

# Simultaneous Machine Translation with Visual Context

Ozan Caglayan<sup>1</sup>, Julia Ive<sup>1</sup>, Veneta Haralampieva<sup>1</sup>, Pranava Madhyastha<sup>1</sup>

Loïc Barrault<sup>2</sup> and Lucia Specia<sup>1,2,3</sup>

Imperial College London<sup>1</sup>, University of Sheffield<sup>2</sup>, ADAPT - Dublin City University<sup>3</sup>

o.caglayan@ic.ac.uk, j.ive@ic.ac.uk, vlh19@ic.ac.uk, pranava@ic.ac.uk

l.barrault@sheffield.ac.uk, l.specia@ic.ac.uk

## Abstract

Simultaneous machine translation (SiMT) aims to translate a continuous input text stream into another language with the lowest latency and highest quality possible. The translation thus has to start with an incomplete source text, which is read progressively, creating the need for anticipation. In this paper, we seek to understand whether the addition of visual information can compensate for the missing source context. To this end, we analyse the impact of different multimodal approaches and visual features on state-of-the-art SiMT frameworks. Our results show that visual context is helpful and that visually-grounded models based on explicit object region information are much better than commonly used global features, reaching up to 3 BLEU points improvement under low latency scenarios. Our qualitative analysis illustrates cases where only the multimodal systems are able to translate correctly from English into gender-marked languages, as well as deal with differences in word order, such as adjective-noun placement between English and French.

## 1 Introduction

Simultaneous machine translation (SiMT) aims to reproduce human interpretation, where an interpreter translates spoken utterances as they are produced. The interpreter has to dynamically find the balance between how much context is needed to generate the translation reliably, and how long the listener has to wait for the translation. In contrast to *consecutive* machine translation where source sentences are available in their entirety before translation, the challenge in SiMT is thus the design of a strategy to find a good trade-off between the quality of the translation and the latency incurred in producing it. Previous work has considered *rule-based strategies* that rely on waiting until some




Figure 1: An illustration of a 1-word latency system that makes use of visual grounding to resolve the gender of the article ‘un’ and to predict the noun ‘chat’ after reading its qualifier ‘white’.  $\rightarrow$  and  $\rightarrow$  denote READ and WRITE, respectively.

constraint is satisfied, which includes approaches based on syntactic constraints (Bub et al., 1997; Ryu et al., 2006), segment/chunk/alignment information (Bangalore et al., 2012) heuristic-based conditions during decoding (Cho and Esipova, 2016) or deterministic policies with pre-determined latency constraints (Ma et al., 2019). An alternative line of research focuses on *learning* the decision policy: Gu et al. (2017) and Alinejad et al. (2018) frame SiMT as learning to generate READ/WRITE actions and employ reinforcement learning (RL) to formulate the problem as a policy agent interacting with its environment (i.e. a pre-trained MT model). Recent work has also explored supervised learning of the policy, by using oracle action sequences predicted by a pre-trained MT using confidence-based heuristics (Zheng et al., 2019) or external word aligners (Arthur et al., 2020) (details in §2).

Thus far, all prior research has focused on unimodal interpretation<sup>1</sup>. In this paper, we explore SiMT for multimodal machine translation (MMT) (Specia et al., 2016), where in addition to

<sup>1</sup>We note that Imankulova et al. (2020) also attempted to explore multimodality in SiMT. However, their paper overestimates the impact of visual cues, and in personal correspondence with the authors about the mismatch in the findings, they discovered critical bugs in their implementation.

the source sentence, we have access to visual information in the form of an image. We believe that having access to a complementary context should help the models anticipate the missing context (Figure 1) by grounding their decisions about ‘when’ and ‘what’ to translate. To test our hypothesis, we explore heuristic-based decoding and fixed-latency wait- $k$  policy<sup>2</sup>, and investigate the effectiveness of different visual representations (§3). Based on experiments on the Multi30k dataset (Elliott et al., 2016) (§4), we provide quantitative and qualitative analyses for English→German and English→French (§5).

Our **contributions** highlight that: (i) visual context offers up to 3 BLEU points improvement for low-latency wait- $k$  policies, and consistently lowers the latency for *wait-if-diff* (Cho and Esipova, 2016) decoding, (ii) explicit object region features are more expressive than commonly used global visual features, (iii) training wait- $k$  MMTs offers remarkably better grounding capabilities than decoding-only wait- $k$  for linguistic phenomena such as gender resolution and adjective-noun ordering, and (iv) with twice the runtime speed of decoder-based visual attention, the encoder-based grounding is promising for application scenarios.

## 2 Related Work

### 2.1 Multimodal Machine Translation (MMT)

MMT aims to improve the quality of automatic translation using auxiliary sources of information (Sulubacak et al., 2020). The most typical framework explored in previous work makes use of the images when translating their descriptions between languages, with the hypothesis that visual grounding could provide contextual cues to resolve linguistic phenomena such as word-sense disambiguation or gender marking.

Existing work often rely on the use of visual features extracted from state-of-the-art CNN models pre-trained on large-scale visual tasks. The methods can be grouped into two branches depending on the feature type used: (i) multimodal attention (Calixto et al., 2016; Caglayan et al., 2016; Libovický and Helcl, 2017; Delbrouck and Dupont, 2017) which implements a soft attention (Bahdanau

et al., 2014) over *spatial* feature maps, and (ii) multimodal interaction between a *pooled* visual feature vector and linguistic representations (Calixto and Liu, 2017; Caglayan et al., 2017a; Elliott and Kádár, 2017; Grönroos et al., 2018).

### 2.2 Simultaneous Neural MT

Simultaneous NMT was first explored by Cho and Esipova (2016) in a greedy decoding framework where heuristic waiting criteria are used to decide whether the model should read more source words or emit a target word. Gu et al. (2017) instead utilised a pre-trained NMT model in conjunction with a reinforcement learning agent whose goal is to learn a READ/WRITE policy by maximising quality and minimising latency. Alinejad et al. (2018) further extended the latter approach by adding a PREDICT action whose purpose is to anticipate the next source word.

A common property of the above approaches is their reliance on *consecutive* NMT models pre-trained on full-sentences. Dalvi et al. (2018) pointed out a potential mismatch between the training and decoding regimens of such approaches and proposed fine-tuning the models using chunked data or prefix pairs. Ma et al. (2019) proposed an end-to-end, fixed-latency framework called ‘wait- $k$ ’ which allows *prefix-to-prefix* training using a deterministic policy: the agent starts by reading a specified number of source tokens ( $k$ ), followed by alternating WRITE and READ actions. Arivazhagan et al. (2019) extended the wait- $k$  framework using an advanced attention mechanism and optimising a differential latency metric (DAL). Recently, Arivazhagan et al. (2020) explored a radically different approach which enriches full-sentence training with prefix pairs (Niehues et al., 2018) and allows *re-translation* of previously committed target tokens to increase the translation quality.

Another line of research focuses on learning adaptive policies in a supervised way by using oracle READ/WRITE actions generated with heuristic or alignment-based approaches. Zheng et al. (2019) extracted action sequences from a pre-trained NMT model with a confidence-based heuristic and used them to train a separate policy network while Arthur et al. (2020) explored jointly training the translation model and the policy with oracle sequences obtained from a word alignment model.

<sup>2</sup>During our initial experiments we also explored the RL-based SiMT policy (Gu et al., 2017) but could not find good hyper-parameter settings, especially settings which were stable across two language pairs. Therefore, we did not proceed with RL for multimodal SiMT.

### 3 Methods

In this section, we first describe the underlying NMT architectures and baseline simultaneous MT approaches, to then introduce the proposed multimodal extensions to SiMT.

#### 3.1 Baseline NMT

Our *consecutive* baseline consists of a 2-layer GRU (Cho et al., 2014) encoder and a 2-layer Conditional GRU decoder (Sennrich et al., 2017) with attention (Bahdanau et al., 2014). The encoder is **unidirectional** as source sentences would be progressively read<sup>3</sup>. Given a source sequence of embeddings  $X=\{x_1, \dots, x_S\}$  and a target sequence of embeddings  $Y=\{y_1, \dots, y_T\}$ , the encoder first computes the sequence of hidden states  $H=\{h_1, \dots, h_S\}$ . At a given timestep  $t$  of decoding, the output layer estimates the probability of the next target word  $y_t$  as follows:

$$\begin{aligned} d_t &= \text{GRU}(y_{t-1}, d'_{t-1}) \\ c_t &= \text{Attention}(H, \text{query} \leftarrow d_t) \quad (1) \\ d'_t &= \text{GRU}'(c_t, d_t) \\ o_t &= \tanh(\mathbf{W}_c c_t + \mathbf{W}_d d'_t + \mathbf{W}_y y_{t-1}) \\ l_t &= \mathbf{W}_o(\mathbf{W}_b o_t + b_b) + b_o \end{aligned}$$

$$P(y_t|X, Y_{<t}) = \text{softmax}(l_t)$$

For a single training sample, we then maximise the joint likelihood of source and target sentences:

$$L(X, Y) = \sum_{t=1}^T \log(P(y_t|X_{\leq g(t)}, Y_{<t})) \quad (2)$$

Following the formulation of Ma et al. (2019),  $g(t)$  in equation 2 is a function which returns the number of source tokens encoded so far when predicting the target token  $y_t$ . In the case of consecutive NMT, since all source tokens are observed before the decoder runs,  $g(t)$  is equal to the length of the source sentence i.e.  $g(t) = |X|$ .

#### 3.2 Incorporating the visual modality

We consider a setting where the visual context is static and is available in its entirety at encoding time. This is a realistic setting in many applications, for example, the simultaneous translation of news, where images (or video frames) are shown before the whole source stream is available. We consider the following ways of integrating visual information.

<sup>3</sup>Although it is possible to encode the growing prefixes bidirectionally, it would incur a quadratic complexity.

**Object classification (OC) features** are *global* image information extracted from convolutional feature maps, which are believed to capture spatial cues. These *spatial* features are extracted from the final convolution layer of a ResNet-50 CNN (He et al., 2016) trained on ImageNet (Deng et al., 2009) for object classification. An image is represented by a feature tensor  $V \in \mathbb{R}^{8 \times 8 \times 2048}$ .

**Object detection (OD) features** are explicit object information where *local* regions in an image detected as objects are encoded by *pooled* feature vectors. These features are generated by the ‘‘bottom-up-top-down (BUTD)’’ (Anderson et al., 2018) extractor<sup>4</sup> which is a Faster R-CNN/ResNet-101 object detector (with 1600 object labels) trained on the Visual Genome dataset (Krishna et al., 2017). For a given image, the detector provides 36 object region proposals and extracts a pooled feature vector from each. An image is thus represented by a feature tensor  $V \in \mathbb{R}^{36 \times 2048}$ . We hypothesise that explicit object information can result in better referential grounding by using conceptually meaningful units rather than global features.

#### 3.3 Multimodal architectures

**Decoder attention (DEC-OC/OD).** A standard way of integrating visual modality into NMT is to apply a secondary attention at each decoding timestep (Calixto et al., 2016; Caglayan et al., 2016). We follow this approach to construct an MMT baseline. Specifically, equation 1 is extended so that the decoder attends to both the source hidden states  $H$  (eq. 3) and the visual features  $V$  (eq. 4), and they are added together to form the *multimodal* context vector  $c_t$  (eq. 5):

$$c_t^{\mathbf{T}} = \text{Attention}^{\mathbf{T}}(H, \text{query}=d_t) \quad (3)$$

$$c_t^{\mathbf{V}} = \text{Attention}^{\mathbf{V}}(V, \text{query}=d_t) \quad (4)$$

$$c_t = c_t^{\mathbf{T}} + c_t^{\mathbf{V}} \quad (5)$$

**Multimodal encoder (ENC-OD).** Instead of integrating the visual modality into the decoder, we propose to ground the source sentence representation within the encoder similar to Delbrouck and Dupont (2017). We hypothesise that early visual integration could be more appropriate for SiMT to fill in the missing context. Our approach differs from Delbrouck and Dupont (2017) in the use of *scaled-dot* attention (Vaswani et al., 2017) and **object detection (OD)** features. The attention layer

<sup>4</sup><https://hub.docker.com/r/airsplay/bottom-up-attention>

receives unidirectional hidden states  $H$  (for source states that were encoded/read so far) as the *query* and the visual features  $V$  as *keys* and *values*, i.e. it computes a mixture  $M$  of region features based on the cross-modal relevance. The final representation that will be used as input to the equation 1 is defined as  $\text{LAYERNORM}(M + H)$  (Ba et al., 2016).

Regardless of the multimodal approach taken, all visual features are first linearly projected into the dimension of textual representations  $H$ . To make modality representations compatible in terms of magnitude statistics, we apply layer normalisation (Ba et al., 2016) on textual representations  $H$  and the previously projected visual representations  $V$ . A dropout (Srivastava et al., 2014) of  $p = 0.5$  follows the layer normalisation.

### 3.4 Simultaneous MT approaches

This section summarises the SiMT approaches explored in this work: (i) the *heuristic-based* decoding approach **wait-if-diff** (Cho and Esipova, 2016), (ii) the *wait- $k$*  policy (Ma et al., 2019), and (iii) the reinforcement learning (RL) policy (Gu et al., 2017). The first approach offers a heuristically guided latency while the second one fixes it to an arbitrary value. The third one learns a stochastic policy to find the desired quality-latency trade-off. But before going into full details of methods, we now introduce the common metrics used to measure the latency of a given SiMT model.

#### 3.4.1 Latency metrics

**Average proportion** (AP) is the very first metric used for latency measurement in the literature (Cho and Esipova, 2016). AP computes a normalised score between 0 and 1, which is the average number of source tokens required to commit a translation:

$$\text{AP}(X, Y) = \frac{1}{|X||Y|} \sum_{t=1}^{|Y|} g(t)$$

AP produces different scores for 2 samples when the underlying latency is actually the same but the source and target sentence lengths differ. To remedy this, Ma et al. (2019) propose **Average Lagging** (AL) which estimates the number of tokens the “writer” is lagging behind the “reader”, as a function of the number of input tokens read.  $\tau$  denotes the timestep where the entire source sentence has been read, as the authors state that the

subsequent timesteps do not incur further delay:

$$\text{AL}(X, Y) = \frac{1}{\tau} \sum_{t=1}^{\tau} g(t) - \frac{t-1}{\gamma} \quad (\gamma = \frac{|Y|}{|X|})$$

Finally, **Consecutive Wait** (CW) (Gu et al., 2017) measures how many source tokens are consecutively read between committing two translations:

$$C_0 = 0$$

$$C_t = \begin{cases} C_{t-1} + 1 & \text{if action is READ} \\ 0 & \text{if action is WRITE} \end{cases}$$

#### 3.4.2 Wait- $k$ training

Ma et al. (2019) propose a simple deterministic policy which relies on training and decoding an NMT model in a *prefix-to-prefix* fashion. Specifically, a *wait- $k$*  model starts by reading  $k$  source tokens and writes the first target token. The model then reads and writes one token at a time to complete the translation process. This implies that the attention layer will now attend to a **partial textual representation**  $H_{\leq g(t)}$  instead of  $H$ , with  $g(t)$  redefined as  $\min(k + t - 1, |X|)$  (eq. 1 and 2).

**Decoding-only mode.** A *wait- $k$*  model is denoted as “trained” if it is both trained and decoded using the algorithm above. It is also possible to take a pre-trained consecutive NMT or MMT model, and apply *wait- $k$*  algorithm at decoding time i.e. during *greedy search*.

#### 3.4.3 Wait-if decoding

Cho and Esipova (2016) propose two decoding algorithms which can be directly applied on a pre-trained *consecutive* NMT model, similar to the consecutive *wait- $k$*  decoding. These algorithms have two hyper-parameters, namely the number of initial source tokens to read ( $k$ ) before starting the translation and the number of further tokens to read ( $\delta$ ) if the algorithm decides to wait for more context. We specifically use the **wait-if-diff (WID)** variant, which reads more tokens if the current most likely target word **changes** when doing so. We intentionally left out the **wait-if-worse (WIW)** approach as it exhibits very high latency.

#### 3.4.4 Reinforcement learning

Gu et al. (2017) frame SiMT as a sequence of READ or WRITE actions and aim to learn a reinforcement learning (RL) strategy with a reward function taking into account both quality and latency. Following standard RL, the framework is composed of



|       | Sents | En→De |       | En→Fr |       |
|-------|-------|-------|-------|-------|-------|
|       |       | T/S   | OOV   | T/S   | OOV   |
| TRAIN | 29K   | 0.96  | –     | 1.10  | –     |
| VAL   | 1014  | 0.98  | 30.4% | 1.09  | 19.0% |
| 2016  | 1000  | 0.94  | 28.6% | 1.09  | 18.1% |
| 2017  | 1000  | 0.97  | 37.1% | 1.14  | 21.3% |
| COCO  | 461   | 0.99  | 38.2% | 1.10  | 21.1% |

Table 1: Multi30k statistics: T/S and OOV are the average target-to-source sentence length ratio and the % of sentence pairs with at least 1 unknown token.

an environment and an agent. The environment is a pre-trained NMT system which is not *updated* during RL training. The agent is a GRU that parameterises a stochastic policy which decides on the action  $a_t$  by receiving as input the observation  $o_t$ . The observation  $o_t$  is defined as  $[c_t; d_t; y_t]$ , i.e. the concatenation of vectors coming from the environment. At each step, the agent receives a reward  $r_t = r_t^Q + r_t^D$  where  $r_t^Q$  is the quality reward (the difference of smoothed BLEU scores for partial hypotheses produced from one step to another) and  $r_t^D$  is the latency reward formulated as:

$$r_t^D = \alpha [\text{sgn}(C_t - C^*) + 1] + \beta [D_t - D^*]_+$$

where  $C_t$  denotes the CW metric introduced here to avoid long consecutive waits and  $D_t$  refers to AP (see § 3.4.1 for metrics).  $D^*$  and  $C^*$  are hyperparameters that determine the expected/target values for AP and CW, respectively. The optimal quality-latency trade-off is achieved by balancing the two reward terms.

## 4 Experimental Setup

### 4.1 Dataset

We use the Multi30k dataset (Elliott et al., 2016)<sup>5</sup> which has been the primary corpus for MMT research across the three shared tasks of the “Conference on Machine Translation (WMT)” (Specia et al., 2016; Elliott et al., 2017; Barrault et al., 2018). Multi30k extends the Flickr30k image captioning dataset (Young et al., 2014) to provide caption translations in German, French and Czech. In this work, we focus on the English→German and English→French (Elliott et al., 2017) language directions (Table 1). We use flickr2016 (2016), flickr2017 (2017) and coco2017 (COCO) for model evaluation. The latter test set is explicitly designed (Elliott et al., 2017) to contain at least

<sup>5</sup><https://github.com/multi30k/dataset>

one ambiguous word per sentence, which makes it appealing for MMT experiments.

**Preprocessing.** We use Moses scripts (Koehn et al., 2007) to lowercase, punctuation-normalise and tokenise the sentences with *hyphen splitting*. We then create word vocabularies on the *training* subset of the dataset. We did not use subword segmentation to avoid its potential side effects on SiMT and to be able to analyse the grounding capability of the models better. The resulting English, French and German vocabularies contain 9.8K, 11K and 18K tokens, respectively.

### 4.2 Reproducibility

**Hyperparameters.** The dimensions of embeddings and GRU hidden states are set to 200 and 320, respectively. The decoder’s input and output embeddings are shared (Press and Wolf, 2017). We use ADAM (Kingma and Ba, 2014) as the optimiser and set the learning rate and mini-batch size to 0.0004 and 64, respectively. A weight decay of  $1e-5$  is applied for regularisation. We clip the gradients if the norm of the full parameter vector exceeds 1 (Pascanu et al., 2013). For the **RL baseline**, we closely follow (Gu et al., 2017)<sup>6</sup>. The agent is implemented by a 320-dimensional GRU followed by a softmax layer and the baseline network – used for variance reduction of policy gradient – is similar to the agent except with a scalar output layer. We use ADAM as the optimiser and set the learning rate and mini-batch size to 0.0004 and 6, respectively. For each sentence pair in a batch, ten trajectories are sampled. For inference, greedy sampling is used to pick action sequences. We set the hyperparameters  $C^*=2$ ,  $D^*=0.3$ ,  $\alpha=0.025$  and  $\beta=-1$ . To encourage exploration, the negative entropy policy term is weighed empirically with 0.1 and 0.3 for En→Fr and En→De directions, respectively.

**Training.** We use `nmtpytorch` (Caglayan et al., 2017b) with PyTorch (Paszke et al., 2019) v1.4 for our experiments<sup>7</sup>. We train each model for a maximum of 50 epochs and early stop the training if validation BLEU (Papineni et al., 2002) does not improve for 10 epochs. We also halve the learning rate if no improvement is obtained for two epochs. On a single NVIDIA RTX2080-Ti GPU, it takes around 35 minutes for the *unimodal* and *multimodal encoder* variants to complete train-

<sup>6</sup><https://github.com/nyu-dl/dl4mt-simul-trans>

<sup>7</sup><https://github.com/ImperialNLP/pysimt>

|        | English→German |       |       | English→French |       |       | Avg    |
|--------|----------------|-------|-------|----------------|-------|-------|--------|
|        | 2016           | 2017  | COCO  | 2016           | 2017  | COCO  |        |
| NMT    | 34.6           | 26.4  | 22.1  | 57.8           | 50.3  | 41.4  | –      |
| ENC-OD | ↓ 0.6          | ↑ 0.3 | ↑ 0.8 | ↑ 0.3          | ↑ 0.1 | ↑ 1.1 | ↑ 0.33 |
| DEC-OC | ↑ 0.4          | ↑ 0.8 | ↑ 0.7 | ↓ 0.3          | ↑ 0.2 | ↑ 0.7 | ↑ 0.52 |
| DEC-OD | ↑ 0.7          | ↑ 1.0 | ↑ 1.7 | ↑ 0.1          | ↑ 0.6 | ↑ 1.2 | ↑ 0.88 |

Table 2: Multimodal gains in BLEU for *consecutive* baselines: the DEC-OD system exhibits the best average improvements.

ing whereas the *decoder attention* variant requires around twice that time. The number of learnable parameters is between 6.9M and 9.4M depending on the language pair and the type of multimodality. For the **RL baseline**, we choose the model that maximises the quality-to-latency ratio (BLEU/AL) on the validation set with patience set to ten epochs. The number of learnable parameters is around 6M.

### 4.3 Evaluation

To mitigate variance in results due to different initialisations, we repeat each experiment **three** times, with random seeds. Following previous work, we decode translations with *greedy search*, using the checkpoint that achieved the lowest perplexity. We report **average** BLEU scores across three runs using sacreBLEU (Post, 2018), which is also used for computing sentence-level scores for the oracle experiments.

## 5 Results

### 5.1 Consecutive baselines

We first present the impact of the visual integration approaches on *consecutive* NMT performance (Table 2). We observe that the decoder-attention using object detection features (DEC-OD) performs better than other variants. We also see that the improvements on flickr2017 ( $\uparrow 0.5$ ) and coco2017 ( $\uparrow 1.03$ ) test sets are higher than flickr2016 ( $\uparrow 0.1$ ) on average. A possible explanation is that flickr2017 and coco2017 are more distant from the training set distribution (higher OOV count, see Table 1) and thus there is more room for improvement with the visual cues. In summary, unlike previous conclusions in MMT where improvements were not found to be substantial (Grönroos et al., 2018; Caglayan et al., 2019), we observe that the benefit of the visual modality is more pronounced here. We believe that this is due to (i) the encoder being now *unidirectional* different from state-of-the-art NMT, (ii)




Figure 2: AL vs BLEU comparison across unimodal SiMT approaches: wait- $k$  systems are “trained”.

the modality representations being passed through layer normalisation (Ba et al., 2016), and (iii) the representational power of OD features.

### 5.2 Unimodal SiMT baselines

We now compare **unimodal** SiMT approaches to get an initial understanding of how they perform on Multi30k. Figure 2 contrasts AL and BLEU for three trained wait- $k$  systems, wait-if-diff (WID) decoding with  $k \in \{1, 2\}$  and  $\delta=1$ , reinforcement learning (RL) and the consecutive NMT. These configurations are chosen particularly to satisfy a low-latency regimen. The results suggest that wait- $k$  models offer good translation quality for fixed latency. The RL based policy, however, is not able to surpass wait- $k$  models. Finally, WID decoding exhibits the worst performance, according to BLEU. Given the difficulty in finding stable hyper-parameters for the RL models, we leave the integration of RL to MMT for future work and explore wait- $k$ , and WID approaches in what follows.

### 5.3 Wait- $k$ training for MMT

We present results with trained wait- $k$  MMTs with  $k \in \{1, 2, 3, 5, 7\}$ . Figure 3 plots a summary of the gains obtained by three MMT variants with respect to the unimodal wait- $k$ . We observe that as  $k$  increases, the gains due to the visual modality decrease globally, in line with the findings of Caglayan et al. (2019). This phenomenon is more visible for German, which exhibits 1 BLEU point




Figure 3: BLEU comparison of trained wait- $k$  MMT systems: the vertical axes of each subplot represent the improvement with respect to the unimodal wait- $k$ .

drop consistently across all models and two test sets at  $k=7$ . We hypothesise that this instability is probably due to the interplay of several factors for German, including the high OOV rates & rich morphology and source sentences being slightly longer than target on average. The latter is a major issue for trained wait- $k$  systems since the source sentences may not have been fully observed during training,<sup>8</sup> preventing the decoder to learn about source  $\langle \text{EOS} \rangle$  markers. For French, the results are much more encouraging as the improvements are larger and still observed with larger  $k$  values.

From a multimodal perspective, like with the consecutive models, the DEC-OD system has the best performance: it is beneficial for all values of  $k$  in French, and it shows the largest gains in German for  $k \in \{1, 2\}$ . From a runtime perspective, encoder-based attention benefits heavily from batched matrix operations and runs almost at the same speed as a unimodal NMT, thus encouraging us to focus more on that in the future.

**Qualitative examples.** Table 3 shows some examples regarding the impact of the visual modality for the wait-1 policy. In the first example, the image assists in predicting the correct article **eine** (feminine ‘a’) instead of **ein** (masculine ‘a’) in German. Upon inspecting the attention over object regions, we observe that the region that obtained the highest probability ( $p=0.2$ ) when predicting **eine** is labelled with ‘woman’ by the object detector. In the second example, we observe a biased anticipation

<sup>8</sup>Ma et al. (2019) proposed optional “catchup” logic for this, but we did not apply it here for the sake of simplicity.

|   |   |
|---|---|
| <p>SRC : a young brunette woman ...<br/> NMT : <u>ein junger</u> <b>brünnette</b> frau ...<br/> MMT : <b>eine</b> junge brünnette frau ...</p>                |  |
| <p>SRC : a black and white bird ...<br/> NMT : <b>un chien</b> (dog) <b>noir et blanc</b> ...<br/> MMT : <b>un oiseau</b> (bird) <b>noir et blanc</b> ...</p> |  |

Table 3: Examples showing the effectiveness of DEC-OD MMT (wait-1) for gender marking (top) and adjective noun placement (bottom). underlined and **bold** represent wrong and correct word choices, respectively.

case where the NMT system had to emit a wrong translation **chien** (‘dog’) before seeing the noun ‘bird’. However, the multimodal model successfully leveraged the visual context for anticipation and correctly handled the adjective-noun placement phenomenon. Once again, the attention distribution confirms that when generating the first two words – **un** and **oiseau** (‘bird’), the model correctly attends to the object regions corresponding to ‘bird’ (with  $p=0.22$  and  $p=0.14$  respectively).

#### 5.4 Trained vs. decoding-only SiMT

We are now interested in how trained wait- $k$  MMTs compare to decoding-only wait- $k$  and wait-if-diff (WID) heuristic under low latency. Figure 4 summarises latency vs. quality trade-off across all languages and test sets. First of all, the translation **quality** of the heuristic WID approach consistently improves with visual information with its latency slightly increasing across the board. Second, the translation **quality** of both trained and decoding-only wait- $k$  policies improve with multimodality.

Interestingly, although Ma et al. (2019) show that trained wait- $k$  models are substantially better than decoding-only ones for a news translation task, here we observe quite the opposite: in almost all cases there exists a shift between these approaches which favours the decoding-only approach for small  $k$  values. Zheng et al. (2019) observed a similar phenomenon for their wait-1 and wait-2 textual SiMT models. To investigate further, we compute an adaptive low-latency wait- $k$  **oracle** for  $k \in \{1, 2, 3\}$ . Specifically, for a given model, we first select a representative hypothesis across the three runs using median<sup>9</sup> sentence-level BLEU. We then pick the hypothesis with best BLEU (and

<sup>9</sup>We use median across the three runs of the same model as a way to smooth out variance related to random initialisations.

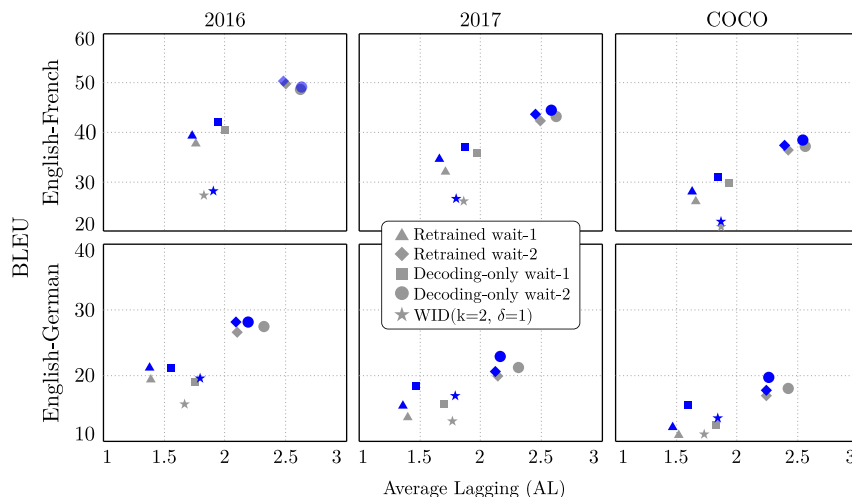


Figure 4: Comparison of trained vs. decoding-only SiMT approaches: light and dark colors denote unimodal and multimodal (DEC-OD) systems, respectively.

---

**Algorithm 1: Multi-run oracle algorithm**


---

**output** : Oracle BLEU  $\rightarrow$  CorpusBLEU( $O$ )  
**output** : Oracle Delay  $\rightarrow$  AverageLag( $O$ )  
1  $N$  : Number of test set sentences  
2  $B$  : Sentence BLEU scores across runs  
3  $C$  : Candidate list of {HYP, BLEU,  $k$ }  
4  $O$  : Final oracle list of  $N$  translations  
5  $O \leftarrow []$   
6 **for**  $n$  *in*  $1 \dots N$  **do**  
7      $C \leftarrow []$   
8     **for**  $k$  *in*  $\{1, 2, 3, 5, 7\}$  **do**  
9          $R \leftarrow$  wait- $k$  hypotheses for input  $n$   
10          $B \leftarrow [\text{BLEU}(R[i]) \forall i \in \{1, 2, 3\}]$   
11          $m \leftarrow$  Index of run with median BLEU  
12         HYP  $\leftarrow R[m]$ ; BLEU  $\leftarrow B[m]$   
13          $C.append(\{\text{HYP}, \text{BLEU}, k\})$   
14     **end**  
15      $h \leftarrow$  Best HYP from  $C$  (lowest  $k$  if tie)  
16      $O.append(h)$   
17 **end**

---

lowest  $k$  in case of a tie) across all wait- $k$  systems of that model, as the oracle translation. Once the hypotheses are collected, we compute corpus-level BLEU and AL (Algorithm 1).

Table 4 suggests that when we let the oracle choose between different  $k$  values, trained wait- $k$  systems are almost always better than decoding-only counterparts. Moreover, this boost in quality is accompanied by slight latency improvements over unimodal NMT across the board. Therefore, we conjecture that the shifts between decoding-only and trained wait- $k$  systems may be due to several factors coming into play such as the length discrepancy issue (§5.3) or the low resourced nature of Multi30k which prevents it from benefiting from prefix-to-prefix training for small  $k$  values.

|        | 2016                         | 2017                 | COCO                 |
|--------|------------------------------|----------------------|----------------------|
|        | English $\rightarrow$ German |                      |                      |
| NMT    | 33.5 (2.52)                  | 26.2 (2.46)          | 21.9 (2.61)          |
|        | 33.7 (2.30)                  | 25.6 (2.32)          | 22.4 ( <b>2.38</b> ) |
| DEC-OD | <b>34.4</b> (2.37)           | <b>27.8</b> (2.24)   | <b>23.8</b> (2.40)   |
|        | 34.3 ( <b>2.23</b> )         | 26.0 ( <b>2.20</b> ) | 22.5 (2.45)          |
|        | English $\rightarrow$ French |                      |                      |
| NMT    | 55.3 (2.78)                  | 48.9 (2.71)          | 41.4 (2.62)          |
|        | 56.9 (2.69)                  | 50.4 (2.69)          | 42.4 (2.68)          |
| DEC-OD | 55.9 (2.71)                  | 50.1 ( <b>2.58</b> ) | 42.3 (2.58)          |
|        | <b>57.9</b> (2.65)           | <b>51.3</b> (2.60)   | <b>43.1</b> (2.57)   |

Table 4: BLEU (AL) oracles for low-latency decoding-only (first line) and trained wait- $k$  (second line).

**Gender resolution accuracy.** Motivated by the qualitative examples in §5.3, we further investigate how accurate English $\rightarrow$ French MMT variants are at choosing the correct indefinite article **une** when translating sentences beginning with ‘a woman’. Table 5 shows that the unimodal NMT has no way of anticipating the context to resolve this kind of gender ambiguity, and therefore always picks the masculine version of the article. This is a clear example of models reflecting biases in the training data. In fact, 69.4% of all training instances starting with an indefinite article in French, have the masculine realisation of the article (**un**) instead of its feminine counterpart (**une**). The results also make it clear that decoding-only wait- $k$  systems are not as successful as their trained counterparts when it comes to incorporating the visual modality, and the explicit object information is more expressive



| wait-1 | DECODING ONLY | TRAINED   |
|--------|---------------|-----------|
| NMT    | 0             | 0         |
| DEC-OC | 10.7          | 49        |
| ENC-OD | 7.9           | 71        |
| DEC-OD | <b>18.7</b>   | <b>72</b> |

Table 5: Gender resolution accuracy of decoding only and retrained wait-1 systems when translating sentences starting with ‘a woman ...’ into French.

than global object features. At  $k = 2$  however, all systems reach 100% accuracy eventually.

## 6 Conclusion

We present the first thorough investigation of the utility of visual context for the task of simultaneous machine translation. Our experiments reveal that integrating visual context lowers the latency for heuristic policies while retaining the quality of the translations. Under low-latency wait- $k$  policies, the visual cues are highly impactful with quality improvements of almost 3 BLEU points compared to unimodal baselines. From a multimodal perspective, we introduce effective ways of integrating visual features and show that explicit object region information consistently outperforms commonly used global features. Our qualitative analysis illustrates that the models are capable of resolving linguistic particularities, including gender marking and word order handling by exploiting the associated visual cues.

We hope that future research continues this line of work, especially by finding novel ways to devise adaptive policies – such as reinforcement learning models with the visual modality. We believe that our work can also benefit research in multimodal speech translation (Niehues et al., 2019) where the audio stream is accompanied by a video stream.

## Acknowledgments

We would like to thank Aizhan Imankulova for being very responsive to our questions regarding their results in Imankulova et al. (2020), and Kyunghyun Cho for answering a question related to wait-if decoding (Cho and Esipova, 2016).

Ozan Caglayan, Julia Ive, Pranava Madhyastha and Lucia Specia received funding from the MultiMT (H2020 ERC Starting Grant No. 678017) and MMVC (Newton Fund Institutional Links Grant, ID 352343575) projects.

## References

- Ashkan Alinejad, Maryam Siahbani, and Anoop Sarkar. 2018. [Prediction improves simultaneous neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3022–3027, Brussels, Belgium. Association for Computational Linguistics.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6077–6086.
- Naveen Arivazhagan, Colin Cherry, Wolfgang Macherey, Chung-Cheng Chiu, Semih Yavuz, Ruoming Pang, Wei Li, and Colin Raffel. 2019. [Monotonic infinite lookback attention for simultaneous machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1313–1323, Florence, Italy. Association for Computational Linguistics.
- Naveen Arivazhagan, Colin Cherry, Wolfgang Macherey, and George Foster. 2020. [Re-translation versus streaming for simultaneous translation](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 220–227, Online. Association for Computational Linguistics.
- Philip Arthur, Trevor Cohn, and Gholamreza Haffari. 2020. Learning coupled policies for simultaneous machine translation. *arXiv preprint arXiv:2002.04306*.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *Computing Research Repository*, arXiv:1409.0473. Version 7.
- Srinivas Bangalore, Vivek Kumar Rangarajan Sridhar, Prakash Kolan, Ladan Golipour, and Aura Jimenez. 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*, pages 437–445, Montréal, Canada. Association for Computational Linguistics.
- Loïc Barrault, Fethi Bougares, Lucia Specia, Chiraag Lala, Desmond Elliott, and Stella Frank. 2018. [Findings of the third shared task on multimodal machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 304–323, Belgium, Brussels. Association for Computational Linguistics.

- Thomas Bub, Wolfgang Wahlster, and Alex Waibel. 1997. *Verbmobil: The combination of deep and shallow processing for spontaneous speech translation*. In *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 71–74. IEEE.
- Ozan Caglayan, Walid Aransa, Adrien Bardet, Mercedes García-Martínez, Fethi Bougares, Loïc Barrault, Marc Masana, Luis Herranz, and Joost van de Weijer. 2017a. *LIUM-CVC submissions for WMT17 multimodal translation task*. In *Proceedings of the Second Conference on Machine Translation*, pages 432–439, Copenhagen, Denmark. Association for Computational Linguistics.
- Ozan Caglayan, Walid Aransa, Yaxing Wang, Marc Masana, Mercedes García-Martínez, Fethi Bougares, Loïc Barrault, and Joost van de Weijer. 2016. *Does multimodality help human and machine for translation and image captioning?* In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 627–633, Berlin, Germany. Association for Computational Linguistics.
- Ozan Caglayan, Mercedes García-Martínez, Adrien Bardet, Walid Aransa, Fethi Bougares, and Loïc Barrault. 2017b. *NMTPY: A flexible toolkit for advanced neural machine translation systems*. *Prague Bull. Math. Linguistics*, 109:15–28.
- Ozan Caglayan, Pranava Madhyastha, Lucia Specia, and Loïc Barrault. 2019. *Probing the need for visual context in multimodal machine translation*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4159–4170, Minneapolis, Minnesota. Association for Computational Linguistics.
- Iacer Calixto, Desmond Elliott, and Stella Frank. 2016. *DCU-UvA multimodal MT system report*. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 634–638, Berlin, Germany. Association for Computational Linguistics.
- Iacer Calixto and Qun Liu. 2017. *Incorporating global visual features into attention-based neural machine translation*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 992–1003, Copenhagen, Denmark. Association for Computational Linguistics.
- Kyunghyun Cho and Masha Esipova. 2016. *Can neural machine translation do simultaneous translation?* *arXiv preprint arXiv:1606.02012*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. *Learning phrase representations using RNN encoder–decoder for statistical machine translation*. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, and Stephan Vogel. 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)*, pages 493–499, New Orleans, Louisiana. Association for Computational Linguistics.
- Jean-Benoit Delbrouck and Stéphane Dupont. 2017. *Modulating and attending the source image during encoding improves multimodal translation*. *Computing Research Repository*, arXiv:1712.03449.
- J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. 2009. *Imagenet: A large-scale hierarchical image database*. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255.
- Desmond Elliott, Stella Frank, Loïc Barrault, Fethi Bougares, and Lucia Specia. 2017. *Findings of the second shared task on multimodal machine translation and multilingual image description*. In *Proceedings of the Second Conference on Machine Translation*, pages 215–233, Copenhagen, Denmark. Association for Computational Linguistics.
- Desmond Elliott, Stella Frank, Khalil Sima'an, and Lucia Specia. 2016. *Multi30K: Multilingual English-German image descriptions*. In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74, Berlin, Germany. Association for Computational Linguistics.
- Desmond Elliott and Ákos Kádár. 2017. *Imagination improves multimodal translation*. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 130–141, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Stig-Arne Grönroos, Benoit Huet, Mikko Kurimo, Jorma Laaksonen, Bernard Merialdo, Phu Pham, Mats Sjöberg, Umut Sulubacak, Jörg Tiedemann, Raphael Troncy, and Raúl Vázquez. 2018. *The MeMAD submission to the WMT18 multimodal translation task*. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 603–611, Belgium, Brussels. Association for Computational Linguistics.
- Jiatao Gu, Graham Neubig, Kyunghyun Cho, and Victor O.K. Li. 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*, pages 1053–1062, Valencia, Spain. Association for Computational Linguistics.

- K. He, X. Zhang, S. Ren, and J. Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Aizhan Imankulova, Masahiro Kaneko, Tosho Hirasawa, and Mamoru Komachi. 2020. Towards Multimodal Simultaneous Neural Machine Translation. *arXiv preprint arXiv:2004.03180*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *Int. J. Comput. Vision*, 123(1):32–73.
- Jindřich Libovický and Jindřich Helcl. 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)*, pages 196–202, Vancouver, Canada. Association for Computational Linguistics.
- Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, Hua Wu, and Haifeng Wang. 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*, pages 3025–3036, Florence, Italy. Association for Computational Linguistics.
- J. Niehues, R. Cattoni, S. Stüker, M. Negri, M. Turchi, E. Salesky, R. Sanabria, L. Barrault, L. Specia, and M. Federico. 2019. The IWSLT 2019 Evaluation Campaign”. In *Proceedings of the 16th International Workshop on Spoken Language Translation (IWSLT 2019)*.
- Jan Niehues, Ngoc-Quan Pham, Thanh-Le Ha, Matthias Sperber, and Alex Waibel. 2018. [Low-latency neural speech translation](#). In *Proc. Interspeech 2018*, pages 1293–1297.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. [On the difficulty of training recurrent neural networks](#). In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1310–1318, Atlanta, Georgia, USA. PMLR.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 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*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Koichiro Ryu, Shigeki Matsubara, and Yasuyoshi Inagaki. 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*, pages 683–690, Sydney, Australia. Association for Computational Linguistics.
- Rico Sennrich, Orhan Firat, Kyunghyun Cho, Alexandra Birch, Barry Haddow, Julian Hirschler, Marcin Junczys-Dowmunt, Samuel Läubli, Antonio Valerio Miceli Barone, Jozef Mokry, and Maria Nădejde. 2017. [Nematus: a toolkit for neural machine translation](#). In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 65–68, Valencia, Spain. Association for Computational Linguistics.
- Lucia Specia, Stella Frank, Khalil Sima’an, and Desmond Elliott. 2016. [A shared task on multimodal machine translation and crosslingual image description](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 543–553, Berlin, Germany. Association for Computational Linguistics.

- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Umut Sulubacak, Ozan Caglayan, Stig-Arne Grönroos, Aku Rouhe, Desmond Elliott, Lucia Specia, and Jörg Tiedemann. 2020. [Multimodal machine translation through visuals and speech](#). *Machine Translation*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. [From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions](#). *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Baigong Zheng, Renjie Zheng, Mingbo Ma, and Liang Huang. 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)*, pages 1349–1354, Hong Kong, China. Association for Computational Linguistics.