
A Case Study on Context Encoding in Multi-Encoder based Document-Level Neural Machine Translation

Ramakrishna Appicharla

appicharla_2021cs01@iitp.ac.in

Baban Gain

gainbaban@gmail.com

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

Santanu Pal

santanu.pal.ju@gmail.com

Wipro AI, Lab45, London, UK

Asif Ekbal

asif@iitp.ac.in

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

Abstract

Recent studies have shown that the multi-encoder models are agnostic to the choice of context, and the context encoder generates noise which helps improve the models in terms of BLEU score. In this paper, we further explore this idea by evaluating with context-aware pronoun translation test set by training multi-encoder models trained on three different context settings *viz.* previous two sentences, random two sentences, and a mix of both as context. Specifically, we evaluate the models on the ContraPro test set to study how different contexts affect pronoun translation accuracy. The results show that the model can perform well on the ContraPro test set even when the context is random. We also analyze the source representations to study whether the context encoder generates noise. Our analysis shows that the context encoder provides sufficient information to learn discourse-level information. Additionally, we observe that mixing the selected context (the previous two sentences in this case) and the random context is generally better than the other settings.

1 Introduction

Document-level neural machine translation (DocNMT) has gained a lot of attention due to the ability to incorporate context through different paradigms such as single encoder (Tiedemann and Scherrer, 2017; Agrawal et al., 2018), multiple encoders (Zhang et al., 2018; Li et al., 2020; Huo et al., 2020), memory networks (Maruf and Haffari, 2018) and pre-trained language models (Donato et al., 2021). This additional context helps to produce more consistent translations (Bawden et al., 2018; Voita et al., 2019) than sentence-level models. Two of the most followed approaches to incorporating context are the concatenation-based and multi-encoder-based approaches. In the concatenation-based method, by concatenating context and current input sentence, a context-aware input sentence is generated (Tiedemann and Scherrer, 2017; Agrawal et al., 2018; Junczys-Dowmunt, 2019; Zhang et al., 2020) and use it as the input to the

encoder. In the multi-encoder approach, to encode the source or target context, an additional encoder is used (Zhang et al., 2018; Voita et al., 2018; Kim et al., 2019; Ma et al., 2020) and the entire model is jointly optimized. Typically, the current sentence’s neighboring sentences (previous or next) are used as the context, whereas models consist of multiple encoders and a single decoder.

Recent studies on Multi-Encoder (MultiEnc) based DocNMT models (Li et al., 2020; Wang et al., 2020; Gain et al., 2022) have shown that the context-encoder is acting as noise generator which improves the robustness of the model and makes the model agnostic to the choice of context. However, the improvement is in terms of BLEU (Papineni et al., 2002), which might not capture the discourse-level phenomenon effectively (Müller et al., 2018). This phenomenon is not studied well in the existing literature. The context encoder might not generate noise if the model can effectively capture any discourse phenomenon, such as pronoun translation from the source to the target language. Modeling the relation between sentences in a given document is essential to capture any discourse phenomenon (Voita et al., 2018). To this end, we hypothesize that, during the training phase, if the model can learn the similarities between all the sentences in a given document given in the form of (*context*, *source*) pairs, the context encoder might not be generating noise since all the sentences in the given document are connected via the context.

In this work, we aim to study the effect of the context in MultiEnc-based DocNMT models and the models’ behavior in random context settings but not to introduce a novel technique. We use the ‘Outside Attention Multi-Encoder’ model (Li et al., 2020) with four different context settings to study the effect of the context. We conduct experiments on News-commentary v14 and TED corpora from English–German direction. We report the results on the ContraPro test set (Müller et al., 2018), a contrastive test set to evaluate models’ performance in translating pronouns. We also report sentence-BLEU (s-BLEU) (Papineni et al., 2002), document-BLEU (d-BLEU) (Liu et al., 2020b; Bao et al., 2021), and COMET (Rei et al., 2020) scores.

To summarize, the specific attributes of our current work are as follows:

- We conduct experiments on multi-encoder based DocNMT models to study if the context encoder is generating noise or not by evaluating the model with ContraPro (Müller et al., 2018) test set.
- We empirically show that the model can learn discourse-level information even when trained with random context.

2 Related Work

The performance of document-level NMT is better than that of sentence-level NMT models due to the encoding of context (Sim Smith, 2017; Voita et al., 2018). Towards this goal to represent context, Tiedemann and Scherrer (2017) concatenate consecutive sentences and use them as input to the single-encoder-based DocNMT model. Agrawal et al. (2018) conducted experiments on varying neighboring contexts and then tied them with the current sentence as input to their model. However, this approach introduced a lot of long-range dependencies. This problem can be alleviated by introducing an additional encoder to encode the context. Towards this, Zhang et al. (2018) and Voita et al. (2018) proposed transformer-based multi-encoder NMT models where the other encoder is used to encode the context. While Miculicich et al. (2018) proposed a hierarchical attention network to encode the context and a more recent approach Kang et al. (2020) proposed a reinforcement learning-based dynamic context selection module for DocNMT. Recent studies (Kim et al., 2019; Li et al., 2020; Wang et al., 2020) have shown that the improvement in the performance of multi-encoder DocNMT models is not due to context encoding but rather the context encoder acting as a noise generator, which improves the

robustness of the DocNMT model. Other approaches such as pre-training (Junczys-Dowmunt, 2019; Donato et al., 2021) and memory-based approaches (Feng et al., 2022) are shown to have improved the performance of DocNMT models. Recent studies (Sun et al., 2022; Post and Junczys-Dowmunt, 2023) have demonstrated that single encoder-decoder models can capture long-range dependencies. Still, special care must be taken to break a large document into smaller fragments for training. Along with document-level translation, multi-encoder models are also commonly used in automatic post-editing (Pal et al., 2019, 2018; Junczys-Dowmunt and Grundkiewicz, 2018; Shin and Lee, 2018), multimodal translation (Libovický et al., 2018; Liu et al., 2020a) and multitask learning (Luong et al., 2015; Zhang et al., 2017; Anastasopoulos and Chiang, 2018) scenarios.

In this work, we study the effect of context in the multi-encoder (Li et al., 2020) based approach. Specifically, we verify if the context encoder is generating noise or not. If the context encoder is generating noise, then the model may not be able to effectively capture the discourse-level phenomenon, such as pronoun translation accuracy, even when the model can perform well on other automatic metrics, such as BLEU.

3 Methodology

3.1 Outside Context Multi-Encoder Model

We conduct all experiments on the ‘Outside Attention Multi-Encoder’ (Li et al., 2020) model. The model (cf. Fig 1) consists of two encoders and one decoder. Both source and context are encoded through two encoders, and the output of these encoders is passed through an attention layer. An element-wise addition is performed on the outputs of the source encoder and the attention layer before passing it to the decoder.




Figure 1: The overview of the Outside Context Multi-Encoder DocNMT architecture. The input to the model consists of $(Context, Source, Target, Label)$. Both the encoders are encoding $Context$ and $Source$. The Context and Source encoder outputs are passed through the Attention layer. Here, ‘ $K-V$ ’ represents Key-Value pairs from the Context encoder, and ‘ Q ’ represents Query from the Source encoder. The output of the Attention layer is element-wise summed with the output of the Source encoder before passing to the Decoder. None of the layers are shared.

3.2 Context-Aware Models

We train context-aware models in four different settings. They are,

1. **MultiEnc-Prev@2**: In this setting, the context consists of the previous two sentences concatenated, with respect to the current source sentence (Zhang et al., 2018), and the

model is trained in this context setting. The same sentence is used as context if the sentence is the first or second sentence in the document.

2. **MultiEnc-Random@2**: In this setting, the context consists of two random sentences sampled from the complete training set.
3. **MultiEnc-Mix@2**: In this setting, 50% of the training set consists of context from ‘MultiEnc-Prev@2’ setting and the remaining 50% consists of context from ‘MultiEnc-Random@2’ setting. Essentially, this setting combines the context settings from the above two approaches.
4. **MultiEnc-Mix-Adapt@2**: This setting is similar to the ‘MultiEnc-Mix@2’ setting but the loss during the training is modified as follows:

$$\mathcal{L} = \alpha \times \mathcal{L} \quad (1)$$

Where ‘ α ’ is a scaling factor which is the fraction of source sentences having the previous two sentences as context over all the sentences in the current batch, and ‘ \mathcal{L} ’ is the loss for the current batch. During the training, we also provide the labels list to facilitate this counting, with 0 indicating random context and 1 indicating the previous two sentences as context. Our motivation in this approach is to penalize the model ¹ based on the number of random context inputs per batch and force the model not to learn from random context.

The validation set consists of the previous two sentences as the context in all four settings.

4 Experimental Setup

This section describes the data sets and the experimental setup used in the experiments.

4.1 Data Statistics

We conduct experiments on English–German corpus obtained from combining ² WMT news-commentary, IWSLT‘17 TED, and Europarl-v7 corpora. For the WMT news-commentary, we use news-commentary v14³ as the train set and newstest2018 as the test set. For IWSLT‘17 TED and Europarl-v7 corpora, we follow the train and test set splits mentioned in the previous work (Maruf et al., 2019)⁴. We use newstest2017 as the validation set for all the models. The models are trained from English to German. Table 1 shows data statistics of the train, validation, and test sets.

4.2 NMT Model Setups

We conduct all the experiments on transformer architecture (Vaswani et al., 2017). All the models are implemented in PyTorch⁵. The models consist of 6-layer encoder-decoder stacks, 8 attention heads, and a 2048-cell feed-forward layer. Positional and token embedding sizes are set to 512. Adam optimizer (Kingma and Ba, 2015) is used for training with a noam learning rate scheduler (Vaswani et al., 2017) and the initial learning rate set to 0.2. The dropout and warmup steps are set to 0.3 (Li et al., 2020; Sun et al., 2022) and 16,000 (Popel and Bojar, 2018) respectively, and we used a mini-batch of 30 sentences. We create joint subword vocabularies of

¹If the entire batch consists of random context, then the loss for that batch will be 0 and vice versa.

²We combine corpora from all three sources into a single corpus and train our models on this corpus.

³<https://data.statmt.org/news-commentary/v14/training/>

⁴<https://github.com/sameenmaruf/selective-attn/tree/master/data>

⁵<https://pytorch.org/>

Data	Corpus	# Sentences	# Documents
Train	News	329,041	8,462
	TED	206,112	1,698
	Europarl	1,666,904	117,855
	Total	2,202,057	128,015
Validation	newstest2017	3,004	130
Test	News	2,998	122
	TED	2,271	23
	Europarl	5,134	360

Table 1: Data statistics of corpora. # **Sentences**, # **Documents** represent the number of sentences and documents, respectively. The train set consists of the corpus obtained by combining News, TED, and Europarl corpora. The models are tested on each test set separately.

size 40,000 by combining source and target parts of the training corpus into a single joint corpus. We use the BPE (Sennrich et al., 2016) to create subword vocabularies with SentencePiece (Kudo and Richardson, 2018) implementation. We also learn the positional encoding of tokens (Devlin et al., 2019), and the maximum sequence length is set to 140 tokens for all models. All models are trained till convergence, and we use perplexity on the validation set as early stopping criteria with the patience of 7 (Popel and Bojar, 2018).

5 Results and Analysis

We test context-aware models with two different contexts *viz.* previous two and random sentences as context. Table 2 shows the s-BLEU (Papineni et al., 2002; Post, 2018)⁶, d-BLEU (Liu et al., 2020b; Bao et al., 2021), and COMET (Rei et al., 2020)⁷ scores. Overall, *MultiEnc-Mix@2* and *MultiEnc-Mix-Adapt@2* models achieve the best overall scores on all test sets. The *MultiEnc-Mix@2* model achieves 23.1 s-BLEU and 25.3 d-BLEU on the News test set in *Prev@2* setting. For Ted test set, *MultiEnc-Mix-Adapt@2* model achieving 20.8 s-BLEU and 24.6 d-BLEU in *Prev@2* setting. For the Europarl test set, *MultiEnc-Mix@2* and *MultiEnc-Mix-Adapt@2* models are achieving 26.5 s-BLEU and the *MultiEnc-Prev@2* model achieving 28.8 d-BLEU on *Random@2* setting. However, the results from *Prev@2* and *Random@2* settings are very similar and not statistically significant when compared to each other.

The COMET scores are also similar in the settings of both *Prev@2* and *Random@2*. On the News test set, *MultiEnc-Mix@2* model achieves a score of 65.8 in *Prev@2* setting. On the Ted test set, *MultiEnc-Mix-Adapt@2* model obtains a score of 71.4 in *Random@2* setting. Similarly, on the Europarl test set, *MultiEnc-Mix@2* model achieves 82.1 in *Prev@2* setting. All the results indicate that the random context might not be random as the performance of the models is similar across both the context setting in terms of BLEU and COMET scores. To verify this, we conduct experiments on the ContraPro test set (Müller et al., 2018) to study the effects of random context and context encoder on pronoun translation accuracy.

5.1 Results on the ContraPro test set

ContraPro test set (Müller et al., 2018) is a test set for contrastive evaluation of models’ performance on translating German pronouns *es*, *er* and *sie*. Contrastive tests test the model’s ability

⁶sacreBLEU signature: nrefs:1|case:mixed|eff:no|tok:13a|smooth:exp|version:2.3.1

⁷COMET model: wmt22-comet-da

Model	s-BLEU			d-BLEU			COMET		
	News	TED	Europarl	News	TED	Europarl	News	TED	Europarl
Prev@2									
MultiEnc-Prev@2	22.9	20.4	26.3	25.2	24.3	28.6	65.3	71.3	81.9
MultiEnc-Random@2	22.7	19.9	26.4	25.0	23.8	28.7	65.4	70.7	81.9
MultiEnc-Mix@2	23.1	20.3	26.4	25.3	24.2	28.7	65.8	71.1	82.1
MultiEnc-Mix-Adapt@2	22.9	20.8	26.4	25.1	24.6	28.7	65.5	71.3	82.0
Random@2									
MultiEnc-Prev@2	22.9	20.5	26.4	25.2	24.4	28.8	65.1	71.3	81.8
MultiEnc-Random@2	22.7	19.9	26.4	25.0	23.8	28.7	65.4	70.7	81.9
MultiEnc-Mix@2	23.0	20.5	26.5	25.2	24.4	28.7	65.6	71.1	82.0
MultiEnc-Mix-Adapt@2	22.7	20.6	26.5	25.0	24.5	28.7	65.7	71.4	82.0

Table 2: s-BLEU, d-BLEU, and COMET scores of the Outside Context DocNMT models, tested with correct and random context. **Prev@2** and **Random@2** denote the previous two and random two-sentence context during the testing. The best scores are shown in bold.

to discriminate between correct and incorrect outputs. In the ContraPro test set, the models’ performance is measured in terms of the model’s accuracy regarding reference pronoun, antecedent location, and antecedent distance. Similar to the training phase, we use the previous two sentences as context for the ContraPro test set, and Table 3 shows the performance of the trained context-aware models. All the models’ performance is similar to reference pronoun translation accuracy, with *MultiEnc-Mix@2* and *MultiEnc-Mix-Adapt@2* achieving a best overall score of 0.47. However, in terms of specific pronouns, *MultiEnc-Random@2*, *MultiEnc-Mix@2*, and *MultiEnc-Mix-Adapt@2* achieved best scores of 0.87, 0.25, and 0.35 for pronouns *es*, *er* and *sie* respectively. This shows that all the models can capture discourse information to translate pronouns, even the model trained with random context (*MultiEnc-Random@2*).

Model	reference pronoun				antecedent location		antecedent distance				
	total	<i>es</i>	<i>er</i>	<i>sie</i>	intra-segmental	external	0	1	2	3	>3
MultiEnc-Prev@2	0.46	0.81	0.21	0.34	0.71	0.39	0.71	0.36	0.45	0.48	0.62
MultiEnc-Random@2	0.46	0.87	0.18	0.33	0.70	0.40	0.70	0.36	0.46	0.49	0.68
MultiEnc-Mix@2	0.47	0.85	0.25	0.31	0.71	0.41	0.71	0.38	0.47	0.48	0.68
MultiEnc-Mix-Adapt@2	0.47	0.83	0.24	0.35	0.73	0.41	0.73	0.38	0.46	0.50	0.64

Table 3: Accuracy on ContraPro test set for Outside Context DocNMT models regarding reference pronoun, antecedent location (within segment vs. outside segment), and antecedent distance of antecedent (in sentences). The best scores are shown in bold.

The results regarding antecedent location show that the *MultiEnc-Mix-Adapt@2* model can perform well when the antecedent occurs in the current segment (intra-segmental). The *MultiEnc-Prev@2* model can perform well when the antecedent is happening within the segment but poorly when the antecedent is outside (external). Similarly, *MultiEnc-Random@2* model is able when the antecedent occurs outside the segment but poorly when the antecedent is within. However, both models *viz.* *MultiEnc-Mix@2* and *MultiEnc-Mix-Adapt@2* perform well in both settings. The results show that mixing some random context is beneficial for the model to learn this discourse phenomenon effectively.

Similarly, the results regarding antecedent distance show that both the *MultiEnc-Mix@2* and *MultiEnc-Mix-Adapt@2* models achieve the best overall performance. Interestingly, the *MultiEnc-Prev@2* model’s performance is good when the antecedent distance is <3 but drops when the distance is >3, but *MultiEnc-Random@2* model is achieving best score when the

distance is >3 but slightly less in other settings than the *MultiEnc-Prev@2* model. The results from *antecedent distance* indicate that the model trained with random context can learn the long distant discourse properties better than the model trained with fixed context (previous two sentences). Based on this, we conclude that the random context might not be random as the model can learn discourse-level information well. These results also indicate that the context encoder might not generate noise as the discourse information outside the current source can only be learned if the context encoder can encode the context well. We further investigate this by analyzing the source sentence embeddings.

5.2 t-SNE Visualization of Source and Target Embeddings

Since the relation between the sentences in a given document is learned through context, the context encoder should be trained sufficiently to capture this aspect. To study this, we take a document and visualize the source sentence representations obtained after performing attention over context and source encoder outputs and combining the resulting output with source encoder output via element-wise addition (cf. Fig 1). Figure 2 shows t-SNE visualization (Van der Maaten and Hinton, 2008)⁸ of the source representations. The document is taken from the train set of *News-commentary v14* corpus and contains 30 sentences. We obtain the source representations in two settings similar to the testing phase *viz.* previous two sentences and randomly sampled two sentences as context.

Figure 2(a) shows the representation when the context consists of the previous two sentences. Interestingly the representations from *MultiEnc-Prev@2* and *MultiEnc-Mix@2* models spread out more than the other two models. The reason might be that the model requires more context to encode the sentences effectively for these two models. The *MultiEnc-Random@2* model can encode the sentences better than the *MultiEnc-Prev@2* model. This might be why the models trained with random context can perform well, as the model can learn sufficient information even from the random context. This shows that the random context might not be random, and it is helping the model to encode sentences well enough to capture the discourse-level information. The *MultiEnc-Mix-Adapt@2* model can learn the source representation better than other models as the sentences are projected in a smaller zone. This shows that the *MultiEnc-Mix-Adapt@2* model can perform well even when the context is limited.

Similarly, Figure 2(c) shows the representation when the context consists of two random sentences. The representation of *MultiEnc-Prev@2* model is adversely affected by the random context, but all other models can learn the representations of sentences well. Interestingly, the spread of *MultiEnc-Mix@2* model is smaller than the spread when the context consists of the previous two sentences. This indicates that the model can learn better representations from random context even though the model is trained by mixing 50% of previous sentence context and 50% of random context. The same is true for the *MultiEnc-Random@2* model also. The representations of the *MultiEnc-Mix-Adapt@2* model are the same as the model trained with the previous sentence context and show that this model is consistent even when the type of context changes.

We also analyze how different types of context affect the decoder. Figure 2(b) and Figure 2(d) shows the target representations when the context consists of the previous two and random two sentences, respectively. We observe that the source context is insufficient to project the target embeddings closer even though the sentences are from the same document. Interestingly, in the random context setting (cf. Figure 2(d)), the representations are projected closer than the correct context setting (cf. Figure 2(b)), indicating that the decoder is mainly unaffected by the choice of the context. This might be due to the context being chosen from the source side and

⁸<https://scikit-learn.org/stable/modules/generated/sklearn.manifold.TSNE.html>

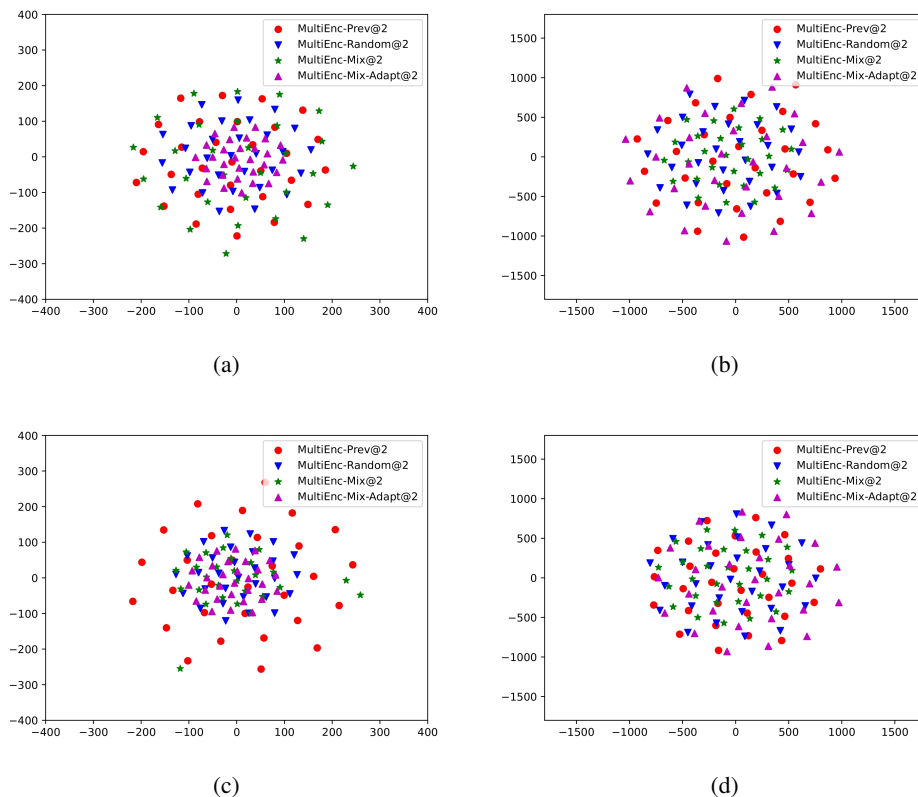


Figure 2: t-SNE visualization of the source and target representations of the Outside Context Multi-Encoder DocNMT models with different context settings. (a) and (b) show when the previous two sentences are used as context, (c) and (d) show when two random sentences are used as context. (a) and (c) show the source representations, (b) and (d) show the target representations. Each point represents a sentence.

leading to the sub-optimal encoding of target sentences. We hypothesize that the model can learn better target representation if trained with target-side context.

Based on the analysis, the model can learn discourse-level properties even when the context is random, indicating that the random context might not adversely affect the context-aware model. This also shows that the context encoder might not be generating noise; instead, it can generate sufficient information to capture discourse-level properties based on the type of context the model is trained with. The t-SNE visualization shows that the source representations are affected by the choice of context, and the target representations are mainly unaffected. This suggests that metrics such as BLEU might not be enough to measure the discourse-level information the system can learn and requires unique discourse-level test sets to evaluate.

5.3 Results of Multi-Encoder models with identical Source and Context

We further study whether the context encoder generates noise by feeding the same source sentence as the context. If the context encoder is generating noise, then the models' performance should be similar, as the inputs are identical to every model. Table 4 shows the s-BLEU scores of the models tested in this setting. Interestingly *MultiEnc-Random@2* models' performance

Model	News	TED	Europarl
MultiEnc-Prev@2	16.9	16.1	22.8
MultiEnc-Random@2	14.1	15.7	20.1
MultiEnc-Mix@2	20.5	19.0	24.4
MultiEnc-Mix-Adapt@2	23.0	20.5	26.5

Table 4: s-BLEU scores of the Outside Context DocNMT models, tested with the same source sentences as the context. The best scores are shown in bold.

is lowest than the other models. This indicates that the context encoder might generate noise when the context is random. Similarly, *MultiEnc-Mix@2* and *MultiEnc-Mix-Adapt@2* models’ performance is better than the other models. Results indicate that mixing the random context with the correct context makes the model robust and results in better-quality translations.

6 Conclusion and Future Work

In this work, we conducted experiments on multi-encoder-based DocNMT systems to study how different types of contexts affect context-aware pronoun translation. Specifically, we consider three different types of context settings *viz.*, previous two sentences, random two sentences, and a mix of both these settings. We use the ContraPro test set as the context-aware test set to analyze the pronoun translation accuracy. Our analysis shows that the multi-encoder models can perform well on pronoun translation even when the context is random. We further conduct experiments to study whether the context encoder is generating noise or not by projecting the sentence representations from a single document using t-SNE. The analysis shows that the context encoder can encode the context sufficiently enough to capture the relation between the sentences, as these sentences are connected only via context. Based on the analysis, we conclude that the random context might not adversely affect the performance of multi-encoder-based DocNMT models. Choosing context is essential for effectively capturing any discourse phenomenon. The context encoder might not be generating noise. Instead, the encoding from the context encoder is dependent on the choice of context. As we observed that mixing selected context (previous two sentences in this case) and random context is performing better than the other settings, we plan to explore effective context encoding through contrastive learning (Hwang et al., 2021) and dynamic context generation based on the source and target pairs which can help during the inference in round-trip-translation (Tu et al., 2017) method.

Acknowledgements

We gratefully acknowledge the support from “NLTM: VIDYAAPATI” project, sponsored by Electronics and IT, Ministry of Electronics and Information Technology (MeitY), Government of India. We also thank the anonymous reviewers for their insightful comments.

References

- Agrawal, R. R., Turchi, M., and Negri, M. (2018). Contextual handling in neural machine translation: Look behind, ahead and on both sides. In *21st Annual Conference of the European Association for Machine Translation*, pages 11–20.
- Anastasopoulos, A. and Chiang, D. (2018). Tied multitask learning for neural speech translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 82–91, New Orleans, Louisiana. Association for Computational Linguistics.

- Bao, G., Zhang, Y., Teng, Z., Chen, B., and Luo, W. (2021). G-transformer for document-level machine translation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3442–3455. Online. Association for Computational Linguistics.
- Bawden, R., Sennrich, R., Birch, A., and Haddow, B. (2018). Evaluating discourse phenomena in neural machine translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1304–1313, New Orleans, Louisiana. Association for Computational Linguistics.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Donato, D., Yu, L., and Dyer, C. (2021). Diverse pretrained context encodings improve document translation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1299–1311. Online. Association for Computational Linguistics.
- Feng, Y., Li, F., Song, Z., Zheng, B., and Koehn, P. (2022). Learn to remember: Transformer with recurrent memory for document-level machine translation. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1409–1420, Seattle, United States. Association for Computational Linguistics.
- Gain, B., Appicharla, R., Chennabasavaraj, S., Garera, N., Ekbal, A., and Chelliah, M. (2022). Investigating effectiveness of multi-encoder for conversational neural machine translation. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 949–954, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Huo, J., Herold, C., Gao, Y., Dahlmann, L., Khadivi, S., and Ney, H. (2020). Diving deep into context-aware neural machine translation. In *Proceedings of the Fifth Conference on Machine Translation*, pages 604–616. Online. Association for Computational Linguistics.
- Hwang, Y., Yun, H., and Jung, K. (2021). Contrastive learning for context-aware neural machine translation using coreference information. In *Proceedings of the Sixth Conference on Machine Translation*, pages 1135–1144. Online. Association for Computational Linguistics.
- Junczys-Dowmunt, M. (2019). Microsoft translator at WMT 2019: Towards large-scale document-level neural machine translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 225–233, Florence, Italy. Association for Computational Linguistics.
- Junczys-Dowmunt, M. and Grundkiewicz, R. (2018). MS-UEdin submission to the WMT2018 APE shared task: Dual-source transformer for automatic post-editing. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 822–826, Belgium, Brussels. Association for Computational Linguistics.
- Kang, X., Zhao, Y., Zhang, J., and Zong, C. (2020). Dynamic context selection for document-level neural machine translation via reinforcement learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2242–2254. Online. Association for Computational Linguistics.

- Kim, Y., Tran, D. T., and Ney, H. (2019). When and why is document-level context useful in neural machine translation? In *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*, pages 24–34, Hong Kong, China. Association for Computational Linguistics.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kudo, T. and Richardson, J. (2018). SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Li, B., Liu, H., Wang, Z., Jiang, Y., Xiao, T., Zhu, J., Liu, T., and Li, C. (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*, pages 3512–3518, Online. Association for Computational Linguistics.
- Librovický, J., Helcl, J., and Mareček, D. (2018). Input combination strategies for multi-source transformer decoder. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 253–260, Brussels, Belgium. Association for Computational Linguistics.
- Liu, J., Luo, L., Ao, X., Song, Y., Xu, H., and Ye, J. (2020a). Meet changes with constancy: Learning invariance in multi-source translation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1122–1132, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., Lewis, M., and Zettlemoyer, L. (2020b). Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Luong, M.-T., Le, Q. V., Sutskever, I., Vinyals, O., and Kaiser, L. (2015). Multi-task sequence to sequence learning. *arXiv preprint arXiv:1511.06114*.
- Ma, S., Zhang, D., and Zhou, M. (2020). A simple and effective unified encoder for document-level machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3505–3511, Online. Association for Computational Linguistics.
- Maruf, S. and Haffari, G. (2018). Document context neural machine translation with memory networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1275–1284, Melbourne, Australia. Association for Computational Linguistics.
- Maruf, S., Martins, A. F. T., and Haffari, G. (2019). Selective attention for context-aware neural machine translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3092–3102, Minneapolis, Minnesota. Association for Computational Linguistics.
- Miculicich, L., Ram, D., Pappas, N., and Henderson, J. (2018). Document-level neural machine translation with hierarchical attention networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2947–2954, Brussels, Belgium. Association for Computational Linguistics.

- Müller, M., Rios, A., Voita, E., and Sennrich, R. (2018). A large-scale test set for the evaluation of context-aware pronoun translation in neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 61–72, Brussels, Belgium. Association for Computational Linguistics.
- Pal, S., Herbig, N., Krüger, A., and van Genabith, J. (2018). A transformer-based multi-source automatic post-editing system. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 827–835, Belgium, Brussels. Association for Computational Linguistics.
- Pal, S., Xu, H., Herbig, N., Krüger, A., and van Genabith, J. (2019). USAAR-DFKI – the transference architecture for English–German automatic post-editing. In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 124–131, Florence, Italy. Association for Computational Linguistics.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Popel, M. and Bojar, O. (2018). Training tips for the transformer model. *arXiv preprint arXiv:1804.00247*.
- Post, M. (2018). A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Post, M. and Junczys-Dowmunt, M. (2023). Escaping the sentence-level paradigm in machine translation.
- Rei, R., Stewart, C., Farinha, A. C., and Lavie, A. (2020). COMET: A neural framework for MT evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Sennrich, R., Haddow, B., and Birch, A. (2016). Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Shin, J. and Lee, J.-H. (2018). Multi-encoder transformer network for automatic post-editing. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 840–845, Belgium, Brussels. Association for Computational Linguistics.
- Sim Smith, K. (2017). On integrating discourse in machine translation. In *Proceedings of the Third Workshop on Discourse in Machine Translation*, pages 110–121, Copenhagen, Denmark. Association for Computational Linguistics.
- Sun, Z., Wang, M., Zhou, H., Zhao, C., Huang, S., Chen, J., and Li, L. (2022). Rethinking document-level neural machine translation. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3537–3548, Dublin, Ireland. Association for Computational Linguistics.
- Tiedemann, J. and Scherrer, Y. (2017). Neural machine translation with extended context. In *Proceedings of the Third Workshop on Discourse in Machine Translation*, pages 82–92, Copenhagen, Denmark. Association for Computational Linguistics.
- Tu, Z., Liu, Y., Shang, L., Liu, X., and Li, H. (2017). Neural machine translation with reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

- Van der Maaten, L. and Hinton, G. (2008). Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Voita, E., Sennrich, R., and Titov, I. (2019). When a good translation is wrong in context: Context-aware machine translation improves on deixis, ellipsis, and lexical cohesion. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1198–1212, Florence, Italy. Association for Computational Linguistics.
- Voita, E., Serdyukov, P., Sennrich, R., and Titov, I. (2018). Context-aware neural machine translation learns anaphora resolution. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics.
- Wang, L., Tu, Z., Wang, X., Ding, L., Ding, L., and Shi, S. (2020). Tencent AI lab machine translation systems for WMT20 chat translation task. In *Proceedings of the Fifth Conference on Machine Translation*, pages 483–491, Online. Association for Computational Linguistics.
- Zhang, J., Liu, Q., and Zhou, J. (2017). Me-md: An effective framework for neural machine translation with multiple encoders and decoders. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 3392–3398.
- Zhang, J., Luan, H., Sun, M., Zhai, F., Xu, J., Zhang, M., and Liu, Y. (2018). Improving the transformer translation model with document-level context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 533–542, Brussels, Belgium. Association for Computational Linguistics.
- Zhang, P., Chen, B., Ge, N., and Fan, K. (2020). Long-short term masking transformer: A simple but effective baseline for document-level neural machine translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1081–1087, Online. Association for Computational Linguistics.