

3AM: An Ambiguity-Aware Multi-Modal Machine Translation Dataset

Xinyu Ma¹, Xuebo Liu^{2,*}, Derek F. Wong^{1,*}, Jun Rao², Bei Li³,
Liang Ding⁴, Lidia S. Chao¹, Dacheng Tao⁴, and Min Zhang²

¹NLP²CT Lab, Department of Computer and Information Science, University of Macau, Macau, China

²Institute of Computing and Intelligence, Harbin Institute of Technology, Shenzhen, China

³Northeastern University, Shenyang, China ⁴The University of Sydney, Sydney, Australia

{yc27434, derekfw, lidiasc}@um.edu.mo, {liuxuebo, zhangmin2021}@hit.edu.cn,

rao7jun@gmail.com, libei_neu@outlook.com, liangding.liam@gmail.com, dacheng.tao@gmail.com

Abstract

Multimodal machine translation (MMT) is a challenging task that seeks to improve translation quality by incorporating visual information. However, recent studies have indicated that the visual information provided by existing MMT datasets is insufficient, causing models to disregard it and overestimate their capabilities. This issue presents a significant obstacle to the development of MMT research. This paper presents a novel solution to this issue by introducing 3AM, an ambiguity-aware MMT dataset comprising 26,000 parallel sentence pairs in English and Chinese, each with corresponding images. Our dataset is specifically designed to include more ambiguity and a greater variety of both captions and images than other MMT datasets. We utilize a word sense disambiguation model to select ambiguous data from vision-and-language datasets, resulting in a more challenging dataset. We further benchmark several state-of-the-art MMT models on our proposed dataset. Experimental results show that MMT models trained on our dataset exhibit a greater ability to exploit visual information than those trained on other MMT datasets. Our work provides a valuable resource for researchers in the field of multimodal learning and encourages further exploration in this area. The data, code and scripts are freely available at <https://github.com/MaxyLee/3AM>.

Keywords: Multimodal datasets, Multimodal Machine translation

1. Introduction

Multimodal machine translation (MMT) (Specia et al., 2016) is a branch of Machine translation (MT) (Liu et al., 2019, 2020; Peng et al., 2023) that involves exploiting complementary information between multiple modalities and language generation (Sulubacak et al., 2020). MMT aims to enhance translation quality by incorporating visual information as input (Specia et al., 2016). Many MMT models have been proposed, which have shown superior performance compared to text-only models (Calixto et al., 2017; Zhang et al., 2020b; Li et al., 2022a; Zuo et al., 2023). However, recent studies have revealed that the visual information in existing MMT datasets contributes only marginally to translation quality. Experiments have shown that replacing images in the input with non-relevant images (Elliott, 2018) or random noise (Wu et al., 2021) has little effect on translation performance. Moreover, some studies have suggested that MMT models are less sensitive to visual information when exposed to complete sentences (Caglayan et al., 2019; Li et al., 2022a). Grönroos et al. (2018); Cho et al. (2021) have also indicated that pre-trained text-only models perform no worse than multimodal models in terms of translation quality.

A similar phenomenon has been observed in other multimodal tasks, such as image captioning (Devlin et al., 2015) and VQA (Goyal et al., 2017).

These observations suggest that natural language can provide a strong prior that can result in good superficial performance, without the underlying models truly understanding the visual content (Agrawal et al., 2016; Zhang et al., 2016). Therefore, we argue that MMT data needs to be more ambiguous and rich in visual concepts to help the models understand visual information and improve translation performance. As illustrated in Figure 1, given a sentence such as “A green gecko is seen on a palm.” the language prior is likely to lead us to interpret the word “palm” as meaning “palm of the hand” whereas the visual information suggests that it actually means “palm tree”.

The exploitation of language priors can create the false impression that MMT models are making progress toward understanding images, when in fact, they are only leveraging language to achieve good performance. This can hinder progress in pushing the state-of-the-art MMT models, underscoring the urgent need for a new and more challenging MMT dataset and to better understand the circumstances under which visual information can be effectively incorporated into MMT models. To address the limitations of existing MMT datasets, we propose 3AM, an ambiguity-aware MMT dataset that contains a larger number of ambiguous examples and a wider range of visual concepts, elevating the role of visual information understanding in MMT. Specifically, the dataset is constructed using a semi-automatic approach, which involves collecting and filtering a diverse range of data from

* Corresponding authors.











| Image | English Sentence | Senses | Possible Chinese Translations |
|---|--|--|--|
|  | A green gecko is seen on a palm . |   |  在 棕榈树 上看到一只绿色壁虎。  在 手掌 上看到一只绿色壁虎。 |
|  | A group of people on skis are being taped . |   |  一群滑雪板上的人正在被 录像 。  一群滑雪板上的人正在被 录音 。 |

Figure 1: Examples of our 3AM dataset, where ambiguous words are shown in **bold**, with **red** and **blue** indicating its incorrect and correct translation respectively. For instance, the word ‘palm’ has two different meanings: “palm tree” and “palm of the hand”. Only from the image can it be distinguished that the correct meaning is the former.

existing vision-and-language (V+L) datasets, scoring the ambiguity of the data using a word sense disambiguation (WSD) model, and then translating the English sentences into Chinese by professional translators according to the image contents. This results in a dataset of approximately 26K pieces of data. Examples from our proposed dataset are shown in Figure 1.

The proposed ambiguous dataset, 3AM, will compel MMT models to prioritize visual information. When a word is ambiguous in a sentence, the only means of resolving its meaning is by referencing the image. Thus, it is contended that the 3AM dataset will preclude MMT models from achieving high performance solely relying on language priors. The dataset is expected to facilitate MMT evaluation that more accurately reflects the models’ ability to comprehend visual information. To verify this hypothesis, a number of MMT models are evaluated based on 3AM and other existing MMT datasets. Extensive experiments are conducted on both text-only models and multimodal models. The experimental results show that the 3AM dataset presents a heightened level of complexity, necessitating the employment of models that effectively capitalize on visual patterns rather than relying solely on textual content.

Our study has three main contributions: 1) we propose an MMT dataset that is specifically designed to improve the understanding of visual information by collecting a diverse set of ambiguous sentences. Compared to existing MMT datasets, our 3AM dataset is more challenging and contains a richer set of concepts; 2) We evaluate the performance of state-of-the-art MMT models on our proposed dataset and show that models that can leverage visual information outperform text-only models. This finding supports our hypothesis that the 3AM dataset can encourage MMT models to better exploit visual information and improve translation results; 3) Our approach to constructing datasets involves collecting ambiguous data, which can also be used for other multimodal learning datasets. We hope that our work will facilitate a better understand-

ing of the role of visual information in multimodal learning and contribute to the advancement of research.

2. Related Work

MMT is a popular area of research that lies at the intersection of computer vision (CV) and natural language processing (NLP) (Specia et al., 2016). Based on the widely-used dataset Multi30K (Elliott et al., 2016), several models were proposed to leverage visual information to improve translation performance (Huang et al., 2016; Caglayan et al., 2016; Calixto and Liu, 2017; Delbrouck and Dupont, 2017; Elliott and Kádár, 2017). However, conflicting results have emerged from subsequent studies on the effectiveness of MMT models in utilizing visual input for translation (Elliott, 2018; Wu et al., 2021), raising questions about their capabilities. Although some work has demonstrated that models can use visual information to generate good translations under limited textual contexts such as some words in the source sentences are masked (Caglayan et al., 2019), the role of visual information remains unexplored when the text is complete. To address this issue, a number of datasets were proposed. Li et al. (2021) proposed Ambiguous Captions (AmbigCaps), a gender-ambiguous dataset in which gender information is concealed by back-translation from a gender-neutral language. Li et al. (2022) constructed VISA, a video-guided machine translation dataset by collecting ambiguous subtitles and corresponding video clips. Moreover, other datasets such as MSCTD (Liang et al., 2022) and MultiSubs (Wang et al., 2022) were proposed, but these do not focus on ambiguity. In this work, we aim to collect more diverse and ambiguous data using automatic methods, resulting in a more challenging MMT dataset. Compared to previous datasets, our dataset has different source data domains and language pairs, enabling greater diversity and facilitating a better understanding of the utilization of visual features in machine translation.




Figure 2: The process for constructing the 3AM dataset involves several steps. Firstly, we extract ambiguous words from existing WSD datasets and create a word sense dictionary using BabelNet. We then use this dictionary to filter sentences that contain ambiguous words and score them using a WSD model. Finally, we rank the sentences according to their scores to obtain the ambiguous data.

| Dataset | VE | COCO | SBU | CC | CC12M |
|------------------|------|------|------|-------|-------|
| # Sent. | 570K | 616K | 1M | 3.33M | 12.4M |
| # Filtered Sent. | 348K | 616K | 786K | 1.59M | 2.28M |

Table 1: Source data statistics.



3. Dataset Construction

This section elaborates on the procedures for constructing the English-Chinese MMT dataset. Our primary aim is to develop an MMT dataset encompassing a wider range of visual concepts and more ambiguous data. To construct the dataset, we use existing V+L datasets as the source data (Sec. 3.1), from which ambiguous sentences are selected. The data selection process can be divided into two steps: Data Filtering (Sec. 3.2.1) and Data Ranking (Sec. 3.2.2). First, we construct a word sense dictionary based on the existing WSD datasets and filter the sentences accordingly to obtain data containing ambiguous words. Then, we use a WSD model to score and rank the data acquired in the previous step to get the most likely ambiguous sentences. The complete process of building the dataset is illustrated in Figure 2.



3.1. Data Source

For the first part, we collect the source data from five existing V+L datasets:

1. **VE.** SNLI-VE (Xie et al., 2018) is a dataset for a classification task to determine the relationship between an image and a natural language statement. We use sentences with the label *entailment*, which correspond to the content of the image, as our data source.
2. **COCO.** COCO (Lin et al., 2014) is a large-scale human-annotated multimodal dataset that contains approximately 123,000 images and 5-way image-caption annotations.
3. **SBU.** SBU Captions (Ordonez et al., 2011) is an automatically constructed dataset that contains 1 million images based on a web-scale collection of captioned images from the internet.
4. **CC.** Conceptual Captions (CC) (Sharma et al., 2018) is an image alt-text dataset that has more than 3 million images with natural-language captions. In contrast with COCO, images and captions of CC are acquired from the web, and therefore contain a wider variety of styles.
5. **CC12M.** Conceptual 12M (CC12M) (Changpinyo et al., 2021) is a web-acquired dataset with about 12 million image-text pairs that covers a more diverse set of visual concepts than CC.



(a) Distributions of unique nouns per caption (b) Distributions of unique verbs per caption



(c) Distributions of caption lengths

(d) Distributions of ambiguity scores

Figure 3: Statistical histogram distributions on Multi30K, MSCTD, and 3AM. Compared with other datasets, 3AM contains longer captions with more unique nouns and verbs and higher ambiguity scores.

Considering that a large portion of the source data is noisy, we filter the data to ensure quality, which is divided into two steps. The first step is a rule-based filtering approach:

1. **Sentence Length Filtering.** We remove sentences with fewer than 5 or more than 30 tokens, as the former are less likely to be ambiguous, while the latter tend to be poorly phrased.
2. **URL Filtering.** As images under the same URL prefix are often of the same kind, we summarize the collection of URLs with poor image quality and filter out the corresponding data.
3. **Keyword Filtering.** We remove sentences that contain certain keywords, which are often of poor quality.
4. **Proper Noun Filtering.** Finally, sentences with more proper nouns in them are also removed.

The second step is language model filtering, where we use a language model to score the text and filter out the poor-quality captions. Specifically, we use the GPT-2 (Radford et al., 2019) model to evaluate the grammatical correctness of the captions. Given a caption T , the score can be formally defined as:

$$\text{GPTScore} = \exp(\text{GPT}(T)) \quad (1)$$

where the $\text{GPT}(\cdot)$ denotes the output of the GPT-2 model. A higher score indicates that the caption is more likely to be grammatically incorrect. After calculating the scores, we filter out the data with scores above a certain threshold based on the distribution of scores in the dataset.




Figure 4: Plot of the most common words in the captions of Multi30K and 3AM, the words in the 3AM dataset are more evenly distributed.

3.2. Data Selection

3.2.1. Data Filtering

The first step is to construct a word sense dictionary using the existing WSD datasets based on which the data containing ambiguous words can be obtained. The construction of the word sense dictionary consists of the following steps:

1. **Vocabulary Collection.** We collect ambiguous words from the following WSD datasets to construct the ambiguous vocabulary: MuCoW (Raganato et al., 2020), OneSeC (Scarlina et al., 2019), Train-o-Matic (Pasini and Navigli, 2017), SemCor (Miller et al., 1994), and OMSTI (Taghipour and Ng, 2015). The definitions of each word in these datasets along with the WordNet (Miller, 1992) sense keys are saved for word sense disambiguation. Since a word may have different senses for different

parts of speech (POS), we add words to the vocabulary with their POS as well following Li et al. (2022). For example, for the noun ‘fountain’, we add ‘fountain.n’ to the vocabulary.

- 2. Word Sense Dictionary.** In this part, we build a dictionary containing ambiguous words with their Chinese translations and corresponding definitions. For each word in the vocabulary, we use BabelNet (Navigli and Ponzetto, 2012), a large-scale multilingual encyclopedic dictionary, to get its Chinese translations and the corresponding senses. Then, we filter out words with only one sense based on the word sense dictionary, which is considered not ambiguous. As illustrated in Figure 2, given a word w , we can get the set of its senses $S = \text{Senses}(w)$ with its Chinese translation from the word sense dictionary.

After constructing the word sense dictionary, we can use it to select sentences containing ambiguous words. First, we use spaCy (Honnibal and Montani, 2017) to tokenize the sentences, convert words to the base form, and get the POS of words. Then we match these sentences with the dictionary to get data containing ambiguous words. Note that a sentence may contain more than one ambiguous word, they are considered as different data entries in the next step.

3.2.2. Word Sense Disambiguation

As the second step, we apply an automatic method to determine whether words are ambiguous in context. We use a WSD model (ConSeC) (Barba et al., 2021) to score the data obtained in the previous step. Given a sentence T and a target word w , the probability distribution P of the set of senses S of the target word in the context can be defined as:

$$P(S|T, w) = \text{WSD}(T, w, S) \quad (2)$$

where the $\text{WSD}(\cdot)$ is the output of the WSD model. Then, following Pasini and Navigli (2017), the ambiguity score of word w in sentence T can be defined as:

$$\text{AmbigScore}(T, w) = P(s_1|T, w) - P(s_2|T, w) \quad (3)$$

where $s_1 = \text{argmax}_{s \in S} P(s|T, w)$ and $s_2 = \text{argmax}_{s \in S \setminus \{s_1\}} P(s|T, w)$. We then sort the data using the ambiguity score in descending order and get a ranked list of data that the WSD model considers most likely to be ambiguous. To ensure the diversity of ambiguous words, we set an upper bound on the number of occurrences of each sense and get about 30,000 pieces of data for annotation.

3.3. Annotation

We engaged professional annotators from translation companies to translate our data. If the sentence does not match the image or is of low quality, the annotators will remove it. To ensure high quality, the annotated data are sampled in batches, and if the error rate is high, the entire batch is checked. A total of approximately 26,000 items of data are annotated, and we manually select 1000 highly ambiguous sentences for inclusion in our validation and test sets.

3.4. Dataset Statistics

3.4.1. Basic Statistics

We conduct a comprehensive analysis of the 3AM dataset and other MMT datasets. The statistical properties of Multi30K, MSCTD, and our proposed 3AM are shown in Figure 3. We calculate the number of unique words (noun and verb) in each sentence across different datasets. As shown in Figure 3a and 3b, 3AM contains more unique words than other datasets, showing that 3AM is of higher lexical richness. In addition, Figure 3c illustrates that the captions of 3AM are longer than others. We also use a WSD model to score the ambiguity of the data. It can be seen from Figure 3d that the 3AM dataset has a higher ambiguity. Besides, we count the most common words and, as shown in Figure 4, the words in our dataset are more evenly distributed. Here we do not compare with MSCTD as its lexical distribution differs considerably from the other datasets.

3.4.2. Diversity Analysis

Additionally, we conduct a quantitative analysis of diversity, which is presented in Table 2. To evaluate text diversity, we calculate the average sentence length and the proportion of the number of distinct n-grams to the total number of n-grams (Dist-n) following Li et al. (2016). The results show that 3AM contains longer captions with a greater number of unique n-grams, demonstrating that it is more lexically rich than other datasets. To assess the diversity of images, we use the Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al., 2018a) and Inception Score (IS) (Salimans et al., 2016). For LPIPS, we randomly sample 1000 images from each dataset and calculate the perceptual distance between each pair of images following Lee et al. (2018). As shown in Table 2, the LPIPS and IS of 3AM is significantly higher, indicating greater image diversity. Moreover, to evaluate the diversity of objects in the images, we propose to use the Entropy of Objects (Ent-Obj) following Zhang et al.

| Dataset | Text | | | | | Image | | |
|----------|-------------|--------|--------|--------|--------|-------------------|--------------|---------|
| | Avg. length | Dist-1 | Dist-2 | Dist-3 | Dist-4 | LPIPS | IS | Ent-Obj |
| Multi30K | 13.06 | 0.25 | 2.29 | 5.26 | 7.31 | 0.80584 ± 0.00010 | 23.25 ± 2.58 | 3.15 |
| MSCTD | 8.40 | 0.17 | 1.38 | 3.16 | 4.07 | 0.74149 ± 0.00011 | 7.85 ± 0.20 | 3.21 |
| 3AM | 13.48 | 0.77 | 5.23 | 8.85 | 9.67 | 0.82975 ± 0.00011 | 29.94 ± 3.75 | 4.35 |

Table 2: Detailed statistics of Multi30K, MSCTD, and 3AM. We demonstrate the average sentence length, the LPIPS (higher means more different), the average number of unique n-grams, and the average number of unique words. Our 3AM dataset is lexically richer than other datasets in general.

(2018b):

$$\text{Ent-Obj} = -\frac{1}{\sum_o F(o)} \sum_{o \in O} F(o) \log\left(\frac{F(o)}{\sum_o F(o)}\right) \quad (4)$$

where O is the set of all objects, and $F(o)$ denotes the frequency of object o . Specifically, we use YOLOv6 (Li et al., 2022b) to detect objects in the images. The results show that, unlike other datasets where the majority of objects are ‘persons’, etc., the objects in 3AM are more evenly distributed.

4. Experiments

4.1. Datasets

For the dataset, in addition to experiments on 3AM, we also conduct experiments on other MMT datasets for comparison.

4.1.1. Multi30K

Multi30K (Elliott et al., 2016) is a commonly used MMT dataset based on Flickr30K (Young et al., 2014) with the manual translation of English captions into German and French with approximately 30,000 pieces of data, including three testsets: test2016, test2017, and mscoco. Here we use test2017 as the test set. For comparison with 3AM, we create an English-Chinese version of Multi30K by translating the English sentences into Chinese by a strong English-Chinese MT model.

4.1.2. MSCTD

MSCTD (Liang et al., 2022) is a multimodal sentiment chat translation dataset that contains 142,871 English-Chinese parallel sentences and 30,370 English-German parallel sentences. We use English-Chinese translation data from MSCTD to compare with our dataset.

4.2. Baseline Models

In order to provide a convincing benchmark for 3AM, we conduct experiments on multiple machine translation models as the baseline models, which can

be divided into two categories: Text-only models and Multimodal models.

4.2.1. Text-only models

1. **Trans.** We use the Transformer (Trans) (Vaswani et al., 2017) model as one of the baseline models. We conduct experiments on the Transformer-Tiny configuration following previous work (Wu et al., 2021; Li et al., 2022a) for a fair comparison with the Selective Attention model.
2. **Bart.** Bart (Lewis et al., 2020) is a Transformer-based model pre-trained by corrupting text with an arbitrary noise and allows the model to reconstruct it. As Chinese is not included in the vocabulary of Bart, we extend it by incorporating Chinese tokens and then extending the token embeddings of the model to the corresponding dimension.
3. **T5.** T5 (Raffel et al., 2020) is a Transformer model pre-trained on a mixture of supervised and unsupervised tasks which are converted into the text-to-text format. We extend T5 using the same approach as Bart.

4.2.2. Multimodal models

1. **SelAttn.** Selective Attention (SelAttn) (Li et al., 2022a) is an MMT model that uses vision features extracted by Transformer-based models such as Vision Transformer (ViT) (Dosovitskiy et al., 2021). As previously mentioned, we use a small model size following Li et al. (2022a).
2. **VL-Bart.** VL-Bart (Cho et al., 2021) is a multimodal pre-trained model that extends the Bart model encoder to a multimodal encoder by incorporating image features as additional input. We employ the same extending approach as we did for Bart and T5.
3. **VL-T5.** Similarly, VL-T5 (Cho et al., 2021) is a multimodal pre-trained model extended from T5 using the same method as that of VL-Bart.

| Multi30K (train) | | | | | | | | | | | | |
|------------------|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Method | Multi30K (test) | | | | MSCTD (test) | | | | 3AM (test) | | | |
| | B ↑ | BS ↑ | M ↑ | T ↓ | B ↑ | BS ↑ | M ↑ | T ↓ | B ↑ | BS ↑ | M ↑ | T ↓ |
| Trans | 42.86 | 74.32 | 65.44 | 47.86 | 2.87 | 34.99 | 15.75 | 108.20 | 10.86 | 49.10 | 29.40 | 88.85 |
| SelAttn | 42.00 | 74.17 | 64.63 | 49.82 | 2.86 | 36.00 | 16.61 | 107.84 | 11.67 | 50.05 | 30.86 | 87.20 |
| Bart | 56.93 | 83.24 | 79.61 | 32.47 | 7.40 | 46.71 | 29.35 | 101.93 | 22.29 | 59.19 | 45.43 | 73.87 |
| VL-Bart | 56.70 | 82.93 | 77.89 | 32.00 | 8.12 | 46.29 | 27.22 | 86.40 | 23.20 | 60.20 | 45.75 | 70.95 |
| T5 | 60.59 | 85.69 | 82.85 | 27.61 | 10.24 | 52.53 | 38.78 | 85.30 | 25.03 | 62.99 | 50.72 | 67.08 |
| VL-T5 | 59.61 | 85.25 | 82.12 | 27.95 | 11.10 | 52.96 | 38.71 | 77.71 | 25.34 | 63.25 | 50.89 | 66.35 |

| MSCTD (train) | | | | | | | | | | | | |
|---------------|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Method | Multi30K (test) | | | | MSCTD (test) | | | | 3AM (test) | | | |
| | B ↑ | BS ↑ | M ↑ | T ↓ | B ↑ | BS ↑ | M ↑ | T ↓ | B ↑ | BS ↑ | M ↑ | T ↓ |
| Trans | 9.89 | 50.43 | 30.75 | 80.68 | 22.97 | 62.93 | 46.43 | 65.40 | 4.51 | 40.69 | 20.10 | 88.37 |
| SelAttn | 6.91 | 46.75 | 25.04 | 85.31 | 20.87 | 62.08 | 44.27 | 65.58 | 5.30 | 41.87 | 21.05 | 108.70 |
| Bart | 22.77 | 65.66 | 51.50 | 59.95 | 32.68 | 69.82 | 56.68 | 52.60 | 14.93 | 56.34 | 38.72 | 74.58 |
| VL-Bart | 18.10 | 60.34 | 44.81 | 65.29 | 30.81 | 68.96 | 55.63 | 54.03 | 13.61 | 54.24 | 36.46 | 77.53 |
| T5 | 29.17 | 72.04 | 59.82 | 51.32 | 29.39 | 70.43 | 54.22 | 54.46 | 18.49 | 59.68 | 44.13 | 70.26 |
| VL-T5 | 28.43 | 71.09 | 58.85 | 52.82 | 29.49 | 70.63 | 54.48 | 54.52 | 17.87 | 59.27 | 43.44 | 70.55 |

| 3AM (train) | | | | | | | | | | | | |
|-------------|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Method | Multi30K (test) | | | | MSCTD (test) | | | | 3AM (test) | | | |
| | B ↑ | BS ↑ | M ↑ | T ↓ | B ↑ | BS ↑ | M ↑ | T ↓ | B ↑ | BS ↑ | M ↑ | T ↓ |
| Trans | 25.95 | 64.51 | 49.88 | 63.92 | 3.53 | 39.23 | 19.02 | 102.93 | 11.33 | 49.51 | 31.34 | 89.68 |
| SelAttn | 27.81 | 67.06 | 52.13 | 59.77 | 4.25 | 40.34 | 19.84 | 100.19 | 13.33 | 51.54 | 33.47 | 87.05 |
| Bart | 48.13 | 80.16 | 76.07 | 39.19 | 13.45 | 54.61 | 38.30 | 84.94 | 31.47 | 65.87 | 55.62 | 63.65 |
| VL-Bart | 50.13 | 80.74 | 76.38 | 36.87 | 16.13 | 56.45 | 39.15 | 74.17 | 33.27 | 66.56 | 55.84 | 61.28 |
| T5 | 50.16 | 81.84 | 79.18 | 35.92 | 15.56 | 59.18 | 48.04 | 77.79 | 33.09 | 68.15 | 57.26 | 60.09 |
| VL-T5 | 52.04 | 82.60 | 79.76 | 34.37 | 17.12 | 59.94 | 48.54 | 73.01 | 34.24 | 68.39 | 59.12 | 58.88 |

Table 3: Performance of MMT models on 3AM and other MMT datasets in terms of BLEU (B), BERT-Score (BS), METEOR (M), and TER (T).

4.3. Implementation Details

4.3.1. Training

Our models are implemented based on PyTorch (Paszke et al., 2017), Fairseq (Ott et al., 2019), and Huggingface Transformers (Wolf et al., 2020). All experiments are run on Nvidia A40 GPUs. The implementation details of different models are as follows:

1. Models trained from scratch. The Transformer and Selective Attention models are trained from scratch, consisting of 4 encoder and decoder layers. We set the hidden size as 128, the filter size of FFN as 256, and 4 heads in the multi-head self-attention. We use the Adam Optimizer (Kingma and Ba, 2015) with $\beta_1 = .9$, $\beta_2 = .98$, and $\epsilon = 10^{-8}$. We train the models with a batch size of 4096 tokens and a learning rate of 6×10^{-4} . We adopt the early stopping training strategy to prevent overfitting following Li et al. (2022a).

2. Pre-trained models. We use the pre-trained T5-base¹ and Bart-base² models from Huggingface Transformers for fine-tuning following Cho et al. (2021). The models are fine-tuned with a batch size of 60 and a learning rate of 5×10^{-4} . We use AdamW Optimizer (Loshchilov and Hutter, 2019) with a weight decay of 0.01.

4.3.2. Evaluation

The model with the lowest validation perplexity is used for model testing. The beam size for inference is set to 5. We use BLEU (Papineni et al., 2002), BERT-Score (Zhang et al., 2020a), METEOR (Banerjee and Lavie, 2005), and TER (Snober et al., 2006) as the automatic evaluation metrics. We use SacreBLEU (Post, 2018) to calculate the BLEU score following Cho et al. (2021) and

¹<https://huggingface.co/t5-base>

²<https://huggingface.co/facebook/bart-base>

| Dataset | C | I | Δ -Awareness |
|----------|-------|------------------|---------------------|
| Multi30K | 74.16 | 74.11 \pm 0.04 | 0.05 \pm 0.04 |
| MSCTD | 62.08 | 62.08 \pm 0.00 | 0.00 \pm 0.00 |
| 3AM | 51.54 | 50.17 \pm 0.09 | 1.36 \pm 0.09 |

Table 4: BERT-Scores under Congruent (**C**) and Incongruent (**I**) settings, and the image awareness results. The Incongruent and Δ -Awareness scores are calculated by taking the mean and standard deviation of the scores obtained from five random permutations of the visual data.

the Huggingface Transformers (Wolf et al., 2020) for the METEOR score and BERT-Score. Specifically, we use the BERT-Large-Chinese³ model for calculating the BERT-Score.

4.4. Results and Discussion

The experiment results are shown in Table 3. Each model is trained on three datasets and the performance is evaluated on each test set. Although MMT models trained on Multi30K perform close to or even worse than their text-only counterparts on both Multi30K and MSCTD test sets, they perform slightly better on the 3AM test set. This result confirms our view that although the current MMT models can exploit visual information, this ability is hardly reflected in the existing datasets as visual information plays a minor role. In addition, the poor performance of the MMT models trained on the MSCTD dataset compared to the text-only models illustrates the inability of MSCTD to allow the model to learn how to utilize visual information. Besides, models trained on the MSCTD dataset perform poorly on the other test sets as it is a dialogue dataset and differs significantly from the other data. Moreover, MMT models trained on the 3AM dataset outperform the text-only models by a large margin, demonstrating that visual information plays a vital role in our 3AM dataset. These findings confirm our hypothesis that our proposed dataset will force the MMT models to leverage visual information to disambiguate and thus generate higher-quality translations.

4.5. Analysis

4.5.1. Visual Awareness

To evaluate the effect that visual information has in translation, we use an adversarial evaluation method following Elliott (2018) to calculate image awareness. Let x denote the source sentence, y denote the target sentence, v denote the congruent image, and \bar{v} denote the incongruent image. The

³<https://huggingface.co/yechen/bert-large-chinese>



Source: A group of people on skis are being **taped**.

Target: 一群滑雪板上的人正在被**录像**。(record video)

T5: 一群踩着滑雪板的人正在被**录音**。(record audio)

VL-T5: 一群滑雪板上的人正在被**录制视频**。(record video)

Figure 5: A case study of the 3AM dataset. The ambiguous word in the source sentence is in **bold**. The **red** and **blue** represent incorrectly and correctly translated words respectively.

overall image awareness of a model \mathcal{M} on dataset \mathcal{D} can be defined as:

$$\Delta\text{-awareness} = \frac{1}{|\mathcal{D}|} \sum_i a_{\mathcal{M}}(x_i, y_i, v_i, \bar{v}_i) \quad (5)$$

where the $a_{\mathcal{M}}(\cdot)$ is the image awareness of model \mathcal{M} on a single instance:

$$a_{\mathcal{M}}(x_i, y_i, v_i, \bar{v}_i) = \varepsilon(x_i, y_i, v_i) - \varepsilon(x_i, y_i, \bar{v}_i) \quad (6)$$

where the ε is an evaluable performance measure, and here we use the BERT-Score.

To determine whether a model passes the evaluation, we conduct a Wilcoxon signed-rank test using the pairs of BERT-scores calculated in the process of computing the image awareness and combine 5 separate p values from each test using Fisher’s method following Elliott (2018). Table 4 shows the evaluation results of the BERT-Scores and the image awareness scores. We find that while the model did not pass the evaluation on both Multi30K ($\chi^2 = 15.48$, $p = 0.1156$) and MSCTD ($\chi^2 = 9.54$, $p = 0.4819$), the visual information improves the performance on the 3AM dataset ($\chi^2 = 249.46$, $p < 0.0001$) by a large margin. This demonstrates that MMT models trained on 3AM utilize visual information to generate better translation.

4.5.2. Case Study

As shown in Figure 5, the word ‘tape’ in the source sentence has two senses in Chinese: ‘record audio’ and ‘record video’. Although it is difficult to distinguish from the text, from the image we can easily see that the real meaning is the latter. In this example, the T5 model incorrectly translates the word, whereas VL-T5 produces the correct translation as it can utilize visual information to disambiguate. This observation further confirms our hypothesis that MMT models trained on the 3AM dataset can effectively exploit visual information.

5. Conclusion

In this paper, we present 3AM, a challenging MMT dataset that contains more ambiguity and diverse visual concepts, consisting of approximately 26K

data points Several MMT models are benchmarked on the proposed dataset and compared with existing MMT datasets. It is found that models with visual input outperform those without visual input, confirming the hypothesis that the ambiguous dataset forces the models to focus on visual information. This enables a more realistic evaluation of the models' performance and demonstrates the effectiveness of using visual information rather than relying solely on language priors. The release of the 3AM dataset is expected to facilitate the advancement of research on multimodal learning.

6. Acknowledgements

This work was supported in part by the Science and Technology Development Fund, Macau SAR (Grant Nos. FDCT/060/2022/AFJ, FDCT/0070/2022/AMJ), National Natural Science Foundation of China (Grant No. 62261160648), Ministry of Science and Technology of China (Grant No. 2022YFE0204900), the Multi-year Research Grant from the University of Macau (Grant No. MYRG-GRG2023-00006-FST-UMDF), Shenzhen College Stability Support Plan (Grant Nos. GXWD20220811173340003, GXWD20220817123150002), and Shenzhen Science and Technology Program (Grant Nos. RCBS20221008093121053, ZDSYS20230626091203008). Xuebo Liu was sponsored by CCF-Tencent Rhino-Bird Open Research Fund. This work was performed in part at SICC which is supported by SKL-IOTSC, and HPCC supported by ICTO of the University of Macau.

7. Ethical Considerations

Our 3AM dataset was developed on the basis of existing datasets and therefore adheres to the corresponding copyrights. We hired a professional translation company for the data annotation, a contract was signed and the staff were duly remunerated.

8. Bibliographical References

- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. 2016. [Analyzing the behavior of visual question answering models](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1955–1960, Austin, Texas. Association for Computational Linguistics.
- Satanjeev Banerjee and Alon Lavie. 2005. [ME-THEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Edoardo Barba, Luigi Procopio, and Roberto Navigli. 2021. [ConSeC: Word sense disambiguation as continuous sense comprehension](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1492–1503, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ozan Caglayan, Loïc Barrault, and Fethi Bougares. 2016. [Multimodal attention for neural machine translation](#). *ArXiv preprint*, abs/1609.03976.
- 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 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.
- Iacer Calixto, Qun Liu, and Nick Campbell. 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)*, pages 1913–1924, Vancouver, Canada. Association for Computational Linguistics.
- Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. [Unifying vision-and-language tasks via text generation](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 1931–1942. PMLR.
- Jean-Benoit Delbrouck and Stéphane Dupont. 2017. [An empirical study on the effectiveness of images in multimodal neural machine translation](#).

- In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 910–919, Copenhagen, Denmark. Association for Computational Linguistics.
- Jacob Devlin, Saurabh Gupta, Ross Girshick, Margaret Mitchell, and C Lawrence Zitnick. 2015. [Exploring nearest neighbor approaches for image captioning](#). *ArXiv preprint*, abs/1505.04467.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Desmond Elliott. 2018. [Adversarial evaluation of multimodal machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2974–2978, Brussels, Belgium. 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.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. [Making the V in VQA matter: Elevating the role of image understanding in visual question answering](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6325–6334. IEEE Computer Society.
- 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.
- Matthew Honnibal and Ines Montani. 2017. [spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing](#). .
- Po-Yao Huang, Frederick Liu, Sz-Rung Shiang, Jean Oh, and Chris Dyer. 2016. [Attention-based multimodal neural machine translation](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 639–645, Berlin, Germany. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. 2018. [Diverse image-to-image translation via disentangled representations](#). In *Proceedings of the European conference on computer vision (ECCV)*, pages 35–51.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Bei Li, Chuanhao Lv, Zefan Zhou, Tao Zhou, Tong Xiao, Anxiang Ma, and JingBo Zhu. 2022a. [On vision features in multimodal machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6327–6337, Dublin, Ireland. Association for Computational Linguistics.
- Chuyi Li, Lulu Li, Hongliang Jiang, Kaiheng Weng, Yifei Geng, Liang Li, Zaidan Ke, Qingyuan Li, Meng Cheng, Weiqiang Nie, et al. 2022b. [Yolov6: A single-stage object detection framework for industrial applications](#). *ArXiv preprint*, abs/2209.02976.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Xuebo Liu, Houtim Lai, Derek F. Wong, and Lidia S. Chao. 2020. [Norm-based curriculum learning for neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 427–436, Online. Association for Computational Linguistics.

- Xuebo Liu, Derek F. Wong, Yang Liu, Lidia S. Chao, Tong Xiao, and Jingbo Zhu. 2019. [Shared-private bilingual word embeddings for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3613–3622, Florence, Italy. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Roberto Navigli and Simone Paolo Ponzetto. 2012. [Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network](#). *Artificial intelligence*, 193:217–250.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. [Automatic differentiation in pytorch](#). In *NIPS 2017 Workshop on Autodiff*.
- Keqin Peng, Liang Ding, Qihuang Zhong, Li Shen, Xuebo Liu, Min Zhang, Yuanxin Ouyang, and Dacheng Tao. 2023. [Towards making the most of ChatGPT for machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5622–5633, Singapore. Association for Computational Linguistics.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. [Improved techniques for training gans](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2226–2234.
- Matthew Snover, Bonnie Dorr, Rich Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA. Association for Machine Translation in the Americas.
- 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.
- 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*, 34:97–147.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in*

Natural Language Processing: System Demonstrations, pages 38–45, Online. Association for Computational Linguistics.

Zhiyong Wu, Lingpeng Kong, Wei Bi, Xiang Li, and Ben Kao. 2021. [Good for misconceived reasons: An empirical revisiting on the need for visual context in multimodal machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6153–6166, Online. Association for Computational Linguistics.

Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2016. [Yin and yang: Balancing and answering binary visual questions](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 5014–5022. IEEE Computer Society.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018a. [The unreasonable effectiveness of deep features as a perceptual metric](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 586–595. IEEE Computer Society.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020a. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. 2018b. [Generating informative and diverse conversational responses via adversarial information maximization](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 1815–1825.

Zhuosheng Zhang, Kehai Chen, Rui Wang, Masao Utiyama, Eiichiro Sumita, Zuchao Li, and Hai Zhao. 2020b. [Neural machine translation with universal visual representation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Yuxin Zuo, Bei Li, Chuanhao Lv, Tong Zheng, Tong Xiao, and JingBo Zhu. 2023. [Incorporating probing signals into multimodal machine translation via visual question-answering pairs](#). In *Findings of the Association for Computational Linguistics:*

EMNLP 2023, pages 14689–14701, Singapore. Association for Computational Linguistics.

9. Language Resource References

Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. 2021. [Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3558–3568.

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.

Jiaoda Li, Duygu Ataman, and Rico Sennrich. 2021. [Vision matters when it should: Sanity checking multimodal machine translation models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8556–8562, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Yihang Li, Shuichiro Shimizu, Weiqi Gu, Chenhui Chu, and Sadao Kurohashi. 2022. [VISA: An ambiguous subtitles dataset for visual scene-aware machine translation](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6735–6743, Marseille, France. European Language Resources Association.

Yunlong Liang, Fandong Meng, Jinan Xu, Yufeng Chen, and Jie Zhou. 2022. [MSCTD: A multimodal sentiment chat translation dataset](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2601–2613, Dublin, Ireland. Association for Computational Linguistics.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. [Microsoft coco: Common objects in context](#). In *European conference on computer vision*, pages 740–755. Springer.

George A. Miller. 1992. [WordNet: A lexical database for English](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*.

George A. Miller, Martin Chodorow, Shari Landes, Claudia Leacock, and Robert G. Thomas. 1994.

- Using a semantic concordance for sense identification. In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Vicente Ordonez, Girish Kulkarni, and Tamara L. Berg. 2011. [Im2text: Describing images using 1 million captioned photographs](#). In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*, pages 1143–1151.
- Tommaso Pasini and Roberto Navigli. 2017. [Train-O-Matic: Large-scale supervised word sense disambiguation in multiple languages without manual training data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 78–88, Copenhagen, Denmark. Association for Computational Linguistics.
- Alessandro Raganato, Yves Scherrer, and Jörg Tiedemann. 2020. [An evaluation benchmark for testing the word sense disambiguation capabilities of machine translation systems](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 3668–3675, Marseille, France. European Language Resources Association.
- Bianca Scarlino, Tommaso Pasini, and Roberto Navigli. 2019. [Just “OneSeC” for producing multilingual sense-annotated data](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 699–709, Florence, Italy. Association for Computational Linguistics.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. [Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, Melbourne, Australia. Association for Computational Linguistics.
- Kaveh Taghipour and Hwee Tou Ng. 2015. [One million sense-tagged instances for word sense disambiguation and induction](#). In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 338–344, Beijing, China. Association for Computational Linguistics.
- Josiah Wang, Josiel Figueiredo, and Lucia Specia. 2022. [MultiSubs: A large-scale multimodal and multilingual dataset](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6776–6785, Marseille, France. European Language Resources Association.
- Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2018. [Visual entailment task for visually-grounded language learning](#). *ArXiv preprint*, abs/1811.10582.
- 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.