



PDF Download
3748311.pdf
18 January 2026
Total Citations: 0
Total Downloads: 431

Latest updates: <https://dl.acm.org/doi/10.1145/3748311>

RESEARCH-ARTICLE

Impact of Visual Context on Noisy Multimodal NMT: An Empirical Study for English to Indian Languages

BABAN GAIN, Indian Institute of Technology Patna, Patna, BR, India

DIBYANAYAN BANDYOPADHYAY, Indian Institute of Technology Patna, Patna, BR, India

SAMRAT MUKHERJEE, Indian Institute of Technology Patna, Patna, BR, India

CHANDRANATH ADAK, Indian Institute of Technology Patna, Patna, BR, India

ASIF EKBAL, Indian Institute of Technology Patna, Patna, BR, India

Open Access Support provided by:

Indian Institute of Technology Patna

Published: 21 August 2025
Online AM: 10 July 2025
Accepted: 31 May 2025
Revised: 21 May 2025
Received: 19 September 2023

[Citation in BibTeX format](#)

Impact of Visual Context on Noisy Multimodal NMT: An Empirical Study for English to Indian Languages

BABAN GAIN, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India

DIBYANAYAN BANDYOPADHYAY, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India

SAMRAT MUKHERJEE, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India

CHANDRANATH ADAK, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India

ASIF EKBAL, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India and School of AI and Data Science IIT Jodhpur, Indian Institute of Technology Jodhpur, Jodhpur, India

Neural Machine Translation (NMT) has made remarkable progress using large-scale textual data, but the potential of incorporating multimodal inputs, especially visual information, remains underexplored in high-resource settings. While prior research has focused on using multimodal data in low-resource scenarios, this study examines how image features impact translation when added to a large-scale, pre-trained unimodal NMT system. Surprisingly, the study finds that images might be redundant in this context. Additionally, the research introduces synthetic noise to assess whether images help the model handle textual noise. Multimodal models slightly outperform text-only models in noisy settings, even when random images are used. The study's experiments translate from English to Hindi, Bengali, and Malayalam, significantly outperforming state-of-the-art benchmarks. Interestingly, the effect of visual context varies with the level of source text noise: no visual context works best for non-noisy translations, cropped image features are optimal for low noise, and full image features perform better in high-noise scenarios. This sheds light on the role of visual context, especially in noisy settings, and opens up a new research direction for Noisy Neural Machine Translation in multimodal setups. The research emphasizes the importance of combining visual and textual information to improve translation across various environments. Our code is publicly available at <https://github.com/babangain/indicMMT>.

CCS Concepts: • **Computing methodologies** → **Machine translation**;

Baban Gain and Dibyanayan Bandyopadhyay gratefully acknowledge the Prime Minister's Research Fellowship (PMRF) program for providing financial support and enabling this research.

Authors' Contact Information: Baban Gain, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India; e-mail: gainbaban@gmail.com; Dibyanayan Bandyopadhyay, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India; e-mail: dibyanayan@gmail.com; Samrat Mukherjee, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India; e-mail: samratpiv123@gmail.com; Chandranath Adak, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India; e-mail: adak32@gmail.com; Asif Ekbal, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna, India and School of AI and Data Science IIT Jodhpur, Indian Institute of Technology Jodhpur, Jodhpur, India; e-mail: asif.ekbal@gmail.com.



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

© 2025 Copyright held by the owner/author(s).

ACM 2375-4699/2025/08-ART79

<https://doi.org/10.1145/3748311>

Additional Key Words and Phrases: Context-aware translation, multimodal translation

ACM Reference Format:

Baban Gain, Dibyanayan Bandyopadhyay, Samrat Mukherjee, Chandranath Adak, and Asif Ekbal. 2025. Impact of Visual Context on Noisy Multimodal NMT: An Empirical Study for English to Indian Languages. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.* 24, 8, Article 79 (August 2025), 27 pages. <https://doi.org/10.1145/3748311>

1 Introduction

Unimodal Neural Machine Translation (UNMT) systems generate reliable outputs across many domains and languages. However, these models often fail to translate ambiguous words that require additional contextual information for disambiguation. This issue is further amplified when the dataset is noisy, containing spelling or grammatical errors. Additionally, UNMT models struggle to produce context-dependent translations. **Multimodal Neural Machine Translation (MMT)** is particularly important in the era of social media.

In MMT, an image is used as contextual information to support the translation process. For example, in the sentence “*The color of the court is red,*” the word *court* can be interpreted in multiple ways. One interpretation is a playground for indoor games like tennis or badminton, while another could refer to a judicial court. Although several techniques have been proposed for MMT, most have been applied to high-resource language pairs, such as English–German or English–French. These languages belong to the same linguistic family, use the Latin script, and often share common vocabulary. Translating between distant language pairs, e.g., English and Hindi, is more challenging, as they share fewer commonalities. Furthermore, Hindi sentences are gender-dependent, whereas English is not [14]. For example, the translation of “*I am going to the market.*” differs depending on whether the subject is male or female. Although there have been a few efforts to incorporate multimodal information for translating English to Hindi, most do not leverage large-scale unimodal pre-training. Existing methods either omit large-scale unimodal data entirely or use it only for limited purposes such as word embedding training. While multimodal translation is inherently different from unimodal translation, training encoders and decoders with large-scale unimodal data may help achieve better translation performance. Although object tags can aid translation quality [26], relying solely on object tags without incorporating additional visual information from the image may be suboptimal for the following reasons:

- Only object tags (such as *person, ball, tree,* etc.) are linked to the source text; other details about the objects and images that could assist in translation are not utilized.
- If there is no additional information about the objects, their locations, or their interactions, the model may use object tags misleadingly. For instance, the model may fail to determine whether the subject is *male* or *female* if another *male* person is present in the image and the subject is actually *female*, as both would appear in the list of detected objects regardless of their actions.
- For each detected object, only one tag with the highest confidence is selected. Misclassified tags can mislead the (**Neural Machine Translation**) NMT model because tag selection is discrete, and the object detection model is prone to errors.
- Due to inadequate training of the object detection model, many objects may remain unclassified.

The aforementioned factors prevented object tag-based models from consistently improving the results. Therefore, it is preferable to use and incorporate visual features that have been extracted from pre-trained image models, e.g., **Vision Transformer (ViT)**, VGG, ResNet [27], and so on. When source texts are hidden, some probing tasks demonstrate how MMT models can retrieve

missing data, such as color, person, and so on, from visual features [49]. Here, some part of the textual input is masked, and the model is trained to translate the sentence with the help of the image, forcing it to focus more on the visual features in order to retrieve information about masked input tokens. However, it may not be a realistic strategy in scenarios involving incomplete information. Any mask token will not be present in real-world datasets, where some of the information is missing from the source text, and it may be challenging to infer that the information is insufficient. In summary, efforts have been made to make use of visual features in the MT system directly. Most of them only use multimodal data, which is scarcely available. Additionally, other techniques that rely on strong baselines only use image object tags when fine-tuning. In this study, we investigate the circumstances in which visual features have a positive impact on MMT systems. We use large pre-trained models as a baseline so that the models are robust enough to distinguish words from intra-sentence context when available. We observe that images have a negligible contribution in non-noisy settings. We further observe that the models are able to utilize cropped image features in low-noisy settings, and full image features are more effective in high-noisy settings. Our study indicates that visual features' impact on the model performance increases with noise in the source text.

The major **contributions** of this article are briefly outlined below.

(i) Noisy MMT: We develop our systems for noisy MMT. To the best of our knowledge, this is the earliest attempt of its kind to translate noisy utterances in multimodal settings. We observe an increased positive impact of visual context in the presence of noise in the dataset.

(ii) Multilingual MMT: We perform multilingual multimodal fine-tuning and utilize a multilingual pre-trained model. This is the first attempt at a multilingual setting that includes Indic languages.

(iii) Probing the Need for Images: By substituting relevant images corresponding to the text with random ones, we investigate the effectiveness of using semantically aligned visual information. This is the first attempt of its kind to conduct such probing in Indian languages.

(iv) Strong UNMT Baseline: We use a strong off-the-shelf UNMT model for fine-tuning. This is essential to determine whether the improved performance of multimodal methods stems from the extracted information in images or from the limitations of the text-only model in handling ambiguity.

(v) Enhanced Visual Features in Indian Settings: We extract rich visual features from ViT as well as CLIP instead of using CNN-based features adopted in contemporary Indic MMT methods.

(vi) State-of-the-Art (SOTA) Results: This work primarily aims to understand the impact of visual features on multimodal English-to-Indic language translation in the presence of noise. While our goal is not solely to achieve SOTA performance, our model obtains SOTA results for all three target Indic languages (Hindi, Bengali, and Malayalam), with English as the source language.

The rest of the article is organized as follows. In Section 2, we discuss related MMT methods along with models for Indic languages and noisy text translation. Thereafter, Section 3 explains the employed dataset and the noise addition procedure. Subsequently, Section 4 mentions our methodology, and Section 5 and Section 6 discuss the results with observations and ablations. Finally, Section 7 concludes this article.

2 Related Works

MMT has seen growing interest due to its ability to incorporate complementary visual context to improve translation quality. However, the progress in this area is constrained by the limited availability of large-scale source-target-image triplet datasets. Despite this, a wide range of approaches have been proposed to address the challenges in MMT through innovations in model architecture,

integration of visual features, and adaptations for low-resource and linguistically diverse settings. This section surveys major developments in general MMT, its application to Indian languages, context-aware translation, and robustness to noisy input.

2.1 Multimodal Translation

One of the earliest attempts for MMT was to integrate visual features with a shared attention mechanism [6]. However, it did not result in performance improvement. Other early attempts include the concatenation of text features with global and regional visual features [29] to generate Multimodal Translation. Here, regional features are treated as tokens of the input sentence. Subsequently, a multi-task-based training framework [18] was employed with a shared encoder where the tasks included standard NMT tasks and visual feature reconstruction. The following attempts to tackle challenges included consideration of visual features as pseudo-text features [8]. First, they projected visual feature vectors into source text embedding space and used them as the first and/or last word of the source sentence. Furthermore, they use them to initialize the hidden state of the encoder instead of using a zero vector. Finally, they use the projected image representations and encoder output to initialize hidden representations of a decoder. They use bidirectional RNN as the encoder. Since visual features are needed in very specific cases, the model tends to learn to ignore visual features, resulting in ignoring them, even when necessary. Therefore, a visual attention grounding mechanism [86] was adopted for MMT that links the visual semantics with the corresponding textual semantics.

In order to enable the MMT systems to adequately handle multimodal information, a translate and refine approach [31] was adopted for MMT, where the first stage decoder generates translation on textual information and second stage decoder refines it with the help of the visual features. To improve the interaction between textual and visual features, a latent variable was added to model [10]. This latent variable can be seen as a stochastic embedding that is used by the decoder for translation, as well as it can be used to predict image features. Attention [79] based multimodal interactions are widely used in MMT. This includes cross-attention between regional image features and decoder self-attention output [85].

Multimodal information generally helped in obtaining better results. However, it was observed that many of the improvements were not related to the presence of visual features in MMT. This indicated a need to probe how and why MMT models generate better translation compared to UNMT models. In a probing task [49], the words that included color information (red, green, etc.) were masked. In addition, words that indicate characters and nouns were also replaced with mask tokens from the source text to simulate a situation where information is insufficient in the source text. They proposed a Selective Attention-based transformer that is able to recover the masked information from the image at the generated translation better than other existing models. Similarly, it was observed that models can integrate the visual modality if they are complementary rather than redundant [7]. Furthermore, they observed that visual features increase the robustness of NMT by acting as input noise.

2.2 Multimodal Translation on Indian Languages

The early attempts for MMT for Indian languages relied on synthetic data [16], which was generated by using a text-based PBSMT model. Then, they use the projected image feature for hidden representations of encoder and decoder [8] to generate Multimodal MT. Initializing the hidden representation of the decoder generated better results compared to that of the encoder. Subsequently, it was observed that using multimodal methods achieves better results [47, 68] compared to UNMT or image captioning. Doubly Attentive decoder [9] architecture, which was proven to generate a better translation for other languages, was adopted for English-Hindi [42, 43]. They use pre-trained

word embedding learned from IITB Corpus [39] to initialize the models and augment datasets by using Giza++ Tool [53], to obtain phrase pairs. Furthermore, the pre-trained VGG19 [72] model is used to extract image features.

It was observed that using text-based pre-trained models achieved better results for Hindi VG Test set compared to the multimodal model trained only on multimodal dataset [23]. However, it obtained lower results for the Challenge set. Furthermore, the text-based model was trained on a larger bilingual corpus, which could be a major contributing factor to better results. Multimodal Transformer architecture was adopted for English-Hindi [71], which improved the results compared to its UNMT baseline. There are other works that utilized multiple captions from the dataset to train the models [73]. It was observed that multiple captions introduce lexical diversity, which helps to generate robust translation.

Only a limited number of studies have employed large-scale pre-training. A model was pre-trained on IITB corpus [39], and it was observed that large-scale pre-training is more important for better translation in contrast to image features [58]. In another study, the multimodal features are found to be less relevant as a good pre-trained model can handle ambiguity issues [5]. A pre-trained mBART model was used to initialize the weights [26]. Then, they used unimodal data for the first stage of fine-tuning. Finally, they concatenated object tags with source text before last-stage fine-tuning. Those methods are entirely unimodal except for the usage of Object detection models for object tag extraction at data pre-processing stages. Thereafter, a similar method was adopted for English-Bengali [59]. Inspired from [25], a sequence-to-sequence model was used for MMT [60]. In this method, the input is the concatenated source and target side, and the output is object tags, which act as dummy image features on unimodal datasets. They follow the object tag concatenation-based method, which resulted in huge BLEU gains compared to the text-only model. However, the text-only models were not trained on very large corpus in required languages. This makes them unable to disambiguate sentences properly.

There were some attempts in the reverse direction, where Spatial attention was used for image features [51] for Hindi to English Multimodal News Translation. In the meantime, some MMT methods were proposed for other low-resource Indian languages. A transliteration-based phrase augmentation method was adopted for English-Assamese [44] with Assamese **VisualGenome** (VG) dataset [45] and English-Bengali [41]. A similar method was used for English-Hindi [46], and this method was able to generate better results compared to UNMT. A comparison of results was performed on the Mizo VG dataset [36] with Transformer, BiLSTM, and PBSMT models [40, 62, 77], where the transformer outperforms the other models. However, it was observed that NMT-based models generate a poor-quality translation for longer sentences. A multilingual multimodal NMT system was introduced for English and Dravidian languages [11], which generated better results compared to UNMT and single-language MMT. All these methods assumed the presence of clean and grammatical training data, which may not be very suitable for real-world use cases.

2.3 Context-aware Translation

Apart from using multimodal information as context, there have been several attempts to use unimodal information as a context. For example, the previous and next sentences can serve as a useful context for document translation. Similarly, previous sentences in a conversation can serve as a context for chat translation. The context is generally incorporated by concatenation [24, 78] or using a different encoder for context. Hierarchical Attention Networks can achieve better results [52] compared to context-agnostic NMT. It was observed that performance gains from multi-encoder models cannot be entirely attributed to the contextual information. Setting the weight of the context to zero during testing does not have a significant negative effect on the results [20, 48]. Similarly, training and testing with random context also improve the results [1].

The standard metrics for NMT were unable to capture the improvements at the discourse level. [19] proposed a discourse-aware framework for performance evaluation.

2.4 Noisy NMT

The presence of noise in the data has adverse effects on the translation quality [3, 35]. Early attempts for Noisy MT relied on Normalization of text [81]. Due to the lack of parallel corpora for Noisy NMT, most of the research relies on synthetically generated noise. This method helped to achieve better results in translating erroneous speech transcripts [75, 83]. Subsequently, the noise was added with texts for different datasets with multiple types of synthetic noise [13, 33]. Generally, noisy text translation can be viewed as a domain adaptation task. Fine-tuning on noisy data improves the robustness of the model [4, 15, 28]. Other methods include the construction of source and target side adversarial examples [12] for noisy NMT. Subsequently, Insertion and Deletion predictors were added to the NMT, which improved results on noisy inputs [82]. It was observed that multi-tasking with translation and denoising objectives helps to generate robust translation [61]. Furthermore, they observed improvements after adding synthetic noise to the dataset and using Dual Channel Decoding. Noisy text translation is an unexplored area in Indian languages. It was observed that fine-tuning with noisy data improves the results for Conversational NMT [21]. However, since the training datasets were synthetically generated, the improvements are mostly due to domain adaptation.

3 Dataset Details

We use the VG datasets containing Bengali [69], Hindi [57], and Malayalam [56] sentences for fine-tuning. The train, validation, test, and challenge subsets of each dataset contain 29,230, 998, 1,595, and 1,400 sentences, respectively. The challenge subset comprises sentences with potential ambiguity (such as “court”) that may or may not be resolved using textual information alone. Each row in a dataset contains the following fields:

- **Source text:** The utterance in English.
- **Target text:** The corresponding utterance in Hindi, Bengali, or Malayalam.
- **Image ID:** The filename of the image relevant to the utterance.
- **Coordinates of image:** The dataset provides the bounding box of the specific image region on which the utterance is based.

A sample of the dataset is presented in Figure 1, and we describe some additional dataset statistics in Table 1. We frequently noticed that the image size within the box’s boundaries is insufficient. More information might be available if the entire image were used rather than just the cropped area contained within the bounding box. A study [22] observed that using the cropped image feature of ResNet-50 yields better results compared to the full image. However, the difference is marginal in most experimental settings, and using full images achieved slightly better results in one setting. Therefore, we perform our experiments in both setups.

3.1 Noisy Data Generation

We generate noise in the datasets as they appear in social media. Social media captions frequently misspell words with two identical consecutive characters, write words without vowels, and omit articles (a, an, the). Since our work intends to find if the visual context can help to translate noisy text, we did not add noise to the images.

3.1.1 Low Noise. We iterate through all the words in a sentence and remove them if we encounter an article with 0.2 probability. If the word is non-article or article but not removed in the previous step, we look for vowel characters in the word and remove them with 0.1 probability. Finally, we



Fig. 1. Example of combined Hindi, Bengali, and Malayalam dataset.

Table 1. Statistics of Hin: Hindi, Beng: Bengali, and Mala: Malayalam VG Datasets

Subset	#Sentences	#Average Words per Utterance			
		English	Hindi	Bengali	Malayalam
<i>Train</i>	28,930	4.95	5.03	3.94	3.70
<i>Valid</i>	998	4.93	4.99	3.94	3.63
<i>Test</i>	1,595	4.92	4.92	4.02	3.57
<i>Challenge</i>	1,400	5.85	6.17	4.76	4.32

check if the word contains the same letters at two consecutive positions. In that case, we remove one of them with 0.2 probability. We started with 0.3 probability for each noise generation operation and randomly selected 20 samples from the dataset. We engaged some human annotators who are active on social media to rate noised sentences in terms of naturalness from 1 to 5. A rating of 5 indicates that these types of noisy sentences often appear in social media, and a rating of 1 indicates that the appearance of these types of noisy sentences is quite unlikely. We keep reducing the probability of each type of noise by 0.1 until we reach at least a 4.5 rating on average.

The statistics of added noise with respect to the actual sentences are reported in Table 2. As we acquire the noisy version of the data through subjective human judgment, we use the BLEU [55] metric to report n-gram matches as a measure of similarity. Similarly, TER [74] indicates the number of edits required to obtain the original source from the corrupted text. chrF2 [64, 65] is a character-level F-score metric. To maintain a consistent distribution of noise throughout training, validation, and evaluation, we introduce noise to all the available subsets of the data. In low noise settings, the resulting BLEU scores typically fall within the range of 65.9 to 67.5, while chrF2 scores

Table 2. Statistics of the Source Text After Addition of Noise when Compared with Original Source Text

Dataset	Low Noise			High Noise		
	BLEU	chrF2	TER	BLEU	chrF2	TER
<i>Train</i>	67.5	87.3	16.2	38.1	71.1	34.9
<i>Valid</i>	66.9	87.1	16.4	36.7	70.6	35.5
<i>Test</i>	67.0	87.0	16.5	36.9	70.6	35.3
<i>Challenge</i>	65.9	87.7	16.5	37.3	71.9	36.0

range from 87.0 to 87.7; a lower TER score, in this case, indicates a lower amount of noise, which ranges from 16.2 to 16.5 in low noise settings.

3.1.2 High Noise. We remove articles with 0.3 probability. Unlike low noise, we do not make further modifications to articles if they are not removed because we are already adding other noises in high amount. Then, we remove vowels and repeat characters with 0.3 probability each. In scenarios with high noise levels, the achieved BLEU scores usually lie between 36.7 to 38.1, and the corresponding chrF2 scores range from 70.6 to 71.9. A lower TER score in such cases suggests a reduced amount of noise, which typically falls within the range of 34.9 to 36.0 in high-noise settings.

3.2 Pre-processing

We use `langid.py` [50], a language identification tool to detect if any sentence pair contains the wrong language. We found some alignment issues in the corpus, such as the presence of source sentences in the target field of the dataset. Some human annotators manually annotated those instances by translating them into the target language, taking the corresponding image into consideration. Furthermore, the source data was unavailable in a few instances of Hindi VG. Since the Bengali VG is aligned with Hindi VG, we copy the corresponding source sentence from Bengali VG. Finally, while Malayalam VG contains the same set of sentences, image pairs are Hindi and Bengali VG; they are not aligned. We align them to match the order of Bengali VG. Aligning the datasets helps to reduce memory usage, as visual features can be common between all languages. Finally, we follow the same pre-processing steps that are performed to train the pre-trained model. Since Hindi, Bengali, and Malayalam are written in different scripts, we convert all target side data to the Devanagari script by a publicly available tool [38]. We use Hindi, Bengali, and Malayalam VG datasets to fine-tune all our models. For cropped experiments, we crop image portions within bounding boxes. We extract (cropped as well as full) image features from `vit_base_patch16_224` model [37] unless mentioned otherwise.

4 Methodology

In the literature, previous probing methods focused on the usability of images when some information from the source was deprived of the model. For example, the words with *color*, *person*, and so on, which can be available in the image, are masked. Although the models function well in these circumstances, masking words in the source sentence explicitly tells the model that some information is hidden. This forces the model to look for pertinent data in the image. However, in most situations, there will not be explicit mask tokens that indicate that a word is absent. On the other hand, simply dropping the words (without mask token) and expecting the models to retrieve them does not align with the MT objective. For example, in the source sentence *chinese characters on a red vehicle*, if we remove the word *red*, then *chinese characters on a vehicle* is also a valid sentence, and the model should not attempt to add information related to color in the generated translation. Doing this would lead to undesirable information in the generated sentence, which is

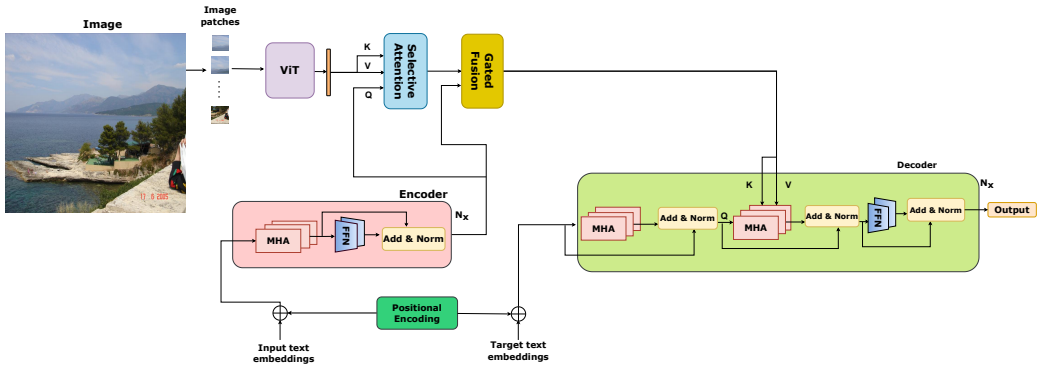


Fig. 2. Selective attention architecture for Multimodal MT. (softcopy after zooming-in exhibits better display).

absent in the source. Therefore, instead of providing any information if the input sentence is noisy, we try to determine whether the models can use information from images in the presence of noise in the source text.

4.1 Baseline

We use a transformer-4x model¹ trained on Samanantar [67] corpus containing 49.6 million sentence pairs between English and 11 Indian languages as our baseline. The model consists of six encoder and six decoder layers. We initialize weights from this model for all our experiments and use the same dictionary for byte-pair-encoding. For the English-Malayalam Challenge set, it achieves a 26.3 BLEU score, whereas the SOTA system [26, 60] achieves a 20.4 BLEU score. This outperforms SOTA by 5.9 BLEU points even before fine-tuning on any data. After fine-tuning the model on only source-target sentence pairs on respective VG data, we observe that it outperforms all the existing multimodal models by a large margin. Considering the additional overhead of handling images and extracting visual features, we attempt to find if they are worth it when we can simply use a text-based model only and achieve better results since text-only data is available in large quantities.

4.2 Unimodal Fine-tuning

First, we follow pre-processing steps as mentioned in Section 3.2 and combine data from Hindi, Bengali, and Malayalam VG datasets. Then, we initialize the encoder and decoder weights from the baseline model (including word embedding) and continue training. This model surpasses all existing baselines on all the languages and both Test and Challenge subsets by large margins. This model does not use any information from images. Unimodal data are widely available, whereas multimodal datasets only have a small number of examples. If multimodal models are not trained on unimodal data, it is difficult to produce reliable outputs.

4.3 Multimodal Fine-tuning with Selective Attention

For Multimodal Translation, we follow the Selective Attention architecture as introduced in [49]. The architecture of the model has been shown in Figure 2. For a given data point of <image, source sentence, target sentence> in the training dataset, at first, we obtain the word and positional embedding for the given source text. The Positional encoding describes the location or position of a token in a sequence in such a way that each position is assigned a unique representation. In Transformers positional encoding, each position is mapped to a vector. As an example, for an input

¹<https://ai4b-public-nlu-nlg.objectstore.e2enetworks.net/en2indic.zip>

sequence of length L , the positional encoding of the k th object within this sequence is given by sine and cosine functions of varying frequencies as follows:

$$P(k, 2i) = \sin\left(\frac{k}{n^{2i/d}}\right); P(k, 2i + 1) = \cos\left(\frac{k}{n^{2i/d}}\right); \quad (1)$$

where, n is a user-defined scalar, and i is used for mapping the indices of the column. The obtained representation is forwarded through the encoder layer of the Transformer [80]. Each encoder layer consists of three major components: **Multi-Head Attention (MHA)** layers, Fully Connected layers and the Add and Norm layer. MHA is a module for attention mechanisms that simultaneously runs through multiple attention mechanisms. After that, the independent attention outputs are combined and transformed linearly into the expected dimension. Intuitively, having multiple attention heads makes it possible to focus on different parts of the sequence simultaneously [80].

After the self-attention mechanism has computed the attention scores between each input token and every other token in the sequence, each token's representation is passed through the feedforward network. The output of the feedforward network is then added back to the original input representation of the token, which helps to capture more complex relationships between tokens in the sequence. The Add and Norm layer acts as a residual connection and performs normalization of the outputs at the corresponding sub-layer.

To obtain the representation H_{img} of the corresponding input image, we pass the image through the pre-trained ViT. It is one of the recent architectures for image classification, which employs a Transformer-like architecture over patches of the image. Any image is first split into fixed-size patches. Each patch is then linearly embedded, after which the position embeddings are added, and the resultant vector is fed to a standard Transformer encoder.

Then, we add an attention layer, where H_{text} is the <query> and H_{img} is used as <key and value> [37]. We perform selective attention on both the H_{img} and H_{text} , where H_{img} acts as <key and value> pair and H_{text} acts as a query for the computation of the attention as follows:

$$H_{img}^{attn} = softmax\left(\frac{H_{text}Q(H_{img}K)^T}{\sqrt{d_k}}\right)H_{img}V; \quad (2)$$

where, Q , K , and V are the layer-specific query, key, and value vectors. Subsequently, text-encoder output H_{text} and selective-attention output H_{img}^{attn} are combined with gated fusion to obtain the final encoder representation Enc_{out} .

$$Enc_{out} = (1 - \lambda) H_{text} + \lambda H_{img}^{attn}; \lambda = sigmoid(W_t H_{text} + W_i H_{img}^{attn}); \quad (3)$$

where, λ indicates the importance of visual features compared to text features. Here, W_t and W_i are trainable parameters.

Finally, Enc_{out} is forwarded to the decoder. In addition to the encoder sublayers, the decoder contains an MHA layer for Encoder-Decoder cross-attention. This layer is similar to encoder-decoder attention in standard transformer architecture, except it attends to the output of gated fusion instead of the output of the text encoder.

We observe that the model was able to achieve only up to 0.25 BLEU improvements in some subsets compared to the text-only model, and it decreased for other test sets by up to 1.38 BLEU. The results indicate that the images are not very useful for generating better translation in high-resource settings. Therefore, we experiment with another architecture for Multimodal Translation.

4.4 Multimodal Transformer

We experiment with the Multimodal Transformer [84] as the alternate architecture for the same task. The source text features $x^{text} \in \mathbb{R}^{m \times d}$ are obtained by enriching the source text embedding

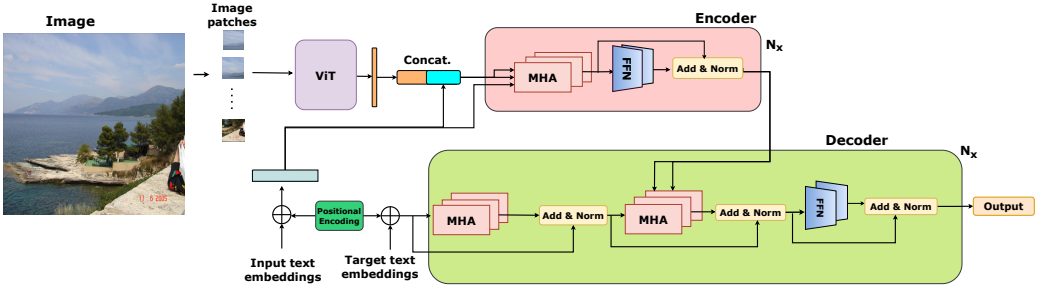


Fig. 3. Multimodal transformer architecture for Multimodal MT. (softcopy after zooming-in exhibits better display).

by the positional embeddings at first. The corresponding image's features x^{img} are extracted by passing the input image through the ViT. Since the image and text features belong to different embedding spaces, we project visual features into the text embedding dimension by a learnable projection matrix, W^{img} to obtain $x^{img}W^{img} \in \mathbb{R}^{n \times d}$. Following this, the features are concatenated with each other to obtain a final input representation x_{com} in $\mathbb{R}^{(m+n) \times d}$, where projected visual features are $x^{img}W^{img} \in \mathbb{R}^{n \times d}$.

In the succeeding phases, we use the combined feature x_{com} as the <key and value> vector for the attention computation in the transformer encoder. The text representation is used as the query vector in the text encoder. The output of the final encoder from the last encoder layer is forwarded as keys and values, while the target text representation is passed as the query vector to the MHA module in the decoder module. The encoder and decoder architectures are inspired by the traditional transformer architecture itself. The detail of the architecture is described in Figure 3.

$$attn_{ij} = \sum_{j=1}^m \text{softmax} \left(\frac{x_i^{com} Q(x_j^{text} K)^T}{\sqrt{d_k}} \right) x_j^{text} V; \quad (4)$$

where, $attn_{ij}$ is the output attention between x_i^{com} and x_j^{text} . Here, Q , K , and V are the layer-specific query, key, and value vectors. We obtain the final encoder representation at the last encoder layer and forward the output to the decoder.

Analogous to the results obtained using the Selective Attention architecture, the Multimodal Transformer architecture was able to achieve only 0.10 BLEU improvements on the test set, on average. However, for the case of the Challenge set, the results actually deteriorated, with a drop of around 0.7 BLEU score.

5 Experiments

This section discusses the results with non-noisy, low-noisy, and high-noisy settings. Furthermore, we report the findings after we replace the actual image with a random image from the dataset; we probe the gate values indicating the relative contribution of visual features compared to textual features. We also report the results by using CLIP as an alternative model for visual feature extraction.

5.1 Non-noisy Settings

We report our results in Table 3 under non-noisy setups. In this setting, both training and testing are done with the original dataset. Text-only baseline achieved 49.08 and 48.42 BLEU scores. While these scores are superior to all the existing models, they are due to large-scale pre-training. Selective attention architecture performs inferior to text-only baseline in non-noisy settings. Although it

Table 3. Results of the Methods with Large-scale Pre-training

Method	Image Type	Hindi		Bengali		Malayalam		Average	
		Test	Chal	Test	Chal	Test	Chal	Test	Chal
SOTA [26, 60]	Object Tags	44.64	51.66	43.9	32.9	41.00	20.4	43.18	34.99
IndicTrans	None	34.74	46.32	26.68	28.91	28.62	26.30	30.01	33.84
+ Text	None	45.79	56.72	50.08	47.78	51.38	40.76	49.08	48.42
+ SelAttn	Crop	46.04	56.36	48.70	47.82	51.56	39.76	48.77	47.98
+ SelAttn	Full	45.30	56.03	49.74	47.48	51.55	39.49	48.86	47.67
+ MMtrans	Crop	45.47	55.61	49.03	47.32	52.40	40.40	48.97	47.78
+ MMtrans	Full	45.10	56.59	50.45	47.24	52.00	39.77	49.18	47.87

Test: Evaluation Subset; **Chal:** Challenge Subset; **Text:** Text-only finetune; **SelAttn:** Selective Attention; **MMtrans:** Multimodal Transformer.

Table 4. Difference in Results Obtained by Multimodal Models and Text-based Model Under Non-noisy Setting

Architecture	Image Feat	Hindi		Bengali		Malayalam		Average	
		Eval	Chal	Eval	Chal	Eval	Chal	Eval	Chal
SelAttn	Crop	0.25	-0.36	-1.38	0.04	0.18	-1.00	-0.31	-0.44
SelAttn	Full	-0.49	-0.69	-0.34	-0.30	0.17	-1.27	-0.22	-0.75
MMtrans	Crop	-0.32	-1.11	-1.05	-0.46	1.02	-0.36	-0.11	-0.64
MMtrans	Full	-0.69	-0.13	0.37	-0.54	0.62	-0.99	0.10	-0.55

Green/Positive indicates that the MMT model achieved better results than the UNMT model.

achieves better results in some subsets, the improvements are insignificant. Subsequently, we report the comparison of the multimodal model with respect to the unimodal model in Table 4. We observe that the image has an overall negative effect on the results. Surprisingly, it performs worse on the challenge subset, which was constructed with ambiguous words. However, the presence of an ambiguous word does not necessarily make a sentence ambiguous, as the proper meaning of the word can be inferred from the surrounding words within the sentence.

The results in Table 4 indicate that the images are mostly redundant when used on top of a strong baseline under non-noisy settings. Our observation is similar to [26], where object tags were concatenated with source tags and trained on top of the mBART model. They achieved 0.52 BLEU improvement on the test set over the text-only model. However, BLEU decreased slightly by 0.06 points on the challenge subset for English-Hindi. Although another study [60] used similar methods and observed huge improvements over text-based models, their baseline was much weaker. Our text-based baseline (without fine-tuning on multimodal data) performs better than their fine-tuned model in some subsets in terms of BLEU score. We found that images do not have a clear positive impact on the results under non-noisy settings. We investigate if the visual context has a positive impact when part of the source text is deprived or distorted via the addition of noise. Specifically, we add noise types that are predominant in social media text such that multimodal models could be useful for text translation in social media.

5.2 Low Noise Settings

The results for low noise settings are reported in Table 5. Selective attention with cropped image features surpassed the text-based models on three subsets (out of six). It is to be noted that the improvements were observed on the test subset, where words were not so ambiguous. It generated lower scores in the challenge subset than the text-based model, which contains ambiguous words.

Table 5. Results of the Methods in Low-noise Settings

Method	Image Type	Hindi		Bengali		Malayalam		Average	
		Test	Chal	Test	Chal	Test	Chal	Test	Chal
Text	None	44.03	50.82	45.29	42.67	47.70	34.76	45.67	42.75
SelAttn	Crop	44.20	50.62	46.07	42.33	49.34	34.19	46.54	42.38
SelAttn	Full	43.48	51.03	45.24	41.72	47.83	34.32	45.52	42.36
MMtrans	Crop	43.25	50.42	45.11	42.20	48.48	34.08	45.61	42.23
MMtrans	Full	42.80	50.68	44.87	41.12	47.91	33.82	45.19	41.87

Test: Evaluation subset; **Chal:** Challenge subset; Text: Text-only fine-tune; SelAttn: Selective Attention; MMtrans: Multimodal Transformer.

Table 6. Difference in Results Obtained by Multimodal Models and UNMT Model Under Noisy Settings

Architecture	Image Feat	Noise	Hindi		Bengali		Malayalam		Average	
			Eval	Chal	Eval	Chal	Eval	Chal	Eval	Chal
SelAttn	Crop	Low	0.17	-0.20	0.78	-0.34	1.64	-0.57	0.87	-0.37
SelAttn	Full	Low	-0.55	0.21	-0.05	-0.95	0.13	-0.44	-0.15	-0.39
MMtrans	Crop	Low	-0.78	-0.40	-0.18	-0.47	0.78	-0.68	-0.06	-0.52
MMtrans	Full	Low	-1.23	-0.14	-0.42	-1.55	0.21	-0.94	-0.48	-0.88
SelAttn	Crop	High	0.07	-0.51	-0.11	-0.78	-0.27	-1.19	-0.1	-0.82
SelAttn	Full	High	-0.04	-0.52	0.44	-0.17	0.44	-0.14	0.28	-0.27

Green/Positive indicates that the MMT model achieved better results than the UNMT model.

Table 7. Results of the Methods in High-noise Settings

Method	Image Type	Hindi		Bengali		Malayalam		Average	
		Test	Chal	Test	Chal	Test	Chal	Test	Chal
Text	None	41.44	45.28	42.49	36.56	44.38	29.35	42.77	37.06
SelAttn	Crop	41.51	44.77	42.38	35.78	44.11	28.16	42.67	36.24
SelAttn	Full	41.40	44.76	42.93	36.39	44.82	29.21	43.05	36.79

Test: Evaluation subset; **Chal:** Challenge subset; Text: Text-only fine-tune; SelAttn: Selective Attention.

This indicates that text-based models are very good at handling most ambiguity when trained on large data.

We report the difference between the multimodal and unimodal models on rows 1–4 of Table 6. The selective attention model with cropped image feature outperformed the text-only model by 0.87 BLEU on average on the test set. While multimodal transformer architecture performed competitively with the text-only baseline in some subsets, the overall results were lower than that of the selective attention model on both the challenge and test sets as depicted in Table 5. We observe that using only cropped image features has only a positive effect on the evaluation set (row 1) on low-noise settings. Similarly, the cropped image features have a higher positive impact on the Malayalam test set as well as a lesser overall negative impact than full images in multimodal transformer architecture. This suggests the superiority of cropped image features under low-noise settings.

5.3 High Noise Settings

Since the Multimodal Transformer model was unable to score the best results in any of the subsets in low-noise settings, we primarily focused on Selective Attention and text-only models. We report the results in Table 7 and the difference with respect to the unimodal model on rows 5–6 of Table 6.

Table 8. Comparison of Using Random Images vs. UNMT Under Noisy Setup

Architecture	Image Feat	Noise	Hindi		Bengali		Malayalam		Average	
			Eval	Chal	Eval	Chal	Eval	Chal	Eval	Chal
SelAttn	Crop	Low	0.26	-0.70	1.13	-0.79	1.27	-0.81	0.89	-0.77
SelAttn	Full	Low	-0.40	-0.36	-0.17	-1.22	-0.93	-2.05	-0.50	-1.21
MMtrans	Crop	Low	-0.69	-0.32	-0.18	-0.49	0.48	-0.08	-0.13	-0.30
MMtrans	Full	Low	-0.79	0.12	-0.72	-1.41	1.01	-0.60	-0.16	-0.63
SelAttn	Crop	High	-0.24	-0.25	-0.32	-0.82	-0.51	-1.29	-0.36	-0.78
SelAttn	Full	High	0.13	-0.41	0.24	-0.49	0.52	-0.81	0.30	-0.57

Green/Positive indicates that random images achieved better results than the unimodal model.

Table 9. Comparison of Using Random Images vs. Actual Images Under Noisy Setup

Architecture	Image Feat	Noise	Hindi		Bengali		Malayalam		Average	
			Eval	Chal	Eval	Chal	Eval	Chal	Eval	Chal
SelAttn	Crop	Low	-0.09	0.50	-0.35	0.45	0.37	0.24	-0.02	0.40
SelAttn	Full	Low	-0.15	0.57	0.12	0.27	1.06	1.61	0.35	0.82
MMtrans	Crop	Low	-0.09	-0.08	0.00	0.02	0.30	-0.60	0.07	-0.22
MMtrans	Full	Low	-0.44	-0.26	0.30	-0.14	-0.80	-0.34	-0.32	-0.25
SelAttn	Crop	High	0.31	-0.26	0.21	0.04	0.24	0.10	0.26	-0.04
SelAttn	Full	High	-0.17	-0.11	0.20	0.32	-0.08	0.67	-0.02	0.30

Green/Positive indicates that actual images achieved better results than random images.

Here, using the cropped image was not found to be very effective compared to the text-based baseline. Selective Attention architecture with cropped images was able to surpass the text-only baseline on only one test set out of six, where the BLEU improvement was only 0.07 points. Using visual features from the full image was more effective than using cropped image features. This is the opposite for low-noise settings. This indicates that the model utilizes visual features more in high-noise settings.

5.4 Probing with Random Images

At first glance, the comparable performance of models using random and relevant images may seem to challenge the assertion that visual context is beneficial for handling noisy inputs. However, a deeper examination reveals a more nuanced interaction between visual context and model behavior. Our experiments demonstrate that incorporating visual context consistently improves performance in noisy settings compared to text-only models, affirming that visual information is indeed “helpful.” However, the observation that random images yield similar results to relevant ones suggests that the observed improvements are not attributable to the semantic relevance of the visual information. Instead, the inclusion of image features—whether random or relevant—appears to enhance model robustness, potentially through regularization or by mitigating the effects of textual noise.

To investigate whether these improvements stem from meaningful visual information, we conducted probing experiments where relevant images were replaced with randomly selected images from the same dataset. The results, as presented in Table 9, reveal that in low-noise settings, random images yield similar performance to relevant images for cropped features. For instance, the Selective Attention architecture achieved comparable results on the test set with random images, with only a marginal drop of -0.40 BLEU on the challenge set. However, when using full images, the use of random images had a more pronounced effect, reducing BLEU by 0.35 and 0.82 on the Test and

Table 10. Results of CLIP Features vs. ViT Feature in High-noise Settings

Method	Image Feature	Hindi		Bengali		Malayalam		Average	
		Test	Chal	Test	Chal	Test	Chal	Test	Chal
Text	None	41.44	45.28	42.49	36.56	44.38	29.35	42.77	37.06
SelAttn	CLIP	41.11	45.21	42.67	36.35	44.61	29.24	42.80	36.93
SelAttn	ViT	41.51	44.77	42.38	35.78	44.11	28.16	42.67	36.24

Test: Evaluation subset; **Chal:** Challenge subset; **Text:** Text-only fine-tune; **SelAttn:** Selective Attention.

Challenge subsets, respectively. This outcome aligns with the observation that textual context alone is often sufficient in low-noise scenarios, as evidenced by the strong performance of text-only models fine-tuned on noisy data.

In high-noise settings, the results with random images remained largely consistent with those obtained using relevant images. Comparisons with text-based baselines, as detailed in Table 8, exhibit trends similar to those in Table 6. These findings indicate that the performance gains achieved with relevant images can also be attained using random images. This suggests that the model does not leverage meaningful visual information but instead utilizes image features as a form of noise, enhancing its robustness during training. Consequently, the absence of correct images during inference has minimal impact on performance.

These results highlight a critical limitation of current multimodal approaches: while visual context improves overall performance, the model appears to rely on generic signal patterns from image features rather than extracting and integrating complementary information. This underscores the need for further research to develop mechanisms that can effectively extract and leverage relevant visual features, ensuring that the visual modality contributes meaningfully rather than merely acting as a source of noise.

5.5 Using CLIP Image Features

We explore whether the relatively low impact of image features in multimodal translation can be attributed to the type of features used in previous experiments. To deepen our understanding, we extend the investigation by incorporating image features derived from CLIP [66], a SOTA model designed for learning transferable visual representations. We chose to perform the experiment on high-noise settings because high-noise settings have been found to be most effective for utilizing image features. We use ViT-B/16 model to extract image features and use it as H_{img} in Equation (2). Our findings, summarized in Table 10, reveal that the use of CLIP-derived image features does not significantly alter the overall observations. The model's performance with these features remains closely aligned with that achieved using ViT-based features from earlier experiments. This consistency underscores a notable insight: the limited influence of image features in multimodal translation may not primarily stem from the feature extraction.

5.6 Gate Values

Given the observed limited contribution of images to the translation quality, we further investigated the behavior of the gating mechanism (λ), as described in Equation (3). The gate value, which dynamically adjusts based on the image and text representations, provides insight into the relative importance assigned to the image modality during translation.

In the high-noise training regime with the full image settings, the average gate values computed over the test set and challenge set were found to be 0.000781 and 0.000742, respectively. Interestingly, when random images were used during inference, the corresponding average gate values were 0.000780 and 0.000738 for the test and challenge sets. These values demonstrate minimal

Table 11. Evaluation Results Comparing Translations Produced with Actual and Random Images on Multi30k Test Set

Option	English-Hindi	English-Bengali
Hypothesis-h1 is much better (enhances adequacy)	12	10
Hypothesis-h1 is better (minor grammatical improvements)	4	2
Both hypotheses are of similar quality	15	25
Hypothesis-h2 is much better (enhances adequacy)	13	11
Hypothesis-h2 is better (minor grammatical improvements)	6	2

variation between the settings, indicating that the model assigns an almost negligible weight to the image modality, regardless of whether the input images are relevant or random. This consistency further supports our hypothesis that the contribution of the image modality is more aligned with regularization rather than providing substantive contextual information to improve translation quality.

5.7 Using Multi30K Dataset for Testing

The primary focus of this work is on multimodal translation from English to Indic languages. The widely used Multi30K dataset [17] has been used in several methods to generate multimodal translation [63], visual pre-training [70], image captioning [76], and so on. However, this dataset supports German, French, and English. Moreover, the pre-trained models employed in our experiments are not tailored for the languages in the Multi30K dataset, making direct comparisons challenging. Despite these constraints, we conducted an exploratory analysis as follows.

We utilized the Multi30k test set to generate translations using our trained models under high-noise settings. Since this dataset lacks references in Indic languages, we relied on human annotations to evaluate the quality of the translations. For each test instance, two hypotheses were generated:

- h1: Translations produced using the actual image associated with the text.
- h2: Translations produced using a randomly sampled image during inference.

A subset of 50 English source sentences from the Multi30K dataset was randomly selected, and translations were evaluated in both the English-Hindi and English-Bengali directions. Two experienced annotators were tasked with comparing the translations for each hypothesis and selecting one of the following options:

- Hypothesis-h1 is much better (enhances adequacy of the translation).
- Hypothesis-h1 is better (with less significant or minor grammatical improvements).
- Both hypotheses are of similar quality.
- Hypothesis-h2 is much better (enhances adequacy of the translation).
- Hypothesis-h2 is better (with less significant or minor grammatical improvements).

The results of the human evaluation are summarized in Table 11. The findings reveal no significant difference between translations generated using actual images and those generated with randomly sampled images. For English-Hindi, translations with actual images are rated better in 12 instances, whereas translations with random images are preferred in 13 instances. Similar trends are observed for English-Bengali. These results support our hypothesis that images primarily act as regularizers rather than providing significant contextual information for improving translation quality.

It is important to note that the overall quality of translations in this analysis was reported to be poor, likely due to the high-noise settings used during training and inference. This limitation restricted the scope of our evaluation and highlighted the need for further investigation to isolate the impact of images. However, such an exploration is beyond the scope of this article.

Table 12. The Results by Removing the Selective Attention Module Entirely (Rows 6–8) and Replacing the Combined Image and Text Feature with Text Features Only, Removing the Image Modality (Rows 4–5)

Method	Training Features	Testing Features	Hindi		Bengali		Malayalam		Average	
			Test	Chal	Test	Chal	Test	Chal	Test	Chal
SelAttn	Actual	Actual	41.40	44.76	42.93	36.39	44.82	29.21	43.05	36.79
SelAttn	Actual	Random	41.53	44.35	43.17	35.90	45.34	28.40	43.35	36.22
SelAttn	Random	Random	41.80	45.60	42.88	36.53	44.86	29.45	43.18	37.19
SelAttn → Text	Actual	None	41.35	44.75	42.84	36.19	45.02	28.59	43.07	36.51
SelAttn → Text	Random	None	41.79	45.60	42.88	36.56	44.73	29.45	43.13	37.20
NoSelAttn	Actual	Actual	41.42	45.31	42.61	35.46	44.32	28.22	42.78	36.33
NoSelAttn	Actual	Random	41.20	45.26	42.71	35.78	43.93	28.03	42.61	36.36
NoSelAttn	Random	Random	42.08	45.36	43.70	36.57	44.67	29.08	43.48	37.00

Test: Evaluation subset; **Chal:** Challenge subset; **Text:** Text-only fine-tune; **SelAttn:** Selective Attention; **NoSelAttn:** Not Selective Attention.

6 Ablation and Discussion

To assess the contribution of different components in our multimodal architecture, we conduct a detailed ablation study, the results of which are presented in Table 12.

6.1 Impact of Removing Visual Features

Rows 4 and 5 in Table 12 present results where the multimodal input is reduced to text-only features by removing the visual modality while retaining the rest of the architecture, including the gating mechanism. Specifically, we replace the fused image-text representation with only text-derived features. We experiment with both actual and randomly sampled training features, and in both cases, we remove visual features entirely at test time. Surprisingly, this simplification results in only marginal changes in BLEU scores across all languages and evaluation settings.

For example, using actual training features and removing image features at test time (row 4), the model achieves an average BLEU score of 43.07 on test sets and 36.51 on challenge sets. This is nearly identical to the performance of the full SelAttn (Selective Attention) mechanism with actual image features (row 1), which scores 43.05 and 36.79, respectively. A similar trend holds when random features are used for training (row 5), with scores of 43.13 (Test) and 37.20 (Challenge). These results suggest that the image modality, in its current integration, contributes little in terms of meaningful semantic enhancement and can be effectively replaced by text-only input without sacrificing performance.

6.2 Impact of Removing Selective Attention

Rows 6–8 of Table 12 evaluate the impact of removing the SelAttn entirely. Here, the gating and fusion modules are retained, but the attention layer that modulates visual information is disabled. Interestingly, this also results in minimal performance degradation. In fact, the best test BLEU average of 43.48 is achieved in the configuration where both training and testing use randomly sampled features without any selective attention (row 8). This slightly outperforms the full SelAttn model, further questioning the efficacy of the attention mechanism in contributing meaningful information from the image stream.

Moreover, when actual features are used for training and testing without SelAttn (row 6), the performance is 42.78 (Test) and 36.33 (Challenge), again showing only a marginal drop from the baseline. These findings imply that the selective attention mechanism offers limited gains in this setup, and its removal does not adversely affect the system's robustness or accuracy.

6.3 Effect of Random Image Features

Another notable pattern across Table 12 is the consistent performance of models that utilize random image features. Whether SelAttn is enabled or not, and regardless of whether training or testing features are randomized, the BLEU scores remain relatively stable. For instance, the SelAttn model with random features at both training and test time (row 3) achieves 43.18 (Test) and 37.19 (Challenge), which are slightly better than the full model with actual features. This observation suggests that the gains observed from using the image stream may not be due to semantic visual information but rather due to the introduction of controlled noise that acts as a regularizer.

6.4 Effectiveness of Context for Machine Translation

Overall, the ablation results indicate that the visual modality and the selective attention mechanism are not strictly necessary for strong performance in our task. Replacing the image features with random noise or removing the attention module altogether does not substantially degrade results. These findings highlight that the benefits attributed to the image modality likely stem from the stochastic regularization effect it introduces rather than from its semantic content. This suggests an alternative use of multimodal inputs: not primarily for semantic grounding, but as a source of structured variability to improve robustness, especially in noisy or low-resource scenarios.

While the observations are surprising, a growing body of research suggests that the actual contribution of context (here, image) may be overstated. A detailed evaluation designed with ambiguous source sentences reveals that although MMT systems can technically process visual information, they often fail to integrate it meaningfully into translation choices, limiting its practical use for disambiguation or grounding [34, 49]. Adding to this perspective, another study systematically replaced relevant images with random ones in a multimodal prompting setup and observed that translation performance remained largely unaffected [2]. This surprising result indicates that visual features may primarily stabilize learning dynamics rather than provide true contextual relevance, aligning with the hypothesis that random images can offer similar benefits.

This phenomenon is not unique to multimodal contexts. In document-level and dialogue translation, where external textual context is incorporated, improvements over sentence-level baselines are frequently modest and inconsistent [32]. Despite leveraging information from surrounding sentences or dialogue turns, models often fail to utilize this additional context effectively. A broader analysis of context-aware systems identifies a similar trend: while these models occasionally handle specific phenomena better, such as pronoun resolution or discourse coherence, their overall performance remains comparable to simpler sentence-level models [30]. This raises questions about how much external context truly contributes to translation quality when averaged across benchmarks. Interestingly, several of these studies have hypothesized that external context might be more beneficial when the input text itself is noisy or ambiguous [7]. However, we observe that even in such cases, image features fail to consistently enhance translations. Note that while the previous study used masked tokens to indicate missing information, which may have helped the model to look for additional information, we provide no such hints to the model to simulate real-life noisy examples. Collectively, these observations support the emerging view that external context in machine translation often functions more as a form of implicit regularization than as a direct source of disambiguating information. The challenge of effectively integrating this context into translation models remains largely unsolved.

6.5 Exploring Dataset Quality

We attempt to find out if the dataset is context-aware enough to facilitate the models to learn from images, and if the test sets are context-aware to truly influence the results if the models are able to generate image-aware translation. We build an annotation web portal to check the adequacy,

Table 13. Details of Annotation for Quality and Image Awareness of the Datasets

		Challenge Set (%)		Test Set (%)		Train Set (%)	
Adequacy	Good	82	74	92	88	86	64
	Medium	8	18	8	8	14	32
	Bad	10	8	0	4	0	4
Fluency	Good	86	74	80	72	62	48
	Medium	14	14	20	22	30	42
	Bad	0	12	0	6	8	10
Need of Image	Yes		6		0		0
	Maybe		4		4		10
	No		84		94		90
	Not Reflected		6		2		0


fluency, and image awareness of the dataset by manual evaluation. We provide the annotators with the source, target, and image. We ask the annotator to rate adequacy and fluency on a scale of three - *good*, *medium*, and *bad*. Furthermore, we ask the annotators if they think the image is needed to generate the translation. The answer should be *No* if they would translate it the same even if the image were absent. The answer should be *Yes* if they would have translated it differently or be in confusion without the availability of the image. The answer should be *maybe* when they would have translated similarly anyway, but the image helped to clear the confusion. Finally, the answer is *not reflected* when they believe that the translation could be better with the information from the image, but the same is not reflected in the reference.

We report our observations in Table 13. While the dataset is mostly fluent and adequate, the source text is usually enough to translate it. In the train set, 90% of the examples do not require an image to translate correctly, while in 10%, it would have been somewhat helpful but can be translated without it. This prevents the models from getting proper training signals to use multimodal information effectively. In the challenge set of VG datasets, only 6% of the dataset would require the image to translate it, whereas 4% would have been somewhat helpful. This is 0% and 4% for the test set. There are many instances where an image is essential for accurately translating a sentence. However, the reference translation does not adequately capture this image-aware context. We illustrate one such example in Figure 4. This occurs in 6% and 2% of the examples in Challenge and Test sets, respectively. We generate this report on the first 50 examples of Train, Challenge, and Test Sets of Hindi VG and Bengal VG. This indicates that only a small percentage of examples require images for correct translation. However, the usage of examples that are difficult to translate would encourage the models to use visual features effectively, even when it is trained on a large unimodal corpus. However, to the best of our knowledge, no such data is publicly available.

6.6 Experimental Setup

We try to keep the hyperparameters the same in every model for a fair comparison. We set the learning rate to 3×10^{-5} , which performed the best for the baseline model [67] during fine-tuning. We set warmup updates to 4,000 and the initial learning rate to 10^{-7} with inverse square root as the learning rate scheduler. Due to the small size of the dataset, we validate and save checkpoints at frequent intervals, at multiples of 500 steps or at the end of an epoch. We continue training with Early Stopping till the validation loss does not improve for five consecutive checkpoints. We use *Label Smoothed Cross Entropy* as the loss function and set 0.1 as the *smoothing_value*. We set the *batch_size* to 256 tokens and *update_frequency* to 2. Effectively, this makes the *batch_size* of 512

Multimodal Translation - Annotation Portal



Source : the characters are chinese Target: চরিত্রগুলি চীনা

Adequacy: ▼

Fluency: ▼

Would you need image to translate this sentence? ▼

Fig. 4. Example of the annotation process. This example is obtained from the Challenge Subset of Bengali VG. It is to be noted that the reference is wrong since the image indicates that *character* refers to words in the banner. However, the reference is referring the word *character* to *protagonist* (of movies, stories, etc.)

tokens per step. We use a single GeForce RTX 2080 Ti with 11 GB GPU RAM. All experiments are conducted on Ubuntu 18.04.1 LTS x86_64. We perform memory-efficient FLOP16 training and use the fairseq [54] library for our implementations. We set 42 as the *random_seed* and 0.7 as the *keeping_probability* of dropout. During testing, we used beam search with a *beam_width* of 5.

7 Conclusion

Multimodal NMT can be useful for translating social media content. Previous studies suggest that multimodal models generate better translations compared to unimodal NMT. However, this does not always hold true, especially when the NMT model is trained on a large corpus. We observe that images are often redundant in such high-resource settings. Our findings show that two well-known MMT models do not consistently outperform a strong text-based baseline. We also find that visual context becomes valuable when the source text is noisy. Experimentally, cropped image features achieve the best results in low-noise settings, while full image features yield the best performance in high-noise settings. This indicates that the impact of visual features increases with the level of noise in the source text. However, we found that the improvements were not due to meaningful information from the images, as similar results were achieved even with random image features. This suggests that current models do not yet use images effectively, even when the text becomes less adequate due to noise. One possible solution is to employ datasets where the image is essential, thereby enabling models to learn to integrate visual features more effectively. In future work, we aim to explore the effectiveness of multilingual multimodal NMT across other languages and datasets under noisy conditions.

References

- [1] Ramakrishna Appicharla, Baban Gain, Santanu Pal, and Asif Ekbal. 2023. A Case Study on Context Encoding in Multi-Encoder based Document-Level Neural Machine Translation. In *Proceedings of Machine Translation Summit XIX, Vol. 1: Research Track, Macau SAR, China*. Asia-Pacific Association for Machine Translation. 160–172. Retrieved from <https://aclanthology.org/2023.mtsummit-research.14/>
- [2] Folco Bertini Baldassini, Mustafa Shukor, Matthieu Cord, Laure Soulier, and Benjamin Piwowarski. 2024. What makes multimodal In-context learning work? In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024 - Workshops*, Seattle, WA, USA. 1539–1550. arXiv:2404.15736. DOI : <https://doi.org/10.1109/CVPRW63382.2024.00161>
- [3] Yonatan Belinkov and Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine translation. In *Proceedings of 6th International Conference on Learning Representations, ICLR 2018*. Vancouver, BC, Canada. Retrieved from <https://openreview.net/forum?id=BJ8vJebC->
- [4] Alexandre Berard, Ioan Calapodescu, and Claude Roux. 2019. Naver labs europe’s systems for the WMT19 machine translation robustness task. In *Proceedings of the 4th Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*. Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, André Martins, Christof Monz, Matteo Negri, Aurélie Névéal, Mariana Neves, Matt Post, Marco Turchi, and Karin Verspoor (Eds.), Association for Computational Linguistics, Florence, Italy, 526–532. DOI : <https://doi.org/10.18653/v1/W19-5361>
- [5] Paritosh Bisht and Arun Solanki. 2022. Exploring Practical Deep learning approaches for english-to-hindi image caption translation using transformers and object detectors. In *Proceedings of the Applications of Artificial Intelligence and Machine Learning*. Bhuvan Unhelker, Hari Mohan Pandey, and Gaurav Raj (Eds.), Springer Nature Singapore, Singapore, 47–60.
- [6] 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 1st Conference on Machine Translation: Volume 2, Shared Task Papers*. Ondřej Bojar, Christian Buck, Rajen Chatterjee, Christian Federmann, Liane Guillou, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Aurélie Névéal, Mariana Neves, Pavel Pecina, Martin Popel, Philipp Koehn, Christof Monz, Matteo Negri, Matt Post, Lucia Specia, Karin Verspoor, Jörg Tiedemann, and Marco Turchi (Eds.), Association for Computational Linguistics, Berlin, Germany, 627–633. DOI : <https://doi.org/10.18653/v1/W16-2358>
- [7] 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)*. Jill Burstein, Christy Doran, and Thamar Solorio (Eds.), Association for Computational Linguistics, Minneapolis, Minnesota, 4159–4170. DOI : <https://doi.org/10.18653/v1/N19-1422>
- [8] 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*. Martha Palmer, Rebecca Hwa, and Sebastian Riedel (Eds.), Association for Computational Linguistics, Copenhagen, Denmark, 992–1003. DOI : <https://doi.org/10.18653/v1/D17-1105>
- [9] 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)*. Regina Barzilay and Min-Yen Kan (Eds.), Association for Computational Linguistics, Vancouver, Canada, 1913–1924. DOI : <https://doi.org/10.18653/v1/P17-1175>
- [10] Iacer Calixto, Miguel Rios, and Wilker Aziz. 2019. Latent variable model for multi-modal translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Anna Korhonen, David Traum, and Lluís Màrquez (Eds.), Association for Computational Linguistics, Florence, Italy, 6392–6405. DOI : <https://doi.org/10.18653/v1/P19-1642>
- [11] Bharathi Raja Chakravarthi, Ruba Priyadarshini, Bernardo Stearns, Arun Jayapal, Sridevy S., Mihael Arcan, Manel Zarrouk, and John P. McCrae. 2019. Multilingual multimodal machine translation for dravidian languages utilizing phonetic transcription. In *Proceedings of the 2nd Workshop on Technologies for MT of Low Resource Languages*. Alina Karakanta, Atul Kr. Ojha, Chao-Hong Liu, Jonathan Washington, Nathaniel Oco, Surafel Melaku Lakew, Valentin Malykh, and Xiaobing Zhao (Eds.), European Association for Machine Translation, Dublin, Ireland, 56–63. Retrieved from <https://aclanthology.org/W19-6809/>
- [12] Yong Cheng, Lu Jiang, and Wolfgang Macherey. 2019. Robust neural machine translation with doubly adversarial inputs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Anna Korhonen, David Traum, and Lluís Màrquez (Eds.), Association for Computational Linguistics, Florence, Italy, 4324–4333. DOI : <https://doi.org/10.18653/v1/P19-1425>

- [13] Yong Cheng, Zhaopeng Tu, Fandong Meng, Junjie Zhai, and Yang Liu. 2018. Towards robust neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Iryna Gurevych and Yusuke Miyao (Eds.), Association for Computational Linguistics, Melbourne, Australia, 1756–1766. DOI: <https://doi.org/10.18653/v1/P18-1163>
- [14] Bernard Comrie. 2009. *The World's Major Languages*. Routledge.
- [15] Raj Dabre and Eiichiro Sumita. 2019. NICT's supervised neural machine translation systems for the WMT19 translation robustness task. In *Proceedings of the 4th Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*. Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, André Martins, Christof Monz, Matteo Negri, Aurélie Névél, Mariana Neves, Matt Post, Marco Turchi, and Karin Verspoor (Eds.), Association for Computational Linguistics, Florence, Italy, 533–536. DOI: <https://doi.org/10.18653/v1/W19-5362>
- [16] Koel Dutta Chowdhury, Mohammed Hasanuzzaman, and Qun Liu. 2018. Multimodal neural machine translation for low-resource language pairs using synthetic data. In *Proceedings of the Workshop on Deep Learning Approaches for Low-Resource NLP*. Reza Haffari, Colin Cherry, George Foster, Shahram Khadivi, and Bahar Salehi (Eds.), Association for Computational Linguistics, Melbourne, 33–42. DOI: <https://doi.org/10.18653/v1/W18-3405>
- [17] 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*. Anya Belz, Erkut Erdem, Krystian Mikolajczyk, and Katerina Pastra (Eds.), Association for Computational Linguistics, Berlin, Germany, 70–74. DOI: <https://doi.org/10.18653/v1/W16-3210>
- [18] Desmond Elliott and Ákos Kádár. 2017. Imagination improves multimodal translation. In *Proceedings of the 8th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Greg Kondrak and Taro Watanabe (Eds.), Asian Federation of Natural Language Processing, Taipei, Taiwan, 130–141. Retrieved from <https://aclanthology.org/I17-1014/>
- [19] Patrick Fernandes, Kayo Yin, Emmy Liu, André Martins, and Graham Neubig. 2023. When does translation require context? A data-driven, multilingual exploration. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.), Association for Computational Linguistics, Toronto, Canada, 606–626. DOI: <https://doi.org/10.18653/v1/2023.acl-long.36>
- [20] Baban Gain, Ramakrishna Appicharla, Soumya Chennabasavaraj, Nikeshe Garera, Asif Ekbal, and Muthusamy Chelliah. 2022. Investigating effectiveness of multi-encoder for conversational neural machine translation. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*. Philipp Koehn, Loïc Barrault, Ondřej Bojar, Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Alexander Fraser, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Paco Guzman, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Tom Kocmi, André Martins, Makoto Morishita, Christof Monz, Masaaki Nagata, Toshiaki Nakazawa, Matteo Negri, Aurélie Névél, Mariana Neves, Martin Popel, Marco Turchi, and Marcos Zampieri (Eds.), Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Hybrid), 949–954. Retrieved from <https://aclanthology.org/2022.wmt-1.90/>
- [21] Baban Gain, Ramakrishna Appicharla, Soumya Chennabasavaraj, Nikeshe Garera, Asif Ekbal, and Muthusamy Chelliah. 2022. Low resource chat translation: A benchmark for hindi-english language pair. In *Proceedings of the 15th Biennial Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*. Kevin Duh and Francisco Guzmán (Eds.), Association for Machine Translation in the Americas, Orlando, USA, 83–96. Retrieved from <https://aclanthology.org/2022.amta-research.7/>
- [22] Baban Gain, Dibyanayan Bandyopadhyay, and Asif Ekbal. 2021. Experiences of adapting multimodal machine translation techniques for hindi. In *Proceedings of the 1st Workshop on Multimodal Machine Translation for Low Resource Languages (MMTLRL 2021)*. Thoudam Doren Singh, Cristina España i Bonet, Sivaji Bandyopadhyay, and Josef van Genabith (Eds.), INCOMA Ltd., Online (Virtual Mode), 40–44. Retrieved from <https://aclanthology.org/2021.mmtlrl-1.7/>
- [23] Baban Gain, Dibyanayan Bandyopadhyay, and Asif Ekbal. 2021. IITP at WAT 2021: System description for english-hindi multimodal translation task. In *Proceedings of the 8th Workshop on Asian Translation (WAT2021)*. Toshiaki Nakazawa, Hideki Nakayama, Isao Goto, Hideya Mino, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Shohei Higashiyama, Hiroshi Manabe, Win Pa Pa, Shantipriya Parida, Ondřej Bojar, Chenhui Chu, Akiko Eriguchi, Kaori Abe, Yusuke Oda, Katsuhito Sudoh, Sadao Kurohashi, and Pushpak Bhattacharyya (Eds.), Association for Computational Linguistics, Online, 161–165. DOI: <https://doi.org/10.18653/v1/2021.wat-1.18>
- [24] Baban Gain, Rejwanul Haque, and Asif Ekbal. 2021. Not all contexts are important: The impact of effective context in conversational neural machine translation. In *Proceedings of the 2021 International Joint Conference on Neural Networks (IJCNN)*. 1–8. DOI: <https://doi.org/10.1109/IJCNN52387.2021.9534444>
- [25] Stig-Arne Grønroos, Benoit Huet, Mikko Kurimo, Jorma Laaksonen, Bernard Meriardo, 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 3rd Conference on Machine Translation: Shared Task Papers*. Ondřej

Bojar, Rajen Chatterjee, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Christof Monz, Matteo Negri, Aurélie Nèveól, Mariana Neves, Matt Post, Lucia Specia, Marco Turchi, and Karin Verspoor (Eds.), Association for Computational Linguistics, Belgium, Brussels, 603–611. DOI: <https://doi.org/10.18653/v1/W18-6439>

- [26] Kshitij Gupta, Devansh Gautam, and Radhika Mamidi. 2021. ViTA: Visual-linguistic translation by aligning object tags. In *Proceedings of the 8th Workshop on Asian Translation (WAT2021)*. Toshiaki Nakazawa, Hideki Nakayama, Isao Goto, Hideya Mino, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Shohei Higashiyama, Hiroshi Manabe, Win Pa Pa, Shantipriya Parida, Ondřej Bojar, Chenhui Chu, Akiko Eriguchi, Kaori Abe, Yusuke Oda, Katsuhito Sudoh, Sadao Kurohashi, and Pushpak Bhattacharyya (Eds.), Association for Computational Linguistics, Online, 166–173. DOI: <https://doi.org/10.18653/v1/2021.wat-1.19>
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA, 2016, 770–778. DOI: [10.1109/CVPR.2016.90](https://doi.org/10.1109/CVPR.2016.90)
- [28] Jindřich Helcl, Jindřich Libovický, and Martin Popel. 2019. CUNI system for the WMT19 robustness task. In *Proceedings of the 4th Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*. Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, André Martins, Christof Monz, Matteo Negri, Aurélie Nèveól, Mariana Neves, Matt Post, Marco Turchi, and Karin Verspoor (Eds.), Association for Computational Linguistics, Florence, Italy, 539–543. DOI: <https://doi.org/10.18653/v1/W19-5364>
- [29] Po-Yao Huang, Frederick Liu, Sz-Rung Shiang, Jean Oh, and Chris Dyer. 2016. Attention-based multimodal neural machine translation. In *Proceedings of the 1st Conference on Machine Translation: Volume 2, Shared Task Papers*. Ondřej Bojar, Christian Buck, Rajen Chatterjee, Christian Federmann, Liane Guillou, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Aurélie Nèveól, Mariana Neves, Pavel Pecina, Martin Popel, Philipp Koehn, Christof Monz, Matteo Negri, Matt Post, Lucia Specia, Karin Verspoor, Jörg Tiedemann, and Marco Turchi (Eds.), Association for Computational Linguistics, Berlin, Germany, 639–645. DOI: <https://doi.org/10.18653/v1/W16-2360>
- [30] Jingjing Huo, Christian Herold, Yingbo Gao, Leonard Dahlmann, Shahram Khadivi, and Hermann Ney. 2020. Diving deep into context-aware neural machine translation. In *Proceedings of the 5th Conference on Machine Translation*. Loïc Barrault, Ondřej Bojar, Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Alexander Fraser, Yvette Graham, Paco Guzman, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, André Martins, Makoto Morishita, Christof Monz, Masaaki Nagata, Toshiaki Nakazawa, and Matteo Negri (Eds.), Association for Computational Linguistics, Online, 604–616. Retrieved from <https://aclanthology.org/2020.wmt-1.71/>
- [31] Julia Ive, Pranava Madhyastha, and Lucia Specia. 2019. Distilling translations with visual awareness. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Anna Korhonen, David Traum, and Lluís Màrquez (Eds.), Association for Computational Linguistics, Florence, Italy, 6525–6538. DOI: <https://doi.org/10.18653/v1/P19-1653>
- [32] Linghao Jin, Jacqueline He, Jonathan May, and Xuezhe Ma. 2023. Challenges in context-aware neural machine translation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Houda Bouamor, Juan Pino, and Kalika Bali (Eds.), Association for Computational Linguistics, Singapore, 15246–15263. DOI: <https://doi.org/10.18653/v1/2023.emnlp-main.943>
- [33] Vladimir Karpukhin, Omer Levy, Jacob Eisenstein, and Marjan Ghazvininejad. 2019. Training on synthetic noise improves robustness to natural noise in machine translation. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*. Wei Xu, Alan Ritter, Tim Baldwin, and Afshin Rahimi (Eds.), Association for Computational Linguistics, Hong Kong, China, 42–47. DOI: <https://doi.org/10.18653/v1/D19-5506>
- [34] Nicholas Kashani Motlagh, Jim Davis, Jeremy Gwinnup, Grant Erdmann, and Tim Anderson. 2024. Assessing the role of imagery in multimodal machine translation. In *Proceedings of the 9th Conference on Machine Translation*. Barry Haddow, Tom Kocmi, Philipp Koehn, and Christof Monz (Eds.), Association for Computational Linguistics, Miami, Florida, USA, 1428–1439. DOI: <https://doi.org/10.18653/v1/2024.wmt-1.130>
- [35] Huda Khayrallah and Philipp Koehn. 2018. On the impact of various types of noise on neural machine translation. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*. Alexandra Birch, Andrew Finch, Thang Luong, Graham Neubig, and Yusuke Oda (Eds.), Association for Computational Linguistics, Melbourne, Australia, 74–83. DOI: <https://doi.org/10.18653/v1/W18-2709>
- [36] Vanlalmuansangi Khenglawt, Sahinur Rahman Laskar, Riyanka Manna, Partha Pakray, and Ajoy Kumar Khan. 2022. Mizo visual genome 1.0 : A dataset for english-mizo multimodal neural machine translation. In *Proceedings of the 2022 IEEE Silchar Subsection Conference (SILCON)*. 1–6. DOI: <https://doi.org/10.1109/SILCON55242.2022.10028882>
- [37] 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 16 × 16 Words: Transformers for image recognition at scale. In *Proceedings of International Conference*

on Learning Representations, ICLR 2021. Virtual Event, Austria. Retrieved from <https://openreview.net/forum?id=YicbFdNTTy>

- [38] Anoop Kunchukuttan. 2020. The IndicNLP Library. Retrieved January 2023 from https://github.com/anoopkunchukuttan/indic_nlp_library/blob/master/docs/indicnlp.pdf
- [39] Anoop Kunchukuttan, Pratik Mehta, and Pushpak Bhattacharyya. 2018. The IIT bombay english-hindi parallel corpus. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC 2018)*, Nicoletta Calzolari, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, H el ene Mazo, Asuncion Moreno, Jan Odijk, Stelios Piperidis, and Takenobu Tokunaga (Eds.), European Language Resources Association (ELRA), Miyazaki, Japan. Retrieved from <https://aclanthology.org/L18-1548/>
- [40] Candy Lalrempuii, Badal Soni, and Partha Pakray. 2021. An improved english-to-mizo neural machine translation. *ACM Transactions on Asian and Low-Resource Language Information Processing* 20, 4, Article 61 (2021), 21 pages. DOI: <https://doi.org/10.1145/3445974>
- [41] Sahinur Rahman Laskar, Pankaj Dadure, Riyanka Manna, Partha Pakray, and Sivaji Bandyopadhyay. 2022. English to bengali multimodal neural machine translation using transliteration-based phrase pairs augmentation. In *Proceedings of the 9th Workshop on Asian Translation*. International Conference on Computational Linguistics, Gyeongju, Republic of Korea, 111–116. Retrieved from <https://aclanthology.org/2022.wat-1.14/>
- [42] Sahinur Rahman Laskar, Abdullah Faiz Ur Rahman Khilji, Darsh Kaushik, Partha Pakray, and Sivaji Bandyopadhyay. 2021. Improved english to hindi multimodal neural machine translation. In *Proceedings of the 8th Workshop on Asian Translation (WAT2021)*. Toshiaki Nakazawa, Hideki Nakayama, Isao Goto, Hideya Mino, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Shohei Higashiyama, Hiroshi Manabe, Win Pa Pa, Shantipriya Parida, Ondrej Bojar, Chenhui Chu, Akiko Eriguchi, Kaori Abe, Yusuke Oda, Katsuhito Sudoh, Sadao Kurohashi, and Pushpak Bhattacharyya (Eds.), Association for Computational Linguistics, Online, 155–160. DOI: <https://doi.org/10.18653/v1/2021.wat-1.17>
- [43] Sahinur Rahman Laskar, Abdullah Faiz Ur Rahman Khilji, Partha Pakray, and Sivaji Bandyopadhyay. 2020. Multimodal neural machine translation for english to hindi. In *Proceedings of the 7th Workshop on Asian Translation*. Toshiaki Nakazawa, Hideki Nakayama, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Win Pa Pa, Ondrej Bojar, Shantipriya Parida, Isao Goto, Hidayat Mino, Hiroshi Manabe, Katsuhito Sudoh, Sadao Kurohashi, and Pushpak Bhattacharyya (Eds.), Association for Computational Linguistics, Suzhou, China, 109–113. DOI: <https://doi.org/10.18653/v1/2020.wat-1.11>
- [44] Sahinur Rahman Laskar, Bishwaraj Paul, Partha Pakray, and Sivaji Bandyopadhyay. 2023. English-assamese multimodal neural machine translation using transliteration-based phrase augmentation approach. *Procedia Computer Science* 218 (2023), 979–988. DOI: <https://doi.org/10.1016/j.procs.2023.01.078>. International Conference on Machine Learning and Data Engineering.
- [45] Sahinur Rahman Laskar, Bishwaraj Paul, Siddharth Paudwal, Pranjit Gautam, Nirmita Biswas, and Partha Pakray. 2021. Multimodal neural machine translation for english–assamese pair. In *Proceedings of the 2021 International Conference on Computational Performance Evaluation (ComPE)*. 387–392. DOI: <https://doi.org/10.1109/ComPE53109.2021.9752181>
- [46] Sahinur Rahman Laskar, Rahul Singh, Md Faizal Karim, Riyanka Manna, Partha Pakray, and Sivaji Bandyopadhyay. 2022. Investigation of english to hindi multimodal neural machine translation using transliteration-based phrase pairs augmentation. In *Proceedings of the 9th Workshop on Asian Translation*. International Conference on Computational Linguistics, Gyeongju, Republic of Korea, 117–122. Retrieved from <https://aclanthology.org/2022.wat-1.15/>
- [47] Sahinur Rahman Laskar, Rohit Pratap Singh, Partha Pakray, and Sivaji Bandyopadhyay. 2019. English to hindi multimodal neural machine translation and hindi image captioning. In *Proceedings of the 6th Workshop on Asian Translation*. Toshiaki Nakazawa, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Nobushige Doi, Yusuke Oda, Ondrej Bojar, Shantipriya Parida, Isao Goto, and Hidayat Mino (Eds.), Association for Computational Linguistics, Hong Kong, China, 62–67. DOI: <https://doi.org/10.18653/v1/D19-5205>
- [48] Bei Li, Hui Liu, Ziyang Wang, Yufan Jiang, Tong Xiao, Jingbo Zhu, Tongran Liu, and Changliang Li. 2020. Does multi-encoder help? A case study on context-aware neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.), Association for Computational Linguistics, Online, 3512–3518. DOI: <https://doi.org/10.18653/v1/2020.acl-main.322>
- [49] Bei Li, Chuanhao Lv, Zefan Zhou, Tao Zhou, Tong Xiao, Anxiang Ma, and Jingbo Zhu. 2022. On vision features in multimodal machine translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (Eds.), Association for Computational Linguistics, Dublin, Ireland, 6327–6337. DOI: <https://doi.org/10.18653/v1/2022.acl-long.438>
- [50] Marco Lui and Timothy Baldwin. 2012. Langid.py: An off-the-shelf language identification tool. In *Proceedings of the ACL 2012 System Demonstrations*. Min Zhang (Ed.), Association for Computational Linguistics, Jeju Island, Korea, 25–30. Retrieved from <https://aclanthology.org/P12-3005/>
- [51] Loitongbam Sanayai Meetei, Salam Michael Singh, Alok Singh, Ringki Das, Thoudam Doren Singh, and Sivaji Bandyopadhyay. 2023. Hindi to english multimodal machine translation on news dataset in low resource setting. *Procedia*

- Computer Science* 218 (2023), 2102–2109. DOI : <https://doi.org/10.1016/j.procs.2023.01.186>. International Conference on Machine Learning and Data Engineering.
- [52] Lesly Miculicich, Dhananjay Ram, Nikolaos Pappas, and James Henderson. 2018. Document-level neural machine translation with hierarchical attention networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.), Association for Computational Linguistics, Brussels, Belgium, 2947–2954. DOI : <https://doi.org/10.18653/v1/D18-1325>
- [53] Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics* 29, 1 (2003), 19–51.
- [54] 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)*. Waleed Ammar, Annie Louis, and Nasrin Mostafazadeh (Eds.), Association for Computational Linguistics, Minneapolis, Minnesota, 48–53. DOI : <https://doi.org/10.18653/v1/N19-4009>
- [55] 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*. Pierre Isabelle, Eugene Charniak, and Dekang Lin (Eds.), Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 311–318. DOI : <https://doi.org/10.3115/1073083.1073135>
- [56] Shantipriya Parida and Ondřej Bojar. 2021. Malayalam Visual Genome 1.0. Retrieved January 2023 from <http://hdl.handle.net/11234/1-3533>
- [57] Shantipriya Parida, Ondřej Bojar, and Satya Ranjan Dash. 2019. Hindi visual genome: A dataset for multimodal english-to-hindi machine translation. *Computación y Sistemas* 23, 4 (2019), 1499–1505.
- [58] Shantipriya Parida, Ondřej Bojar, and Petr Motliceck. 2019. IIdiap NMT system for WAT 2019 multimodal translation task. In *Proceedings of the 6th Workshop on Asian Translation*. Toshiaki Nakazawa, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Nobushige Doi, Yusuke Oda, Ondřej Bojar, Shantipriya Parida, Isao Goto, and Hidayat Mino (Eds.), Association for Computational Linguistics, Hong Kong, China, 175–180. DOI : <https://doi.org/10.18653/v1/D19-5223>
- [59] Shantipriya Parida, Subhadarshi Panda, Satya Prakash Biswal, Ketan Kotwal, Arghyadeep Sen, Satya Ranjan Dash, and Petr Motliceck. 2021. Multimodal neural machine translation system for english to bengali. In *Proceedings of the 1st Workshop on Multimodal Machine Translation for Low Resource Languages (MMTLRL 2021)*. Thoudam Doren Singh, Cristina España i Bonet, Sivaji Bandyopadhyay, and Josef van Genabith (Eds.), INCOMA Ltd., Online (Virtual Mode), 31–39. Retrieved from <https://aclanthology.org/2021.mmtlrl-1.6/>
- [60] Shantipriya Parida, Subhadarshi Panda, Stig-Arne Grönroos, Mark Granroth-Wilding, and Mika Koistinen. 2022. Silo NLP's participation at WAT2022. In *Proceedings of the 9th Workshop on Asian Translation*. International Conference on Computational Linguistics, Gyeongju, Republic of Korea, 99–105. Retrieved from <https://aclanthology.org/2022.wat-1.12/>
- [61] Peyman Passban, Puneeth Saladi, and Qun Liu. 2021. Revisiting robust neural machine translation: A transformer case study. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.), Association for Computational Linguistics, Punta Cana, Dominican Republic, 3831–3840. DOI : <https://doi.org/10.18653/v1/2021.findings-emnlp.323>
- [62] Amarnath Pathak, Partha Pakray, and Jereemi Benthani. 2018. Englishmizo machine translation using neural and statistical approaches. *Neural Computing and Applications* 31 (2018), 7615–7631. DOI : <https://doi.org/10.1007/s00521-018-3601-3>
- [63] Ru Peng, Yawen Zeng, and Jake Zhao. 2022. Distill the image to nowhere: Inversion knowledge distillation for multimodal machine translation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.), Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 2379–2390. DOI : <https://doi.org/10.18653/v1/2022.emnlp-main.152>
- [64] Maja Popović. 2015. chrF: Character n-gram F-score for automatic MT evaluation. In *Proceedings of the 10th Workshop on Statistical Machine Translation*. Ondřej Bojar, Rajan Chatterjee, Christian Federmann, Barry Haddow, Chris Hokamp, Matthias Huck, Varvara Logacheva, and Pavel Pecina (Eds.), Association for Computational Linguistics, Lisbon, Portugal, 392–395. DOI : <https://doi.org/10.18653/v1/W15-3049>
- [65] Maja Popović. 2016. chrF deconstructed: Beta parameters and n-gram weights. In *Proceedings of the 1st Conference on Machine Translation: Volume 2, Shared Task Papers*. Ondřej Bojar, Christian Buck, Rajen Chatterjee, Christian Federmann, Liane Guillou, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Aurélie Nèveól, Mariana Neves, Pavel Pecina, Martin Popel, Philipp Koehn, Christof Monz, Matteo Negri, Matt Post, Lucia Specia, Karin Verspoor, Jörg Tiedemann, and Marco Turchi (Eds.), Association for Computational Linguistics, Berlin, Germany, 499–504. DOI : <https://doi.org/10.18653/v1/W16-2341>
- [66] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models

from natural language supervision. *Proceedings of the 38th International Conference on Machine Learning, Virtual Event, Proceedings of Machine Learning Research* 139 (2021), 8748–8763. Retrieved from <http://proceedings.mlr.press/v139/radford21a.html>

- [67] Gowtham Ramesh, Sumanth Doddapaneni, Aravindh Bheemaraj, Mayank Jobanputra, Raghavan A. K., Ajitesh Sharma, Sujit Sahoo, Harshita Diddee, Mahalakshmi J., Divyanshu Kakwani, Navneet Kumar, Aswin Pradeep, Srihari Nagaraj, Kumar Deepak, Vivek Raghavan, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh Shantadevi Khapra. 2022. Samanantar: The largest publicly available parallel corpora collection for 11 indic languages. *Transactions of the Association for Computational Linguistics* 10 (2022), 145–162. DOI : https://doi.org/10.1162/tacl_a_00452
- [68] Loitongbam Sanayai Meetei, Thoudam Doren Singh, and Sivaji Bandyopadhyay. 2019. WAT2019: English-hindi translation on hindi visual genome dataset. In *Proceedings of the 6th Workshop on Asian Translation*. Toshiaki Nakazawa, Chenchen Ding, Raj Dabre, Anoop Kunchukuttan, Nobushige Doi, Yusuke Oda, Ondřej Bojar, Shantipriya Parida, Isao Goto, and Hidayat Mino (Eds.), Association for Computational Linguistics, Hong Kong, China, 181–188. DOI : <https://doi.org/10.18653/v1/D19-5224>
- [69] Arghyadeep Sen, Shantipriya Parida, Ketan Kotwal, Subhadarshi Panda, Ondřej Bojar, and Satya Ranjan Dash. 2022. Bengali visual genome: A multimodal dataset for machine translation and image captioning. In *Proceedings of the Intelligent Data Engineering and Analytics*. Springer, 63–70.
- [70] Bin Shan, Yaqian Han, Weichong Yin, Shuohuan Wang, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2022. ERNIE-UniX2: A Unified Cross-lingual Cross-modal Framework for Understanding and Generation. arXiv:2211.04861. Retrieved from <https://arxiv.org/abs/2211.04861>
- [71] Xiayang Shi and Zhenqiang Yu. 2022. Adding visual information to improve multimodal machine translation for low-resource language. *Mathematical Problems in Engineering* 2022 (2022), 1–9. Retrieved from <https://ideas.repec.org/a/hin/jnlmpe/5483535.html>
- [72] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *Proceedings of 3rd International Conference on Learning Representations (ICLR 2015)*. Computational and Biological Learning Society, 1–14. Retrieved from <https://ora.ox.ac.uk/objects/uuid:60713f18-a6d1-4d97-8f45-b60ad8aebbce>
- [73] Salam Michael Singh, Loitongbam Sanayai Meetei, Thoudam Doren Singh, and Sivaji Bandyopadhyay. 2021. Multiple captions embellished multilingual multi-modal neural machine translation. In *Proceedings of the 1st Workshop on Multimodal Machine Translation for Low Resource Languages (MMLRL 2021)*. Thoudam Doren Singh, Cristina España i Bonet, Sivaji Bandyopadhyay, and Josef van Genabith (Eds.), INCOMA Ltd., Online (Virtual Mode), 2–11. Retrieved from <https://aclanthology.org/2021.mmlrl-1.2/>
- [74] 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*. Association for Machine Translation in the Americas, Cambridge, Massachusetts, USA, 223–231. Retrieved from <https://aclanthology.org/2006.amta-papers.25/>
- [75] Matthias Sperber, Jan Niehues, and Alex Waibel. 2017. Toward robust neural machine translation for noisy input sequences. In *Proceedings of the 14th International Conference on Spoken Language Translation*. Sakriani Sakti and Masao Utiyama (Eds.), International Workshop on Spoken Language Translation, Tokyo, Japan, 90–96. Retrieved from <https://aclanthology.org/2017.iwslt-1.13/>
- [76] Yoad Tewel, Yoav Shalev, Idan Schwartz, and Lior Wolf. 2022. ZeroCap: Zero-shot image-to-text generation for visual-semantic arithmetic. In *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 17897–17907. DOI : <https://doi.org/10.1109/CVPR52688.2022.01739>
- [77] Zaitinkhuma Thihlum, Vanlalmuansangi Khenglawt, and Somen Debnath. 2020. Machine translation of english language to mizo language. In *Proceedings of the 2020 IEEE International Conference on Cloud Computing in Emerging Markets (CEEM)*. 92–97. DOI : <https://doi.org/10.1109/CEEM50674.2020.00028>
- [78] Jörg Tiedemann and Yves Scherrer. 2017. Neural machine translation with extended context. In *Proceedings of the 3rd Workshop on Discourse in Machine Translation*. Bonnie Webber, Andrei Popescu-Belis, and Jörg Tiedemann (Eds.), Association for Computational Linguistics, Copenhagen, Denmark, 82–92. DOI : <https://doi.org/10.18653/v1/W17-4811>
- [79] 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 *Proceedings of the Advances in Neural Information Processing Systems*. I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30, Curran Associates, Inc. Retrieved from https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- [80] 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 *Proceedings of the Advances in Neural Information Processing Systems*. I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30, Curran Associates, Inc. Retrieved from <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>

- [81] Pidong Wang and Hwee Tou Ng. 2013. A beam-search decoder for normalization of social media text with application to machine translation. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Lucy Vanderwende, Hal Daumé III, and Katrin Kirchhoff (Eds.), Association for Computational Linguistics, Atlanta, Georgia, 471–481. Retrieved from <https://aclanthology.org/N13-1050/>
- [82] Tao Wang, Chengqi Zhao, Mingxuan Wang, Lei Li, Hang Li, and Deyi Xiong. 2021. Secoco: Self-correcting encoding for neural machine translation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.), Association for Computational Linguistics, Punta Cana, Dominican Republic, 4639–4644. DOI : <https://doi.org/10.18653/v1/2021.findings-emnlp.396>
- [83] Ziang Xie, Sida I. Wang, Jiwei Li, Daniel Lévy, Aiming Nie, Dan Jurafsky, and Andrew Y. Ng. 2017. Data noising as smoothing in neural network language models. In *Proceedings of the International Conference on Learning Representations*. Retrieved from <https://openreview.net/forum?id=H1VyHY9gg>
- [84] Shaowei Yao and Xiaojun Wan. 2020. Multimodal transformer for multimodal machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.), Association for Computational Linguistics, Online, 4346–4350. DOI : <https://doi.org/10.18653/v1/2020.acl-main.400>
- [85] Yuting Zhao, Mamoru Komachi, Tomoyuki Kajiwara, and Chenhui Chu. 2022. Region-attentive multimodal neural machine translation. *Neurocomputing* 476 (2022), 1–13. DOI : <https://doi.org/10.1016/j.neucom.2021.12.076>
- [86] Mingyang Zhou, Runxiang Cheng, Yong Jae Lee, and Zhou Yu. 2018. A visual attention grounding neural model for multimodal machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.), Association for Computational Linguistics, Brussels, Belgium, 3643–3653. DOI : <https://doi.org/10.18653/v1/D18-1400>

Received 19 September 2023; revised 21 May 2025; accepted 31 May 2025