. This is illustrated in Table 3 where the Czech model translates “woman” as “Muž” (man) whereas the German and French models are able to read the source word “woman” just before translating it⁵. We believe that the incorporation of the visual modality would allow simultaneous MT models to handle such cases in a better way.

Finally, although the lack of context affects all three methods at decoding time, the reason why WAIT-K training lags dramatically behind the other two for Czech is related to the fact that, WAIT-K training is ineffective in terms of leveraging the training resources, especially when the vocabulary is sparse. This makes Czech the most challenging pair across the pairs explored as its vocabulary has 22K unique tokens compared to 11K for French.

4.2 Multimodal Simultaneous MT

Based on the more promising results of the prefix training regime for the unimodal MT experiments, from now on we focus on prefix training across all multimodal experiments. We begin by analysing the impact of incorporating visual information into our Transformer-based simultaneous NMT models, following the architecture described in § 3.3.2. Table 4 compares METEOR scores between the baseline unimodal NMT and the MMT models across all language pairs and test sets.

Overall, our encoder-based cross-modal interaction method improves METEOR scores up to +1.5 points (EN-FR 2016) when compared to unimodal NMT models. This shows that the addition of the visual modality is beneficial for simultaneous MT, especially in low latency scenarios. In particular, consistent quality improvements are observed on the `testCOCO` test split, which (i) does not come from the Flickr30k distribution and (ii) contains more ambiguous words by construction.

For certain language pairs such as English→Czech and English→German, the additional modality brings little to no improvements, especially on the `test2016` test set. However, we will observe different trends once we add supervision (§ 4.2.1). Finally, the **En-**

5. Although the same issue also affects the article choices for gender-marked languages such as French and German, the arbitrary choices are more likely to match the correct article in the target language, since the number of possible translations for articles is much smaller.

		EN-Cs	EN-DE			EN-FR		
		2016	2016	2017	CC	2016	2017	CC
$k=1$	BASELINE NMT	22.4	48.8	43.2	39.6	63.0	58.0	53.7
	+ IMAGES	21.9	48.8	43.3	39.7	64.5	58.2	54.2
$k=2$	BASELINE NMT	24.1	52.4	46.3	42.4	68.1	62.5	58.4
	+ IMAGES	24.4	52.1	46.4	43.1	69.0	63.2	58.6
$k=3$	BASELINE NMT	26.5	53.9	47.8	43.6	71.5	65.2	60.1
	+ IMAGES	26.5	53.5	47.8	44.1	71.8	65.7	60.4

Table 4: METEOR comparison of **unimodal** and **multimodal** simultaneous MT under the **prefix training** framework: the scores are averages of three runs with different seeds. Best systems are indicated in **bold** typeface.

glish→**French** language pair benefits the most from the visual modality on average, in terms of METEOR scores. We find it interesting that as the lexical diversity of the target language (i.e. sparsity of the training set vocabulary) increases, the improvements due to multi-modality decrease. This hints at the fact that the improvements can be more pronounced if the models are (pre-)trained on larger linguistic resources, an exploration that we leave to future work.

4.2.1 SUPERVISED ATTENTION

Our next set of experiments aim to understand whether guiding the alignment between the visual features and source words during training helps improve the translation and grounding ability of the MMT models. We carry out two different experiments: first, we train a multimodal simultaneous translation model from **SCRATCH**, using both the translation and attention losses. As an alternative, we also experiment with a **FINE-TUNED** version of the best (unsupervised visual attention) MMT checkpoints with a reduced learning rate (§3.3.3).

The METEOR scores for supervision experiments are reported in Table 5. Overall, the scores highlight that (i) both variants of supervision are beneficial to the MMT models, which now show more consistent improvements over both the baselines and the MMT models with unsupervised visual attention and (ii) the **FINE-TUNED** model produces consistently higher scores, with improvements of up to 1.9 METEOR points (EN-FR 2016) with respect to the unimodal baseline.

Interestingly, the **FINE-TUNED** variant always outperforms the **SCRATCH** models. One possible explanation for this is that training the MMT model from scratch with both objectives might bias the encoder towards producing representations that align well with image regions, but not necessarily optimal for the MT decoder. To get a better insight on this, we performed a small study where we repeated the supervision experiments by lowering the alignment loss weight from $\beta = 1$ to $\beta \in \{0.1, 0.5\}$. The results showed that although there are cases where tuning the coefficient yields improved scores for some language pairs, the

		EN-Cs	EN-DE			EN-Fr		
		2016	2016	2017	CC	2016	2017	CC
$k=1$	BASILINE NMT	22.4	48.8	43.2	39.6	63.0	58.0	53.7
	+ IMAGES	21.9	48.8	43.3	39.7	64.5	58.2	54.2
	+ SUP (SCRATCH)	21.8	48.8	43.1	39.3	64.3	58.6	53.8
	+ SUP (FINE-TUNED)	22.5	49.2	43.5	39.9	64.9	58.9	54.8
$k=2$	BASILINE NMT	24.1	52.4	46.3	42.4	68.1	62.5	58.4
	+ IMAGES	24.4	52.1	46.4	43.1	69.0	63.2	58.6
	+ SUP (SCRATCH)	24.4	52.3	46.1	42.3	69.2	63.1	57.9
	+ SUP (FINE-TUNED)	24.4	52.9	46.8	43.0	69.5	63.9	58.9
$k=3$	BASILINE NMT	26.5	53.9	47.8	43.6	71.5	65.2	60.1
	+ IMAGES	26.5	53.5	47.8	44.1	71.8	65.7	60.4
	+ SUP (SCRATCH)	26.6	53.7	47.7	43.1	71.9	65.4	59.7
	+ SUP (FINE-TUNED)	27.0	54.3	48.4	43.8	72.5	66.4	60.5

Table 5: The impact of attention supervision with respect to baseline NMT and MMT systems: METEOR scores are averages of three runs with different seeds. Best systems are indicated in **bold** typeface.

scores are not as consistent as $\beta = 1$ across different WAIT-K decodings and test sets. We leave further exploration of this to future work.

4.2.2 PREFIX TRANSLATION ACCURACY

Automatic evaluation metrics are often not very suitable for tasks such as image captioning or MMT, as they can fail to capture non-systematic & minor changes (Caglayan et al., 2020b). This issue is even more pronounced in simultaneous translation, where models are more likely to make mistakes early in the translation (due to limited source context), which could have a negative impact in the overall translation quality. In this section, we provide an analysis focused on the translation accuracy of source prefixes. Specifically, we count the number of matches between the first n words of each translation and its reference sentence, and divide it by the number of test set sentences.

Figure 6 shows the unigram prefix accuracy of each WAIT-K decoded simultaneous MT model, across the three language pairs explored. The results reveal several things: (i) the supervision is globally beneficial, although its contribution diminishes as k increases, and (ii) although not reflected by the METEOR scores (Table 5), the incorporation of the visual modality (and further supervision) boosts the accuracies for **English**→**Czech** substantially (especially for WAIT-1). For the latter, this amounts to a 25% relative improvement by the FINE-TUNED supervision with respect to the baseline NMT. Similar trends were observed for bigram and trigram accuracy for **English**→**Czech** and **English**→**French**.




Figure 6: Unigram prefix accuracy of simultaneous WAIT-K variants: the numbers are obtained from the `test2016` test set, averaged across three runs.

4.2.3 COMPARISON TO STATE-OF-THE-ART

Finally, we conclude our quantitative analyses by comparing the performance of our best MMT models to state-of-the-art in simultaneous multimodal MT. Table 6 provides a summary of the results in terms METEOR and also BLEU (Papineni et al., 2002), since some previous work only reports the latter. Recall that (i) Imankulova et al. (2020) and Caglayan et al. (2020a) rely on recurrent models rather than Transformers, (ii) use different visual features and (iii) train their models using either consecutive or WAIT-K approaches rather than our best *prefix training* methodology.

According to Table 6, the scores clearly show that when combined with the FINE-TUNED attention supervision variant, our Transformer-based MMT models achieve state-of-the-art BLEU and METEOR scores across the board, with the exception of English→German with larger k . For this language direction, our WAIT-2 and WAIT-3 systems slightly lag behind Caglayan et al. (2020a) in terms of BLEU. This is probably due to our reliance on METEOR for early-stopping and best checkpoint selection rather than the perplexity, which was used by Caglayan et al. (2020a).

4.3 Qualitative Insights

In this section, we conduct further analysis focusing on qualitative aspects of the trained MMT models in terms of visual understanding.

4.3.1 GROUNDING ABILITY

We begin our analyses by measuring the grounding ability of our MMT systems on the `test2016` test set, for which we have ground-truth region-phrase annotations from Flickr30k dataset (§ 3.4). More formally, we obtain the cross-modal attention weights from our models during the translation decoding of the test set sentences. Next, similar to Rohrbach et al. (2016) and Wang and Specia (2019), we compute the intersection over union (IoU) between the most-attended region and the ground-truth bounding box. If multiple ground-truth boxes exist for a given phrase, we take the ground-truth box that yields the maximum IoU score. Table 7 provides a summary of the results across the language pairs explored.

		EN-Cs		EN-DE		EN-FR	
		BL	MT	BL	MT	BL	MT
$k=1$	Imankulova et al. (2020)	9.1	-	19.9	-	32.5	-
	Caglayan et al. (2020a) (DEC ATT)	-	-	21.2	45.9	42.1	64.5
	Caglayan et al. (2020a) (ENC ATT)	-	-	21.3	44.7	41.7	64.6
	OUR BEST	12.5	22.5	21.4	49.4	42.9	64.9
$k=2$	Imankulova et al. (2020)	-	-	-	-	-	-
	Caglayan et al. (2020a) (DEC ATT)	-	-	28.1	50.4	49.2	68.7
	Caglayan et al. (2020a) (ENC ATT)	-	-	27.6	49.7	49.2	69.1
	OUR BEST	16.1	24.4	28.0	52.9	51.5	69.5
$k=3$	Imankulova et al. (2020)	19.4	-	28.8	-	44.0	-
	Caglayan et al. (2020a) (DEC ATT)	-	-	32.2	53.7	54.6	71.8
	Caglayan et al. (2020a) (ENC ATT)	-	-	31.4	53.1	54.8	72.0
	OUR BEST	21.8	27.0	31.3	54.3	56.2	72.5

Table 6: Comparison of our best MMT system (fine-tuned supervised attention) with other state-of-the-art simultaneous MMT models on the `test2016` set. Best systems are indicated in **bold** typeface. BL and MT denote BLEU and METEOR, respectively.

Unsurprisingly, we notice that the MMT models with unsupervised visual attention obtain the lowest grounding accuracy across the board. The supervision from SCRATCH achieves the highest accuracy, surpassing the models with unsupervised visual attention by 18% on average. The FINE-TUNED supervision approach seems also quite helpful in terms of grounding, with accuracies slightly lagging (around 2%) behind the SCRATCH systems. Another interesting observation is that the grounding ability does not seem to depend on the choice of target language, e.g. all FINE-TUNED MMT models achieve around 89% accuracy. This is to be expected, since during training, the source sentences and the provided regions never change across different language directions.

4.3.2 PERFORMANCE ON BUTD REGIONS

The previous experiment relied on regional features that were extracted from the ground-truth bounding boxes at training, inference and evaluation. To understand whether the grounding ability generalises across features extracted from **non-gold** regions, we now switch to the 36 region proposals that we extracted from the pre-trained BUTD object detector (§ 3.3.1). Instead of working in the visual coordinate space, which may be noisy due to the fine-grained nature of bounding box annotations, given that BUTD also provides a wider range of object categories, we also measure the grounding ability in the **linguistic space**. Specifically, instead of the coordinates of the BUTD and ground-truth regions, we use the predicted object label \hat{y} (from the BUTD detector) and the ground-truth noun annotation y , respectively. We then compute exact match accuracy (EM) and cosine similarity⁶

6. For the cosine similarity approach, we ensure maximum correspondence by lemmatising both words using the NLTK toolkit (Bird & Loper, 2004).

	EN-Cs	EN-DE	EN-Fr
MMT	73.8% (± 0.7)	72.1% (± 0.8)	75.4% (± 1.0)
+ SUP (FINE-TUNED)	89.8% (± 0.6)	89.4% (± 0.7)	89.0% (± 2.0)
+ SUP (SCRATCH)	92.2% (± 0.2)	92.3% (± 0.1)	91.9% (± 0.2)

Table 7: Grounding ability of WAIT-1 simultaneous MMT systems using features from the ground-truth region annotations. The scores presented are intersection over unions (IoU) on the `test2016` set, averaged across three runs.

	EN-Cs			EN-DE			EN-Fr		
	IoU	Cos	EM	IoU	Cos	EM	IoU	Cos	EM
MMT	23.1%	0.416	9.8%	20.6%	0.403	8.4%	28.5%	0.488	18.3%
+ SUP (FINE-TUNED)	42.5%	0.581	26.6%	42.0%	0.578	26.1%	41.4%	0.572	25.4%
+ SUP (SCRATCH)	44.3%	0.593	27.6%	44.3%	0.590	27.4%	44.4%	0.589	27.2%

Table 8: Grounding ability of WAIT-1 simultaneous MMT systems using the 36 BUTD region proposals: IoU, Cos and EM denote intersection-over-union, cosine similarity and exact match accuracy, respectively. The scores are obtained from the `test2016` test set and averaged across three runs.

(Cos) between the lowercased versions of these words, by using pre-trained GloVe (Pennington et al., 2014) word embeddings.

Table 8 reports all three metrics averaged across three runs of each model. We observe that all three metrics show the same trend across model types. Regardless of the variant type, the incorporation of attention supervision dramatically improves the ratio of exact matches by up to 17% in absolute terms. We also see that, unlike the previous experiment where the models were provided the human-labelled regions (Table 7), here we have at least the **English**→**French** MMT model with unsupervised visual attention which obtains much better grounding scores than the other two language directions. Now that the models are given way more regions (36) than the human-labelled scenario (~ 4.3), we think these results are more representative for the final grounding ability of these models. As expected, these results are a lot lower than Table 7 since (i) we had a lot fewer bounding boxes per image in that case and (ii), the BUTD bounding boxes are different from the gold ones, meaning that even if the model focused on a similar part of the image, the IoU score will be lower than 1.

4.3.3 VISUALISING THE ATTENTION

Our main motivation to supervise the visual attention was to guide the model on learning cross-modal visual-linguistic alignments. After quantifying their ability to do so using metrics such as IoU and cosine similarity, we now visualise the impact of supervision on two real examples. Figure 7 shows that the MMT models with unsupervised visual attention (on the left) are already good at attending on relevant regions to some extent. However,

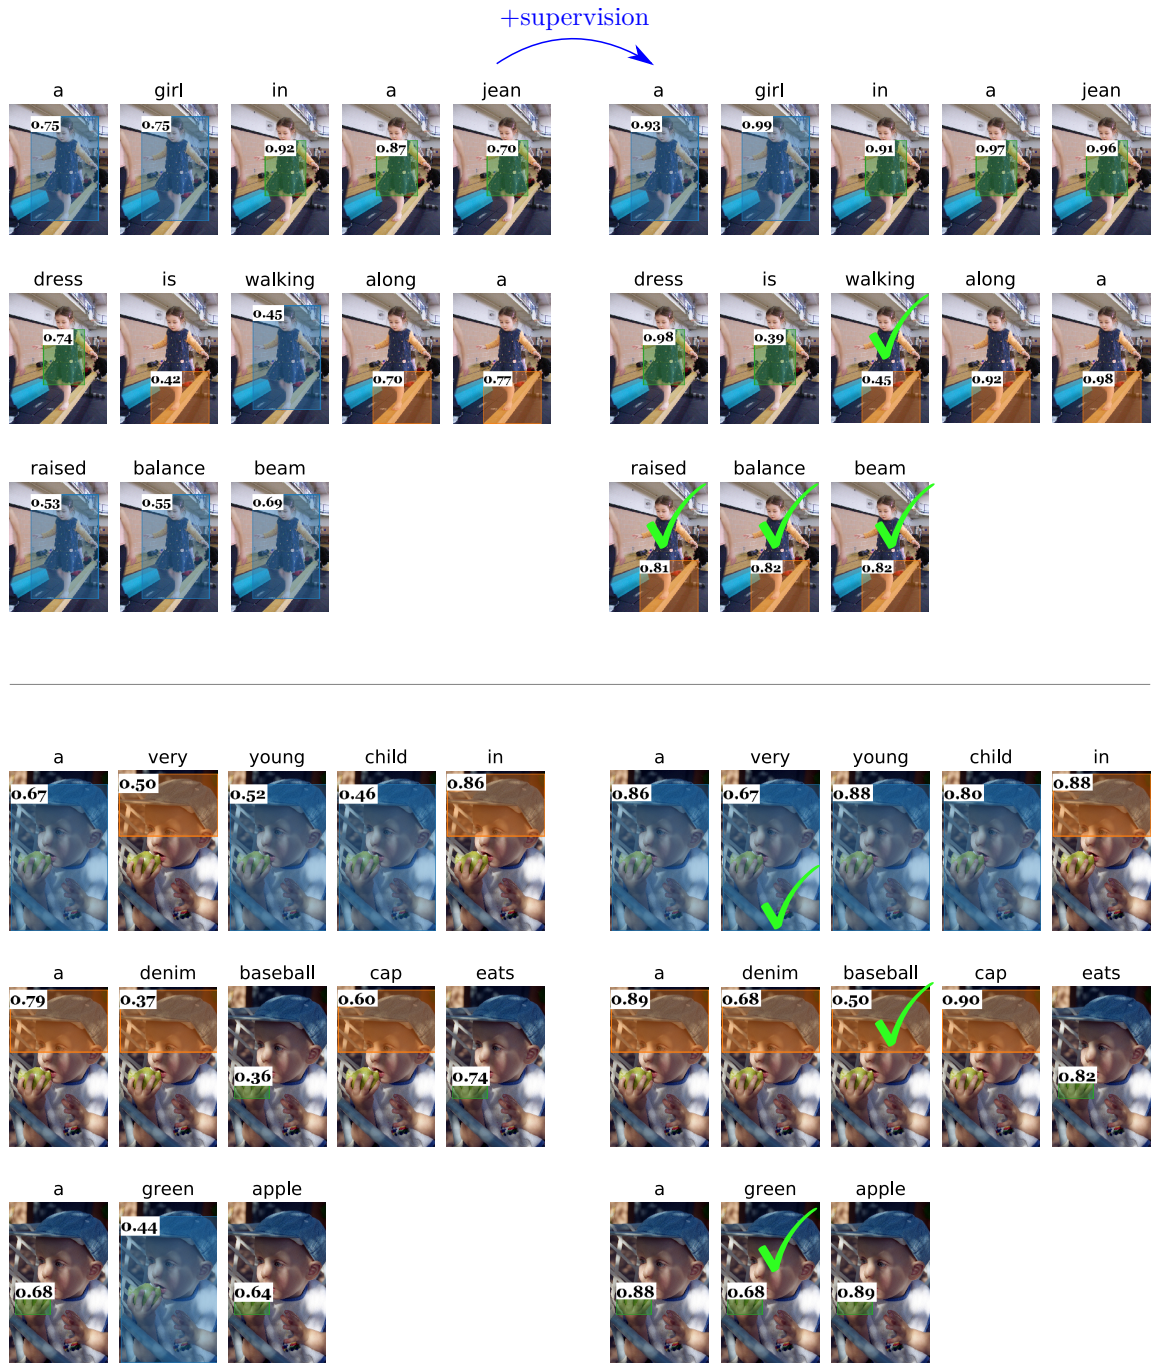


Figure 7: Regional attention examples produced by the MMT model with unsupervised visual attention (left) and the FINE-TUNED variant (right): Green check marks show cases where supervision fixes incorrect word & region alignments.

there are still cases of misalignment: in the first example, none of the words in the phrase “*raised balance beam*” are aligned with the actual region showing the *beam*. Similar skews also happen in the second example such as the word *baseball* getting aligned with the region *apple*. The FINE-TUNED variants (on the right) not only fix these cases of misalignments but also provide much more confident attention distributions. This is particularly evident in the attention probabilities of nouns which, unlike connectives or verbs, have corresponding ground-truth bounding boxes.

4.4 Translation Examples

Finally, we now present qualitative translation examples from our simultaneous MT systems, to highlight both the strengths and weaknesses of the models explored. Figure 8 shows an example for each language direction, using the WAIT-1 decoded NMT, MMT and supervised MMT systems. For the first example in **Czech**, we see that the NMT system hopelessly relies on dataset biases to always generate “Muž (man)” when the definite article “a” is the only observed word at source side. Although the MMT with unsupervised visual attention

	SRC	a woman is playing volleyball.
	NMT	muž se dívá na volejbal. (man looks at volleyball.)
	MMT	žena v bílém oblečení hraje volejbal. (woman in white clothes playing volleyball.)
	SUP	žena jde po ulici hraje volejbal. (woman goes down the street playing volleyball.)
	SUP+	žena, která hraje volejbal. (woman playing volleyball.)
	SRC	a woman with long hair is at a graduation ceremony.
	NMT	ein frau mit langen haaren sitzt bei einem olympischen spielen bei einer konferenz . (a woman with long hair sitting at an olympic game at a conference.)
	MMT	ein frau mit langen haaren ist an einem partner-konzert . (a woman with a long hair is in a partner concert.)
	SUP	eine frau mit langen haaren hält sich an einer konferenz vor. (a woman with a long hair is sitting in front of a conference.)
	SUP+	eine frau mit langen haaren steht an einem hutscher-bahnhof . (a woman with a long hair stands at a German railway.)
	SRC	a woman is holding a small white statue.
	NMT	un femme est en train de tenir une petite statue blanche . (a woman is holding a small white statue.)
	MMT	un femme fait de la petite femme blanche . (a woman makes the little white woman.)
	SUP	une femme fait un signe à un petit homme blanc . (a woman makes a sign to a small white man.)
	SUP+	une femme est en train de tenir une petite statue blanche . (a woman is holding a small white statue.)

Figure 8: WAIT-1 translation examples from all model variants: SUP and SUP+ denote the SCRATCH and FINE-TUNED variants of attention supervision, respectively. **blue** and **red** colors indicate correct and wrong lexical choices. “Google Translate” output for each non-English hypothesis is shown in grey.

and the supervision from scratch (SUP) both pick the correct translation “Žena (woman)”, they somewhat hallucinate the continuation. In contrast, the fine-tuned supervision (SUP+) generates the correct translation without any errors. For the **German** example, all models struggle to translate “graduation ceremony” into German. Interestingly, they all pick words that refer to places such as “concert” or “conference” rather than purely random lexical choices, hinting at the fact that the models may be relying more on visual information than the language signal. In addition, NMT and the non-supervised MMT also commit an error when translating “a” into German i.e. they both pick the masculine form “ein” instead of “eine”. The supervised models are better to integrate visual features in this case, as they correctly pick “eine”. Finally, the **French** example shows similar patterns to the German one as well: only the supervised models are able to pick the correct article form “une” and it is the fine-tuned variant that translates the whole sentence in a correct way. If we further look at examples of WAIT-3 simultaneous MT models and contrast the SCRATCH and FINE-TUNED variants (Figure 9), we observe that both variants are more or less equivalent when



Figure 9: WAIT-3 translation examples comparing two variants of attention supervision: SUP and SUP+ denote the SCRATCH and FINE-TUNED variants, respectively. blue and red colors indicate correct and wrong lexical choices. “Google Translate” output for each non-English hypothesis is shown in grey.

translating the early parts of the sentences but the SCRATCH models commit critical errors in overall, when the final translations are taken into account.

5. Conclusions

In this paper, we proposed the first Transformer-based simultaneous multimodal MT model and extended it with an auxiliary supervision signal that guides the visual attention mechanism using ground-truth phrase-region alignments. We performed extensive experiments on the three language pairs of the Multi30k dataset and upon conducting thorough quantitative and qualitative analyses, we showed that (i) supervised visual attention consistently improves the translation quality of the MMT models, and (ii) fine-tuning the MMT with the supervision loss results in better performance than training the MMT from scratch.

Acknowledgments

This article follows from the MSc. Thesis by Veneta Haralampieva, co-supervised by Lucia Specia and Ozan Caglayan. Lucia Specia and Ozan Caglayan were supported by the MultiMT project (H2020 ERC Starting Grant No. 678017). Lucia Specia also received support from the Air Force Office of Scientific Research (under award number FA8655-20-1-7006).

References

- Alinejad, A., Siahbani, M., & Sarkar, A. (2018). Prediction improves simultaneous neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3022–3027, Brussels, Belgium. Association for Computational Linguistics.
- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., & Zhang, L. (2018). Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*.
- Arivazhagan, N., Cherry, C., Macherey, W., Chiu, C.-C., Yavuz, S., Pang, R., Li, W., & Raffel, C. (2019). Monotonic infinite lookback attention for simultaneous machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1313–1323.
- Arivazhagan, N., Cherry, C., Macherey, W., & Foster, G. (2020). Re-translation versus streaming for simultaneous translation. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pp. 220–227, Online. Association for Computational Linguistics.
- Arthur, P., Cohn, T., & Haffari, G. (2021). Learning coupled policies for simultaneous machine translation using imitation learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 2709–2719, Online. Association for Computational Linguistics.
- Ba, J. L., Kiros, J. R., & Hinton, G. E. (2016). Layer normalization. *arXiv preprint arXiv:1607.06450*, 1(1).

- Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference on Learning Representations*.
- Bangalore, S., Rangarajan Sridhar, V. K., Kolan, P., Golipour, L., & Jimenez, A. (2012). Real-time incremental speech-to-speech translation of dialogs. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 437–445, Montréal, Canada. Association for Computational Linguistics.
- Barrault, L., Bougares, F., Specia, L., Lala, C., Elliott, D., & Frank, S. (2018). Findings of the third shared task on multimodal machine translation. In *Proceedings of the Third Conference on Machine Translation, Volume 2: Shared Task Papers*, pp. 308–327, Belgium, Brussels. Association for Computational Linguistics.
- Bird, S., & Loper, E. (2004). NLTK: The natural language toolkit. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pp. 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Bub, T., Wahlster, W., & Waibel, A. (1997). Verbmobil: The combination of deep and shallow processing for spontaneous speech translation. In *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 1, pp. 71–74. IEEE.
- Caglayan, O. (2019). *Multimodal Machine Translation*. Theses, Université du Maine.
- Caglayan, O., Aransa, W., Bardet, A., García-Martínez, M., Bougares, F., Barrault, L., Masana, M., Herranz, L., & van de Weijer, J. (2017). LIUM-CVC Submissions for WMT17 Multimodal Translation Task. In *Proceedings of the Second Conference on Machine Translation*, pp. 432–439.
- Caglayan, O., Aransa, W., Wang, Y., Masana, M., García-Martínez, M., Bougares, F., Barrault, L., & van de Weijer, J. (2016). Does multimodality help human and machine for translation and image captioning?. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pp. 627–633, Berlin, Germany. Association for Computational Linguistics.
- Caglayan, O., Ive, J., Haralampieva, V., Madhyastha, P., Barrault, L., & Specia, L. (2020a). Simultaneous machine translation with visual context. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2350–2361, Online. Association for Computational Linguistics.
- Caglayan, O., Madhyastha, P., & Specia, L. (2020b). Curious case of language generation evaluation metrics: A cautionary tale. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 2322–2328, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Calixto, I., Elliott, D., & Frank, S. (2016). DCU-UvA multimodal MT system report. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pp. 634–638.
- Calixto, I., & 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*, pp. 992–1003.

- Calixto, I., Liu, Q., & 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 (Volume 1: Long Papers)*, pp. 1913–1924.
- Chen, X., Fang, H., Lin, T.-Y., Vedantam, R., Gupta, S., Dollár, P., & Zitnick, C. L. (2015). Microsoft COCO captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 1.
- Cho, K., & Esipova, M. (2016). Can neural machine translation do simultaneous translation?. *arXiv preprint arXiv:1606.02012*, 1.
- Dalvi, F., Durrani, N., Sajjad, H., & Vogel, S. (2018). Incremental decoding and training methods for simultaneous translation in neural machine translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pp. 493–499, New Orleans, Louisiana. Association for Computational Linguistics.
- Delbrouck, J.-B., & Dupont, S. (2017). Modulating and attending the source image during encoding improves multimodal translation. *arXiv preprint arXiv:1712.03449*, 1(1).
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee.
- Denkowski, M., & Lavie, A. (2014). Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pp. 376–380. Association for Computational Linguistics.
- Elbayad, M., Besacier, L., & Verbeek, J. (2020). Efficient Wait-k Models for Simultaneous Machine Translation. In *Proc. Interspeech 2020*, pp. 1461–1465.
- Elliott, D., Frank, S., Barrault, L., Bougares, F., & Specia, L. (2017). Findings of the second shared task on multimodal machine translation and multilingual image description. In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*, pp. 215–233, Copenhagen, Denmark. Association for Computational Linguistics.
- Elliott, D., Frank, S., Sima'an, K., & Specia, L. (2016). Multi30K: Multilingual English-German image descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, pp. 70–74, Berlin, Germany. Association for Computational Linguistics.
- Elliott, D., & Kádár, Á. (2017). Imagination improves multimodal translation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 130–141, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Garg, S., Peitz, S., Nallasamy, U., & Paulik, M. (2019). Jointly learning to align and translate with transformer models. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong.
- Gu, J., Neubig, G., Cho, K., & Li, V. O. (2017). Learning to translate in real-time with neural machine translation. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pp. 1053–1062, Valencia, Spain. Association for Computational Linguistics.

- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Imankulova, A., Kaneko, M., Hirasawa, T., & Komachi, M. (2020). Towards multimodal simultaneous neural machine translation. In *Proceedings of the Fifth Conference on Machine Translation*, pp. 594–603, Online. Association for Computational Linguistics.
- Ive, J., Li, A. M., Miao, Y., Caglayan, O., Madhyastha, P., & Specia, L. (2021). Exploiting multimodal reinforcement learning for simultaneous machine translation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 3222–3233, Online. Association for Computational Linguistics.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 1.
- Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., & Herbst, E. (2007). Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pp. 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.-J., Shamma, D. A., et al. (2017). Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1), 32–73.
- Libovický, J., & Helcl, J. (2017). Attention strategies for multi-source sequence-to-sequence learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 196–202.
- Libovický, J., Helcl, J., & Mareček, D. (2018). Input combination strategies for multi-source transformer decoder. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 253–260.
- Liu, L., Utiyama, M., Finch, A., & Sumita, E. (2016). Neural machine translation with supervised attention. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 3093–3102, Osaka, Japan. The COLING 2016 Organizing Committee.
- Lu, J., Batra, D., Parikh, D., & Lee, S. (2019). Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pp. 13–23.
- Lu, J., Yang, J., Batra, D., & Parikh, D. (2016). Hierarchical question-image co-attention for visual question answering. In *Advances in neural information processing systems*, pp. 289–297.
- Ma, M., Huang, L., Xiong, H., Zheng, R., Liu, K., Zheng, B., Zhang, C., He, Z., Liu, H., Li, X., et al. (2019). STACL: Simultaneous Translation with Implicit Anticipation and

- Controllable Latency using Prefix-to-Prefix Framework. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3025–3036.
- Ma, X., Pino, J. M., Cross, J., Puzon, L., & Gu, J. (2020). Monotonic multihead attention. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Mi, H., Wang, Z., & Ittycheriah, A. (2016). Supervised attentions for neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2283–2288, Austin, Texas. Association for Computational Linguistics.
- Niehues, J., Pham, N.-Q., Ha, T.-L., Sperber, M., & Waibel, A. (2018). Low-latency neural speech translation. In *Proc. Interspeech 2018*, pp. 1293–1297.
- Nishihara, T., Tamura, A., Ninomiya, T., Omote, Y., & Nakayama, H. (2020). Supervised visual attention for multimodal neural machine translation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 4304–4314, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Papineni, K., Roukos, S., Ward, T., & Zhu, W.-J. (2002). BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pp. 311–318. Association for Computational Linguistics.
- Pascanu, R., Gulcehre, C., Cho, K., & Bengio, Y. (2014). How to construct deep recurrent neural networks: Proceedings of the second international conference on learning representations (iclr 2014). In *2nd International Conference on Learning Representations, ICLR 2014*.
- Pennington, J., Socher, R., & Manning, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Plummer, B. A., Wang, L., Cervantes, C. M., Caicedo, J. C., Hockenmaier, J., & Lazebnik, S. (2015). Flickr30k Entities: Collecting Region-to-Phrase Correspondences for Richer Image-to-Sentence Models. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 2641–2649.
- Press, O., & Wolf, L. (2017). Using the output embedding to improve language models. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 157–163.
- Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pp. 91–99.
- Rohrbach, A., Rohrbach, M., Hu, R., Darrell, T., & Schiele, B. (2016). Grounding of textual phrases in images by reconstruction. In *European Conference on Computer Vision*, pp. 817–834. Springer.

- Ryu, K., Matsubara, S., & Inagaki, Y. (2006). Simultaneous English-Japanese spoken language translation based on incremental dependency parsing and transfer. In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pp. 683–690, Sydney, Australia. Association for Computational Linguistics.
- Satija, H., & Pineau, J. (2016). Simultaneous machine translation using deep reinforcement learning. In *ICML 2016 Workshop on Abstraction in Reinforcement Learning*.
- Specia, L., Frank, S., Sima'an, K., & Elliott, D. (2016). A shared task on multimodal machine translation and crosslingual image description. In *Proceedings of the First Conference on Machine Translation*, pp. 543–553, Berlin, Germany. Association for Computational Linguistics.
- Specia, L., Wang, J., Jae Lee, S., Ostapenko, A., & Madhyastha, P. (2021). Read, spot and translate. *Machine Translation*, 35(1), 145–165.
- Sulubacak, U., Caglayan, O., Grönroos, S.-A., Rouhe, A., Elliott, D., Specia, L., & Tiedemann, J. (2020). Multimodal machine translation through visuals and speech. *Machine Translation*, 34(2), 97–147.
- Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826.
- Tan, H., & Bansal, M. (2019). LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 5100–5111, Hong Kong, China. Association for Computational Linguistics.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008.
- Wang, J., & Specia, L. (2019). Phrase localization without paired training examples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, Seoul, South Korea. IEEE.
- Wang, Q., Li, B., Xiao, T., Zhu, J., Li, C., Wong, D. F., & Chao, L. S. (2019). Learning deep transformer models for machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1810–1822, Florence, Italy. Association for Computational Linguistics.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4), 229–256.
- Young, P., Lai, A., Hodosh, M., & Hockenmaier, J. (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.

- Zheng, B., Zheng, R., Ma, M., & Huang, L. (2019). Simpler and faster learning of adaptive policies for simultaneous translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 1349–1354.
- Zhou, M., Cheng, R., Lee, Y. J., & 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*, pp. 3643–3653, Brussels, Belgium. Association for Computational Linguistics.